Upgrade from FAS2820

Upgrade controllers

NetApp
February 10, 2026

Table of Contents

Upgrade from FAS2820

Upgrade from FAS2820 by converting to DS212C drive shelf

Perform a nondisruptive upgrade of a NetApp FAS2820 system by converting each node to a DS212C drive shelf. Then, connect the drive shelves to the replacement nodes. This moves the FAS2820 onboard storage to the replacement system.

Before you begin

Review the general upgrade scenarios and upgrade considerations for upgrading by moving volumes or storage:

- Decide whether to upgrade by moving volumes or storage
- Considerations for upgrading controller hardware

About this task

The FAS2820 high-availability (HA) pair controllers are node1 and node2 and the replacement HA pair controllers are node3 and node4.

1. **Migrate LIFs and data aggregates on node2 to node1**

Before converting FAS2820 node2 to a drive shelf, migrate the logical interfaces (LIFs) and data aggregates on node2 to node1.

2. **Convert node2 to a drive shelf and connect to node4**

Convert FAS2820 node2 to a DS212C drive shelf and then connect to node4 before reassigning drives from node2 to node4.

3. **Reassign drives from node2 to node4**

After converting FAS2820 node2 to a DS212C drive shelf and connecting to node4, reassign the drives that previously belonged to node2 to node4.

4. **Migrate data aggregates, epsilon, and LIFs on node1 to node4**

Before converting FAS2820 node1 to a drive shelf, migrate the data aggregates, epsilon, and LIFs on node1 to node4.

5. **Convert node1 to a drive shelf and connect to node3**

Convert FAS2820 node1 to a DS212C drive shelf and connect to node3 before reassigning drives from node1 to node3.

6. **Reassign drives from node1 to node3**

After converting FAS2820 node1 to a DS212C drive shelf and connecting to node3, reassign the drives that previously belonged to node1 to node3.

7 Migrate LIFs and data aggregates on node4 to node3

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

Migrate LIFs and data aggregates on FAS2820 node2 to node1

Migrate the logical interfaces (LIFs) and data aggregates on FAS2820 node2 to node1 before converting node2 to a drive shelf.

Before you begin

Verify that you meet the following requirements:

- The FAS2820 and replacement controllers have the same ONTAP release and patch version, when possible. Refer to the NetApp Hardware Universe for supported ONTAP releases.

  - You need to netboot and install the ONTAP version on node3 and node4 to match the ONTAP version on the FAS2820 systems. Node3 and node4 are the replacement controllers.
  - Both the primary and backup boot image of the node3 and node4 controllers must have the same ONTAP version.
  - You need to clear any residual cluster configuration on node3 and node4 by performing a `wipeconfig` from the boot menu.

- Both replacement controllers are on standby at the LOADER prompt.
- All required cabling is available.

About this task

The following steps are performed on FAS2820 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 their administrative status 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 FAS2820 node2 to a drive shelf and connect to node4

Convert FAS2820 node2 to a DS212C drive shelf. Connect to node4 before reassigning drives from node2 to node4. Node1 and node2 are the controllers inside the DS212C shelf.

Steps

1. Disconnect all the network cables from node2.

2. Remove node2 from the FAS2820 chassis.
3. Insert the IOM12 or IOM12B module into the bay of node2.
4. Cable the node4 serial-attached SCSI (SAS) port to an available port on the IOM12 or IOM12B module. See the Hardware Universe to verify the SAS ports for your system.
5. Create temporary cluster connections by connecting node1 ports e0a and e0b to any two 25GbE ports on node4.

> If node4 only supports 10GbE cluster connections, you must have 10GbE cabling to create temporary cluster connections.

What's next?

Reassign drives from node2 to node4

Reassign drives from FAS2820 node2 to node4

Reassign the drives that previously belonged to FAS2820 node2 to node4.

Before you begin

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

About this task

Perform the following steps on node4.

Steps

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

```
boot_ontap maint
```

   The Maintenance Mode prompt appears.

2. Display all attached drives:

```
disk show -v
```

3. 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.
4. Reassign all drives from node2 to node4.

If you're using whole disks

Run the following command:

```
disk reassign -s <node2_system_ID> -d <node4_system_ID>
```

If you're using partitioned disks

Run the following command:

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

5. 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.

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

```
aggr status
```

7. Exit maintenance mode:

```
halt
```

What's next?

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

Migrate data aggregates, epsilon, and LIFs on FAS2820 node1 to node4

Migrate the data aggregates, epsilon, and logical interfaces (LIFs) on the FAS2820 node1 to node4.

About this task

Perform the following steps on node4.

Steps

1. At the LOADER prompt for node4, set partner-system ID:

```
setenv partner-sysid <system_ID_of_node1>
```

2. Check the partner system-ID:

```
printenv partner-sysid
```

3. Save the changes:

```
saveenv
```

4. Boot the node into the boot menu:

```
boot_ontap menu
```

5. At the boot menu, 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.

6. 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.

7. Migrate the cluster LIFs:

```
set -privilege advanced
```

```
network port show
```

   If the system cluster ports are not similar when upgrading a FAS2820 to a replacement controller, 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>
```

8. Wait for the cluster to reach quorum, then verify that the cluster nodes are healthy:

```
- cluster show
```

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

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

```
network interface modify -vserver Cluster -lif <cluster_LIF> -home-node
<node4> -home-port <port_name>
```

10. Complete this step only if interface groups and data VLANs are used on the FAS2820 cluster you are upgrading. Otherwise, go to Step 11.

    The physical network port names on the replacement controller differ from those on the FAS2820. This might lead to displaced VLANs and incorrectly configured interface groups on node4.

    a. Show the displaced VLANs:

    ```
    cluster controller-replacement network displaced-vlans show
    ```

    b. Restore the displaced VLANs:

    ```
    cluster controller-replacement network displaced-vlans restore
    ```

    c. Fix the incorrectly configured interface groups. The port names between the FAS2820 and the controllers you are upgrading might be different. Update the interface groups with the correct member ports:

    ```
    ifgrp remove-port -node <node2> -ifgrp <ifgrp_name> -port <port_name>
    ```

    ```
    ifgrp add-port -node <node2> -ifgrp <ifgrp_name> -port <port_name>
    ```

11. 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
```

12. Show data aggregate status:

```
storage aggregate show
```

13. Migrate the epsilon by removing it from node1 and 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>
```

14. Show cluster status to verify the epsilon change:

```
cluster show
```

15. Display all data network LIFs:

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

16. Migrate all data LIFs to node4:

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

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

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

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

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

19. Migrate the cluster management LIF:

```
network interface migrate -vserver <vserver_name> -lif
<cluster_mgmt_lif> -destination-node <node4> -destination-port
<port_name>
```

20. Display the status of the cluster management LIF:

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

21. 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 FAS2820 node1 to a drive shelf and connect to node3

Convert FAS2820 node1 to a DS212C drive shelf. Connect to node3 before reassigning drives from node1 to node3..

Steps

1. Disconnect all network cables from node1.
2. Remove node1 from the FAS2820 chassis.
3. Insert the IOM12 or IOM12B module into the bay of node1.
4. Cable the node3 SAS port to an available port on the IOM12 or IOM12B module. See the Hardware Universe to verify the SAS ports for your system. See the Hardware Universe to verify the SAS ports for your system.
5. Create temporary cluster connections by connecting the node4 cluster ports to any cluster ports on node3.

If node3 only supports 10GbE cluster connections, you must have 10GbE cabling to create temporary cluster connections.

What's next?

Reassign drives from node1 to node3

Reassign drives from FAS2820 node1 to node3

Reassign the drives that were previously assigned to FAS2820 node1 to node3.

About this task

Perform the following steps on node3.

Steps

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

   ```
   boot_ontap maint
   ```

   The Maintenance Mode prompt appears.

2. Display all attached drives:

   ```
   disk show -v
   ```

3. Record the local system ID value; this is the system ID of node3. Also record the system IDs of node1 and node4 from the "OWNER" column.
4. Reassign all drives from node1 to node3:

   If you're using whole disks

   Run the following command:

   ```
   disk reassign -s <node1_system_ID> -d <node3_system_ID>
   ```

   If you're using partitioned disks

   Run the following command:

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

5. 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.

6. Exit Maintenance Mode:

```
halt
```

What's next?

Migrate LIFs and data aggregates on node4 to node3

Migrate LIFs and data aggregates on FAS2820 node4 to node3

To complete the upgrade, you connect FAS2820 node3 to node4 and then migrate the data logical interfaces (LIFs) and data aggregates on node4 to node3.

About this task

Perform the following steps on node3.

Steps

1. At the LOADER prompt for node3, 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 node3.

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

3. Enter `y` to continue.
4. 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.

5. Verify that the cluster and HA ports are connected between node3 to node4.
6. Display the cluster and HA ports on node3 and node4:

```
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 planned home ports:

```
network interface migrate -vserver <vserver_name> -lif <lif_name>
-destination-node <node_name> -destination-port <port_name>
```

11. Display all cluster network LIFs:

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

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

```
network interface modify -vserver <vserver_name> -lif <lif_name> -home
-port <port_name>
```

13. Migrate all data LIFs meant for node3 back to node3:

```
network interface migrate -vserver <vserver_name> -lif <lif_name>
-destination-node <node3> -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 their administrative status to `up` by entering the following command once for each LIF:

```
network interface modify -vserver <vserver_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 <vserver_name> -lif
<cluster_mgmt_lif> -destination-node <node3> -destination-port
<port_name>
```

17. Display the status of the cluster management LIF:

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

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

```
storage aggregate show
```

19. Enable cluster high availability in the two-node cluster:

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

20. Enable and verify storage failover for node3 and node4:

```
storage failover modify -node <node3> -enabled true
```

```
storage failover modify -node <node4> -enabled true
```

```
storage failover show
```

21. Migrate data aggregates owned by node4 that should be owned by node3:

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

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

```
storage aggregate show
```

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

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

24. Enable storage failover automatic giveback:

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

25. Display the cluster status:

```
cluster show
```

26. 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, perform a takeover and giveback from both sides of the cluster.

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

```
storage aggregate show
```

Copyright information

Trademark information