Adding a second controller to create an HA pair

Upgrading a stand-alone controller module to a two-node cluster in an HA pair is a multistep process involving both hardware and software changes that must be performed in the proper order.

Before you begin

- A controller module must already be installed, configured, and operating in ONTAP 8.2 or later.
  This controller module is referred to as the *existing* controller module. The examples in this procedure have the console prompt `existing_ctlr>`.
  The controller module that is being added is referred to as the *new* controller module. The examples in this procedure have the console prompt `new_ctlr>`.
  The new controller module and the existing controller module should be of the same model.
- The new controller module must be received from NetApp as part of the upgrade kit.
  This procedure does not apply to moving a controller module from a preexisting system or a system that was previously in an HA pair. However, if you populate the new controller module with PCIe cards from existing inventory at the customer site, you must verify that they are compatible with and supported by the new controller module.
  NetApp Hardware Universe
- Your system must have an empty slot available for the new controller module when upgrading to a single-chassis HA pair (an HA pair in which both controller modules reside in the same chassis).

  Note: This configuration is not supported on all systems.
- You must have rail kits, rack space, and cables for the new controller module when upgrading to a dual-chassis HA pair (an HA pair in which the controller modules reside in separate chassis).

  Note: This configuration is not supported on all systems.
- Each controller module must be connected to the management network through the e0M port (wrench port).

About this task

This procedure can take over an hour, with additional time needed to initialize the disks. The time to initialize the disks depends on the size of the disks.

Steps

Preparing for the upgrade

Before upgrading to an HA pair, you must verify that your system meets all requirements and that you have all of the necessary information.

Steps

1. Verify that your system has enough available disks or partitions for the new controller module.

   You need to identify unassigned disks or spare disks with available partitions that you can assign to the new controller module.

   Physical Storage Management Guide

   ONTAP 9 Disks and Aggregates Power Guide

   - If your system does not use root-data partitioning, check disk ownership: **`storage disk show -ownership`**
     Two parity disks and one data disk, plus one spare, are required for root aggregate creation.

     Note: You must set the `auto_assign` option to **`off`** on the existing controller module before adding any new disks.

   - If your system uses root-data partitioning, determine the system spare disks and available partition space:

     a. Identify the unassigned disks: **`storage disk show`**

```
clust01::> storage disk show

       Usable              Disk    Container   Container
Disk   Size     Shelf Bay  Type    Type        Name      Owner
------ -------- ----- --- ------- ----------- --------- --------
1.0.0         -    0   0    SAS     unassigned  -         -
1.0.1  408.2GB    0   1    SAS     shared      -
1.0.2  408.2GB    0   2    SAS     shared      -         clust01-01
1.0.3  408.2GB    0   3    SAS     shared      -         clust01-01
1.0.4  408.2GB    0   4    SAS     shared      -         clust01-01
1.0.5  408.2GB    0   5    SAS     shared      -         clust01-01
1.0.6  408.2GB    0   6    SAS     shared      -         clust01-01
1.0.7  408.2GB    0   7    SAS     shared      -         clust01-01
.
.
.
.
```

     b. Identify the disks with usable root and data space: **`storage aggregate show-spare-disks`**
     You should look for usable space in the `Local Data Usable` and `Local Root Usable` columns for available partition space.

```
clust01::> storage aggregate show-spare-disks
.
.
.
Original Owner: existing_ctlr
Pool0
  Partitioned Spares

                                          Local   Local
                                          Data    Root     Physical
Disk    Type     RPM    Checksum         Usable  Usable   Size     Status
-----   -----    -----   ---------       ------ --------  ------   -------
2.3.9    BSAS     7200    advanced_zoned  0B      73.89GB  931.5GB  zeroed
3.0.6    BSAS     7200    block           0B      0B       828.0GB  offline
3.5.1    BSAS     7200    block           354.3GB 413.8GB  828.0GB  zeroed
```

```
.
.
.
.
```

- If the system has 12 internal disks, five partitions are required for root aggregate creation and a sixth partition is required as a spare.
- If the system has 24 internal disks, 10 partitions are required for root aggregate creation and two disks are required as spares.
- If the system is an All Flash FAS (AFF) system, 22 partitions are required for the root aggregate creation and two disks are required for spares.

2. Based on the results of the previous step, perform either of the following:

If the result showed...	Then...
Enough spare disks available for the new controller module	Go to the next step.

Not enough spares for the new controller module on a system without root-data partitioning	Complete the following substeps:
	a. Determine where the aggregates for the existing node are located:
	`storage aggregate show`
	Note: If you do not have enough free disks for your system, you need to add more storage. Contact technical support for more information.
	b. If disk ownership automatic assignment is on, turn it off:
	`storage disk option modify -node node_name -autoassign off`
	c. Remove ownership on disks that do not have aggregates on them:
	`storage disk removeowner disk_name`
	d. Repeat the previous step for as many disks as you need for the new node.

Not enough spares for the new controller module on a system with root-data partitioning	Perform either of the following steps: • Add more storage to the system. • Identify disk partitions that are not part of existing aggregates, so that you can use them when assigning disks.

3. Verify that you have cables ready for the following connections:
   - Cluster connections
     If you are creating a two-node switchless cluster, you require two cables to connect the controller modules. Otherwise, you require a minimum of four cables, two for each controller module connection to the cluster-network switch. Other systems (like the 80xx series) have defaults of either four or six cluster connections.
   - HA interconnect connections, if the system is in a dual-chassis HA pair
4. Verify that you have a serial port console available for the controller modules.
5. Verify that your environment meets the site and system requirements.
   NetApp Hardware Universe
6. Gather all of the IP addresses and other network parameters for the new controller module.

Preparing to add a controller module when using Storage Encryption

If the existing controller module is configured for Storage Encryption, you must gather information from the system and rekey the self-encrypting disks (SEDs) before adding the new controller module.

About this task

You must enter the commands in the steps below in the nodeshell. For more information about the nodeshell, see the *System Administration Reference*.

Steps

1. Enter the following command and note the key IDs on all disk drives that are using Storage Encryption:

   `disk encrypt show`

   #### Example

   The command displays the status of each self-encrypting disk:

   ```
   storage-system> disk encrypt show
   Disk       Key ID                                                           Locked?
   0c.00.1    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.0    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.3    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.4    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.2    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   0c.00.5    080CF0C8000000000100000000000000A948EE8604F4598ADFFB185B5BB7FED3  Yes
   ```

2. Enter the following command and note all the necessary certificate files (`client.pem`, `client_private.pem`, and *`ip_address_key_server`*`_CA.pem`) that have been installed:

   `keymgr list cert`

   Later in the procedure you need to install these same certificate files on the new partner controller module.

3. Enter the following command to identify the IP address of the key servers:

   `key_manager show`

   All external key management servers associated with the storage system are listed. Later in the procedure you need to add these same key servers on the new partner controller module.

   #### Example

   The following command displays all external key management servers associated with the storage system:

   ```
   storage-system> key_manager show
     172.18.99.175
   ```

4. Enter the following command and check that the key IDs listed match those shown by the `disk encrypt show` command in step 1:

   `key_manager query`

   #### Example

   The following command checks the status of all key management servers linked to the storage system and displays additional information:

```
storage-system> key_manager query

Key server 172.18.99.175 reports 4 keys.

Key tag                           Key ID
--------                          -------
storage-system                    080CF0C80...
storage-system                    080CF0C80...
storage-system                    080CF0C80...
storage-system                    080CF0C80...
```

5. Back up all data on all aggregates using standard methods for your site.
6. Enter the following command to reset the authentication key on the drives using Storage Encryption to their Manufacturing System ID (MSID):

   `disk encrypt rekey 0x0 *`

7. Examine the CLI command output to ensure that there are no `disk encrypt rekey` failures.

Preparing cluster ports on an existing controller module

Before installing a new controller module, you must configure cluster ports on the existing controller module so that the cluster ports can provide cluster communication with the new controller module.

About this task

If you are creating a two-node switchless cluster (with no cluster network switches), you must enable the switchless cluster networking mode.

For detailed information about port, LIF, and network configuration in ONTAP, see the *Clustered Data ONTAP Network Management Guide*.

ONTAP 9 Network Management Guide

Steps

1. Determine which ports should be used as the node's cluster ports.

   For a list of the default port roles for your platform, see the *Hardware Universe* at *hwu.netapp.com*

   The *Installation and Setup Instructions* for your platform on the NetApp Support Site contains information about the ports for cluster network connections.

2. For each cluster port, identify the port roles: **`network port show`**

Example

In the following example, ports e0a, e0b, e0c, and e0d must be changed to cluster ports:

```
cluster_A::> network port show

Node: controller_A_1
Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link  MTU    Admin/Oper  Status
--------- ------------ ---------------- ----  ----   ----------- --------
e0M       Default      mgmt_bd_1500     up    1500   auto/1000   healthy
e0a       Default      Default          up    1500   auto/10000  healthy
e0b       Default      Default          up    1500   auto/10000  healthy
e0c       Default      Default          up    1500   auto/10000  healthy
e0d       Default      Default          up    1500   auto/10000  healthy
e0i       Default      Default          down  1500   auto/10     -
e0j       Default      Default          down  1500   auto/10     -
e0k       Default      Default          down  1500   auto/10     -
e0l       Default      Default          down  1500   auto/10     -
e2a       Default      Default          up    1500   auto/10000  healthy
```

```
e2b       Default      Default          up    1500   auto/10000  healthy
e4a       Default      Default          up    1500   auto/10000  healthy
e4b       Default      Default          up    1500   auto/10000  healthy
13 entries were displayed.
```

3. Based on the version of ONTAP your system is running, set the port roles to the **cluster**:

If your system is running...	Then...
Data ONTAP 8.2.x or earlier	For each port that you identified as a cluster port, modify the role to cluster, and then set the MTU to the default value of 9000: **network port modify -node local -port *port_name* -role cluster -mtu 9000 -ipspace Cluster**
ONTAP 8.3 or later	If the port roles are not set to **cluster**: **a.** You must change each incorrect port role to the correct role: **network port broadcast-domain remove-ports -ipspace *Default* -broadcast-domain *Default* -ports *node_name:port_name*** **Note:** Your domain name might be different from the name that is shown in the example. **b.** You must add the port to the cluster domain: **network port broadcast-domain add-ports -ipspace *Cluster* -broadcast-domain *Cluster* -ports *node_name:port_name*** **c.** You must modify the MTU of the cluster broadcast domain: **network port broadcast-domain modify -broadcast-domain *Cluster* -ipspace *Cluster* -mtu *Cluster* 9000**

4. Verify that the port roles have changed: **network port show**

Example

The following example shows that ports e0a, e0b, e0c, and e0d are now cluster ports:

```
Node: controller_A_1
Speed(Mbps) Health
Port      IPspace      Broadcast Domain Link  MTU    Admin/Oper  Status
--------- ------------ ---------------- ----  ----   ----------- --------
e0M       Default      mgmt_bd_1500     up    1500   auto/1000   healthy
e0a       Cluster      Cluster          up    9000   auto/10000  healthy
e0b       Cluster      Cluster          up    9000   auto/10000  healthy
e0c       Cluster      Cluster          up    9000   auto/10000  healthy
e0d       Cluster      Cluster          up    9000   auto/10000  healthy
e0i       Default      Default          down  1500   auto/10 -
e0j       Default      Default          down  1500   auto/10 -
e0k       Default      Default          down  1500   auto/10 -
e0l       Default      Default          down  1500   auto/10 -
e2a       Default      Default          up    1500   auto/10000  healthy
e2b       Default      Default          up    1500   auto/10000  healthy
e4a       Default      Default          up    1500   auto/10000  healthy
e4b       Default      Default          up    1500   auto/10000  healthy
13 entries were displayed.
```

5. For each cluster port, change the home port of any of the data LIFs on that port to a data port: **network interface modify**

Example

The following example changes the home port of a data LIF to a data port:

```
cluster1::> network interface modify -lif datalif1 -vserver vs1 -home-port e1b
```

6. For each LIF that you modified, revert the LIF to its new home port: **network interface revert**

Example

The following example reverts the LIF datalif1 to its new home port e1b:

```
cluster1::> network interface revert -lif datalif1 -vserver vs1
```

7. If your system is part of a switched cluster, create cluster LIFs on the cluster ports: **`network interface create`**

Example

The following example creates a cluster LIF on one of the node's cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address.

```
cluster1::> network interface create -vserver Cluster -lif clus1 -role cluster -home-node
node0 -home-port e1a -auto true
```

8. If you are creating a two-node switchless cluster, enable the switchless cluster networking mode:
   a. Change to the advanced privilege level from either node: **`set -privilege advanced`**

      You can respond **`y`** when prompted whether you want to continue into advanced mode. The advanced mode prompt appears (`*>`).

   b. Enable the switchless cluster networking mode: **`network options switchless-cluster modify -enabled true`**

   c. Return to the admin privilege level: **`set -privilege admin`**

Important: Cluster interface creation for the existing node in a two-node switchless cluster system is completed after cluster setup is completed through a netboot on the new controller module.

Preparing the netboot server to download the image

When you are ready to prepare the netboot server, you must download the correct ONTAP netboot image from the NetApp Support Site to the netboot server and note the IP address.

Before you begin

- You must be able to access an HTTP server from the system before and after adding the new controller module.
- You must have access to the NetApp Support Site to download the necessary system files for your platform and your version of ONTAP.
  NetApp Support Site
- Both controller modules in the HA pair must run the same version of ONTAP.

Steps

1. Download and extract the `netboot.tgz` file from the NetApp Support Site.

   The `netboot.tgz` file is used for performing a netboot of your system. You should download the file contents to a web-accessible directory.

   a. Download the `netboot.tgz` file from the NetApp Support Site to a web-accessible directory.

   b. Switch to the web-accessible directory.

   c. Extract the contents of the `netboot.tgz` file to the target directory `tar -zxvf netboot.tgz`.

Your directory listing should contain the following directory:

```
netboot/
```

2. Download the `image.tgz` file from the NetApp Support Site to the web-accessible directory.

   Your directory listing should contain the following file and directory:

```
image.tgz
netboot/
```

3. Determine the IP address of the existing controller module.

   This address is referred to later in this procedure as *`ip-address-of-existing controller`*.

4. Ping *`ip-address-of-existing controller`* to verify that the IP address is reachable.

Setting the mode on the existing controller module

You must use the `storage failover modify` command to set the mode on the existing controller module. The mode value is enabled later after you reboot the controller module.

Step

1. Enter the following command for the existing node, either *`ha`* or *`mcc`*:

   `storage failover modify -mode` ***`ha_state`*** **`-node`** ***`existing_node_name`***

Shutting down the existing controller module

You must perform a clean shutdown of the existing controller module to verify that all of the data has been written to disk. You must also disconnect the power supplies.

Steps

1. Halt the node from the existing controller module prompt: **`halt local -inhibit-takeover true`**

   If you are prompted to continue the halt procedure, enter **`y`** when prompted, and then wait until the system stops at the LOADER prompt.

   Attention: You must perform a clean system shutdown before replacing the system components to avoid losing unwritten data in the NVRAM or NVMEM.

   - In a 32xx system, the NVMEM LED is located on the controller module to the right of the network ports, marked with a battery symbol.
   - In a 62xx system, the NVRAM LED is located on the controller module to the right of the network ports, marked with a battery symbol.
   - In an 80xx system, the NVRAM LED is located on the controller module to the right of the network ports, marked with a battery symbol.

   This LED blinks if there is unwritten data in the NVRAM. If this LED is flashing amber after you enter the `halt` command, you need to reboot your system and try halting it again.

2. If you are not already grounded, properly ground yourself.

3. Turn off the power supplies and disconnect the power, using the correct method for your system and power-supply type:

If your system uses...	Then...
AC power supplies	Unplug the power cords from the power source, and then remove the power cords.
DC power supplies	Remove the power at the DC source, and then remove the DC wires, if necessary.

Installing and cabling the new controller module

You must physically install the new controller module in the chassis, and then cable it.

Steps

1. If you have an I/O expansion module (IOXM) in your system and are creating a single-chassis HA pair, you must uncable and remove the IOXM.

   You can then use the empty bay for the new controller module. However, the new configuration will not have the extra I/O provided by the IOXM.

2. Physically install the new controller module and, if necessary, install additional fans:

If you are adding a controller module...	Then perform these steps...
To an empty bay to create a single-chassis HA pair and the system belongs to one of the following platforms: • 6210 • 6220	**a.** Install three additional fans in the chassis to cool the new controller module: **i.** Remove the bezel by using both hands to hold it by the openings on each side, and then pull the bezel away from the chassis until it releases from the four ball studs on the chassis frame. **ii.** Remove the blank plate that covers the bay that will contain the new fans. **iii.** Install the fans as described in the *Replacing a fan module* document for your system on the NetApp Support Site at *mysupport.netapp.com*. **b.** Remove the blank plate in the rear of the chassis that covers the empty bay that will contain the new controller module. **c.** Gently push the controller module halfway into the chassis.
To an empty bay to create a single-chassis HA pair and the system belongs to one of the following platforms: • FAS22xx • FAS25xx • FAS2600 series • 32xx • 8020 • 8040 • 8060	**a.** Remove the blank plate in the rear of the chassis that covers the empty bay that will contain the new controller module. **b.** Gently push the controller module halfway into the chassis. To prevent the controller module from automatically booting, do not fully seat it in the chassis until later in this procedure.

If you are adding a controller module...	Then perform these steps...
In a separate chassis from its HA partner to create a dual-chassis HA pair when the existing configuration is in a controller-IOX module configuration. • FAS32xx • FAS62xx • FAS8080	Install the new system in the rack or system cabinet.

3. If you have a dual-chassis HA pair, cable the HA interconnect.
4. Cable the cluster network connections, as necessary:
   a. Identify the ports on the controller module for the cluster connections.

      Note: AFF systems are not supported with array LUNs.

      Installation and Setup Instructions for AFF A300 Systems

      Installation and Setup Instructions for AFF A700 and FAS9000

      Installation and Setup Instructions for FAS8200 Systems

      Installation and Setup Instructions FAS8040/FAS8060 Systems

      Installation and setup Instructions FAS80xx Systems with I/O Expansion Modules

      Installation and Setup Instructions FAS8020 systems

      Installation and Setup Instructions 62xx Systems

      Installation and Setup Instructions 32xx Systems

   b. If you are configuring a switched cluster, identify the ports that you will use on the cluster network switches.

      See the *Clustered Data ONTAP Switch Setup Guide for Cisco Switches*, *NetApp 10G Cluster-Mode Switch Installation Guide* or *NetApp 1G Cluster-Mode Switch Installation Guide*, depending on what switches you are using.

   c. Connect cables to the cluster ports:

If the cluster is...	Then...
A two-node switchless cluster	Directly connect the cluster ports on the existing controller module to the corresponding cluster ports on the new controller module.
A switched cluster	Connect the cluster ports on each controller to the ports on the cluster network switches identified in substep b.

5. Power up the existing controller module.
6. Depending on your configuration, power up the new controller module and interrupt the boot process:

If the new controller module is...	Then...
In the same chassis as the existing controller module	**a.** Push the controller module firmly into the bay. When fully seated, the controller module receives power and automatically boots. **b.** Interrupt the boot process by pressing Ctrl-C. **c.** Tighten the thumbscrew on the cam handle. **d.** Install the cable management device. **e.** Bind the cables to the cable management device with the hook and loop strap.
In a separate chassis from the existing controller module	**a.** Turn on the power supplies on the new controller module. **b.** Interrupt the boot process by pressing Ctrl-C.

The system displays the LOADER prompt (`LOADER>`, `LOADER-A>`, or `LOADER-B>`).

Note: If there is no LOADER prompt, record the error message and contact technical support. If the system displays the boot menu, reboot and attempt to interrupt the boot process again.

Configuring and cabling CNA ports (80xx, FAS2600 series, FAS2552, and FAS2554 systems only)

If you are adding a controller module to an 80xx, FAS2600 series, FAS2552, or FAS2554 system, you must check the configuration of the CNA ports on the new controller module and, if necessary, change the default port configuration to match the CNA port configuration of the existing controller module.

Before you begin

You must have the SFP+ modules for the CNA ports. If the ports are set to a 10 GbE personality, you can use twinax cables.

Steps

1. Boot to Maintenance mode on the new node, if it is not in Maintenance mode, by entering **`halt`** to go to the LOADER prompt.

   If you are running ONTAP 8.2.1 or later, enter **`boot_ontap maint`** at the LOADER prompt and enter **`y`** to continue when prompted.

2. On the existing controller module console, check how the ports are currently configured: **`system node hardware unified-connect show`**

   ### Example

   The system displays output similar to the following example:

   ```
   node_name::> system node hardware unified-connect show
                Current Current   Pending  Pending
   Node  Adapter  Mode    Type      Mode     Type      Status
   ----  -------  ------- --------- -------  -------   ------
   f-a    0e      fc      initiator -        -         online
   f-a    0f      fc      initiator -        -         online
   f-a    0g      cna     target    -        -         online
   ...
   ```

3. On the console of the new node, display the port settings: **`ucadmin show`**

Example

The system displays output similar to the following example:

```
*> ucadmin show
                Current Current   Pending  Pending
Node  Adapter  Mode    Type      Mode     Type      Status
----  -------  ------- --------- -------  -------   ------
f-a    0e      fc      initiator -        -         online
f-a    0f      fc      initiator -        -         online
f-a    0g      cna     target    -        -         online
...
```

4. If the current SFP+ module does not match the desired use, replace it with the correct SFP+ module.
5. If the current configuration does not match the existing node's configuration, change the configuration as required:

If the desired use is for...	Then enter the following command...
FC initiator	**ucadmin modify -t initiator *adapter_name***
FC target	**ucadmin modify -t target *adapter_name***
Ethernet	**ucadmin modify -m cna *adapter_name***

Note: If you changed the port configuration, it will take effect when the new node is booted. To confirm the configuration change, you must verify the settings after the boot.

6. Cable the port.

Verifying and setting the HA state of the controller module and chassis

You must verify the **HA** state of the chassis and controller modules, and, if necessary, update the state to indicate that the system is in an HA pair or a MetroCluster configuration. If your system is in a MetroCluster configuration, you must have ONTAP 8.3 or later installed. If you have a FAS20xx, 31xx, or 60xx system, you can skip this task.

Steps

1. If you are not already in Maintenance mode, boot to Maintenance mode by entering **halt** to go to the LOADER prompt:

   If you are running ONTAP 8.2.1 or later, enter **boot_ontap maint** at the LOADER prompt and enter **y** to continue when prompted.

2. In Maintenance mode, display the **HA** state of the new controller module and chassis from the existing controller module: **ha-config show**

   The **HA** state should be the same for all of the components.

If your system is...	The HA state for all components should be...
In an HA pair	ha
In a MetroCluster configuration	mcc
Stand-alone	non-ha

3. If the displayed system state of the controller does not match your system configuration, set the **HA** state for the controller module: **ha-config modify controller [ha | non-ha]**

If your system is...	Then enter this command...
In an HA pair	**ha-config modify controller ha**
In a MetroCluster configuration	**ha-config modify controller mcc**
Stand-alone	**ha-config modify controller non-ha**

4. If the displayed system state of the chassis does not match your system configuration, set the **HA** state for the chassis: **ha-config modify chassis [ha | non-ha]**

If your system is...	Then enter this command...
In an HA pair	**ha-config modify chassis ha**
In a MetroCluster configuration	**ha-config modify chassis mcc**
Stand-alone	**ha-config modify chassis non-ha**

5. Retrieve the system ID for the current node: **sysconfig**

   You should make a note of the system ID. You require the system ID when you assign disks to the new node.

6. Exit Maintenance mode: **halt**

7. If you are using root-data partitioning, set the partner system ID for each controller module:

   a. On the existing controller module, set the partner system ID to that of the new controller module: **setenv partner-sysid *sysID_of_new_controller***

   b. On the new controller module, set the partner system ID to that of the existing controller module: **setenv partner-sysid *sysID_of_existing_controller***

Assigning disks to the new controller using root-data partitioning

For systems using root-data partitioning, you must assign root partitions and data partitions to the new controller module before you complete the initial configuration of the new controller module through netboot.

Before you begin

You must have made sure that there are enough spares, unassigned disks, or assigned disks that are not part of an existing aggregate available that were identified in *Preparing for the upgrade* on page 2.

About this task

These steps are performed on the existing controller module.

Steps

1. Enter advanced mode on the existing controller module:

   set -privilege advanced

   Enter **y** when you are prompted.

2. Assign a root partition belonging to the container disk assigned in the previous step to the new controller module:

   storage disk assign -disk *disk_name* -root true -sysid *new_controller_sysID* -force true

   The system ID of the new controller module was identified in *Verifying and setting the HA state of the controller module and chassis* on page 12.

Example

For example, the following command assigns the root partition of disk 2.3.9 to the new controller module:

```
storage disk assign -disk 2.3.9 -root true -sysid 1873758094 -force true
```

In the following example, the system ID of the new controller module is 1873758094.

3. Assign the same container disk from Step 2 to the new controller module:

   ```
   storage disk assign -disk container_disk_name -sysid new_controller_sysID -force true
   ```

Example

For example, the following command assigns container disk 2.3.9 to the new controller module:

```
storage disk assign -disk 2.3.9 -sysid 1873758094 -force true
```

4. Assign a spare data partition to the new controller module:

   ```
   storage disk assign -disk disk_name -data true -sysid new_controller_sysID -force true
   ```

   The system ID of the new controller module was identified in *Verifying and setting the HA state of the controller module and chassis* on page 12.

   Note: Available spare disks and partitions were identified in *Preparing for the upgrade* on page 2.

Example

For example, the following command assigns the data partition of disk 3.5.1 to the new controller module:

```
storage disk assign -disk 3.5.1 -data true -sysid 1873758094 -force true
```

In the following example, the system ID of the new controller module is 1873758094.

5. Repeat the preceding steps until all required partitions have been assigned to the new controller module.
6. Verify that the disk assignments are correct by examining the output from the **storage disk show -partition-ownership** command and correcting as needed.
7. Exit advanced mode:

   ```
   set -privilege admin
   ```

Netbooting and setting up Data ONTAP on the new controller module

You must perform a specific sequence of steps to boot and install the operating system on the new controller module.

About this task

This procedure includes initializing disks. The amount of time you need to initialize the disks depends on the size of the disks.

If your system does not use disk partitioning, the system automatically assigns two disks to the new controller module.

Physical Storage Management Guide

ONTAP 9 Disks and Aggregates Power Guide

Steps

1. If you are running Data ONTAP 8.2.x or earlier, set the following boot environment variable at the LOADER prompt (LOADER>, LOADER-A>, or LOADER-B>) on the target node console: **setenv bootarg.init.boot_clustered true**

   Note: This step applies only to systems running Data ONTAP.

2. Configure the IP address of the new controller module at the LOADER prompt:

If DHCP is...	Then enter...
Available	**ifconfig e0M -auto**
Not available	**ifconfig e0M -addr=*filer_addr* -mask=*netmask* -gw=*gateway* -dns=*dns_addr* -domain=*dns_domain*** *filer_addr* is the IP address of the storage system. *netmask* is the network mask of the storage system. *gateway* is the gateway for the storage system. *dns_addr* is the IP address of a name server on your network. *dns_domain* is the DNS (Domain Name System) domain name. If you use this optional parameter, you do not need a fully qualified domain name in the netboot server URL; you need only the server's host name. **Note:** Other parameters might be necessary for your interface. For details, use the `help ifconfig` command at the LOADER prompt.

3. Perform netboot at the LOADER prompt: **netboot http://*path_to_web-accessible_directory*/netboot/kernel**
4. Select the **Install new software first** option from the displayed menu.
5. Enter **y** when prompted regarding the installation of the alternate software, from which the node is not running currently.
6. Enter the path to the `image.tgz` file when prompted.

Example

```
What is the URL for the package?
http://path_to_web-accessible_directory/image.tgz
```

7. Enter **n** to skip the backup recovery when prompted:

Example

```
****************************************************************
*             Restore Backup Configuration                     *
*  This procedure only applies to storage controllers that     *
*  are configured as an HA pair.                               *
*                                                              *
*  Choose Yes to restore the "varfs" backup configuration      *
*  from the SSH server.  Refer to the Boot Device Replacement  *
*  guide for more details.                                     *
*  Choose No to skip the backup recovery and return to the     *
*  boot menu.                                                  *
****************************************************************

Do you want to restore the backup configuration
now? {y|n} n
```

8. Reboot by entering **y** when prompted to do so.

```
The node must be rebooted to start using the newly installed software. Do you want to
reboot now? {y|n}y
```

9. If necessary, select the option to **Clean configuration and initialize all disks** after the node has booted.

Because you are configuring a new controller module and the disks for the new controller module are empty, you can respond **y** when the system warns you that selecting this option erases data in all of the disks.

Note: The amount of time needed to initialize the disks depends on the size of your disks and configuration.

10. Respond to the applicable wizard:

If you are running ...	Then...
Data ONTAP 8.2.x or earlier	The Cluster Setup wizard starts after the disks are initialized: **a.** Enter the node management information, as prompted by the wizard. **b.** Go to step *11* on page 17.

If you are running ...	Then...
ONTAP 8.3 or later	The Node Setup wizard starts after the disks are initialized: **a.** Enter the node management LIF information on the console, as shown in the following example: **b.** Manually enter the admin login ID when prompted to do so. **c.** Manually start the Cluster Setup wizard at the prompt: **`cluster setup`**

```
Welcome to node setup.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question
clarified,
  "back" - if you want to change previously answered
questions, and
  "exit" or "quit" - if you want to quit the cluster
setup wizard.
     Any changes you made before quitting will be saved.

To accept a default or omit a question, do not enter a
value.

This system will send event messages and weekly reports
to NetApp
Technical Support. To disable this feature, enter
"autosupport
modify -support disable" within 24 hours. Enabling
AutoSupport can
significantly speed problem determination and resolution
should
a problem occur on your system.
For further information on AutoSupport, please see:
http://support.netapp.com/autosupport/


Type yes to confirm and continue{yes}: yes

Enter the node management interface port [e0M]:
Enter the node management interface IP address:
10.98.230.86
Enter the node management interface netmask: 255.255.240.0
Enter the node management interface default gateway:
10.98.224.1
A node management interface on port e0c with IP address
10.98.230.86 has been created.

This node has its management address assigned and is
ready for cluster setup.

To complete cluster setup after all nodes are ready,
download and run the System Setup utility from the
NetApp Support Site and use it to discover the configured
nodes.

For System Setup, this node's management address is:
10.98.230.86
```

11. With the **Cluster Setup** wizard running, join the node to the cluster: **`join`**

Example

```
Welcome to the cluster setup wizard.

You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
     Any changes you made before quitting will be saved.

You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.


Do you want to create a new cluster or join an existing cluster? {create, join}: join
```

Note: The cluster join fails if the existing node is configured for a two-port cluster and in some cases when the new node uses the default port (8040, 8060, and 8080 systems, which have four cluster ports by default). You need to exit cluster setup, and then run the `network port modify -ipspace` command so that only the ports intended for the cluster network are checked during setup.

12. Respond **`yes`** when prompted to set storage failover to **`HA`** mode.

Example

```
Non-HA mode, Reboot node to activate HA

Warning: Ensure that the HA partner has started disk initialization before
rebooting this node to enable HA.
Do you want to reboot now to set storage failover (SFO) to HA mode?
{yes, no} [yes]: yes

Rebooting now
```

After the node reboots, the Cluster Setup wizard displays “Welcome to node setup” and prompts you to complete the node setup.

13. Respond to the remaining prompts as appropriate for your site.

The *Clustered Data ONTAP Software Setup Guide* for your version of ONTAP contains additional details.

14. If the system is a two-node switchless cluster configuration, create the cluster interfaces on the existing node to create cluster LIFs on the cluster ports: **`network interface create`**

Example

The following example shows how to create a cluster LIF on one of the node's cluster ports. The `-auto` parameter configures the LIF to use a link-local IP address:

```
cluster1::> network interface create -vserver Cluster -lif clus1 -role cluster -home-node node0 -
home-port e1a -auto true
```

15. After setup is complete, verify that the node is healthy and eligible to participate in the cluster: **`cluster show`**

Example

The following example shows a cluster after the second node (cluster1-02) has been joined to it:

```
cluster1::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
cluster1-01           true    true
cluster1-02           true    true
```

You can access the Cluster Setup wizard to change any of the values that you entered for the admin Storage Virtual Machine (SVM) or node SVM by using the `cluster setup` command.

Installing licenses for the new controller module

You must add licenses for the new controller module for any ONTAP services that require standard (node-locked) licenses. For features with standard licenses, each node in the cluster must have its own key for the feature.

About this task

For detailed information about licensing, see the knowledgebase article 3013749: Data ONTAP 8.2 Licensing Overview and References on the NetApp Support Site and the *System Administration Reference*.

Steps

1. If necessary, obtain license keys for the new node on the NetApp Support Site in the My Support section under Software licenses.

   If the site does not have the license keys you need, contact your sales or support representative.

2. Issue the following command to install each license key:

   `system license add -license-code` ***`license_key`***

   The *`license_key`* is 28 digits in length.

   Repeat this step for each required standard (node-locked) license.

Enabling storage failover on both controller modules and enabling cluster HA

You must enable storage failover on both controller modules and separately enable cluster HA.

Before you begin

The MetroCluster configuration must have previously been refreshed using the `metrocluster configure -refresh true` command.

About this task

This task must be performed on each MetroCluster site.

Steps

1. Enable storage failover:

   `storage failover modify -enabled true -node` ***`existing-node-name`***

   The single command enables storage failover on both controller modules.

2. Verify that storage failover is enabled:

   `storage failover show`

Example

The output should be similar to the following:

```
Node           Partner        Possible State Description
-------------- -------------- -------- -------------------------------------
old-ctlr        new-ctlr      true      Connected to new-ctlr
new-ctlr        old-ctlr      true      Connected to old-ctlr
2 entries were displayed.
```

3. Enable cluster HA:

   `cluster ha modify -configured true`

   Cluster high availability (HA) must be configured in a cluster if it contains only two nodes and it differs from the HA provided by storage failover.

Installing the firmware after adding a controller module

After adding the controller module, you must install the latest firmware on the new controller module so that the controller module functions properly with ONTAP.

Step

1. Download the most current version of firmware for your system and follow the instructions for downloading and installing the new firmware.

   NetApp Downloads: System Firmware and Diagnostics

Adding the new node as a LUN mapping's reporting node

You must manually add the second node as a reporting node to make the Multipath I/O (MPIO) operational.

Steps

1. Change the LUN mapping reporting nodes to include the new node: **`lun mapping add-reporting-nodes -vserver * -path * -igroup * -local-nodes true`**
2. Verify that the newly added node is the reporting node in addition to the existing nodes: **`lun mapping show -fields reporting-nodes`**

Example

```
grkna-cm-t3::*> lun mapping show -fields reporting-nodes

vserver   path                            igroup     reporting-nodes
--------- ------------------------------ ---------- ------------------------------
testvs    /vol/testvol1/testlun1         testig     ste-s8080-01a,ste-s8080-01b
```

Setting up Storage Encryption on the new controller module

If the existing system used Storage Encryption, you must configure the new controller module for Storage Encryption, including installing and setting up the key managers, certificates, and servers.

About this task

This procedure includes steps that are performed on both the existing controller module and the new controller module. Be sure to enter the command on the correct system.

You must enter the commands in the steps below in the nodeshell. For more information about the nodeshell, see the *System Administration Reference*.

Steps

1. On the *existing* controller module, enter the following commands to verify that the key server is still available:

   key_manager status

   key_manager query

   Example

   The following command checks the status of all key management servers linked to the storage system:

   ```
   storage-system> key_manager status
   Key server                       Status
   172.18.99.175                    Server is responding
   ```

   The following command checks the status of all key management servers linked to the storage system and displays additional information:

   ```
   storage-system> key_manager query
   Key server 172.18.99.175 is responding.

   Key server 172.18.99.175 reports 4 keys.

   Key tag                           Key ID
   --------                          -------
   storage-system                    080CF0C80...
   storage-system                    080CF0C80...
   storage-system                    080CF0C80...
   storage-system                    080CF0C80...
   ```

2. On the *new* controller module, complete the following steps to install the same SSL certificates that are on the existing controller module:

   a. Copy the certificate files to a temporary location on the storage system.

   b. Install the public certificate of the storage system by entering the following command at the storage system prompt:

      keymgr install cert *path/*client.pem

   c. Install the private certificate of the storage system by entering the following command at the storage system prompt:

      keymgr install cert */path/*client_private.pem

   d. Install the public certificate of the key management server by entering the following command at the storage system prompt:

      keymgr install cert */path/key_management_server_ipaddress*_CA.pem

e. If you are linking multiple key management servers to the storage system, repeat the preceding steps for each public certificate of each key management server.

3. On the *new* controller module, run the Storage Encryption setup wizard to set up and install the key servers.

   You must install the same key servers that are installed on the existing controller module.

   a. Enter the following command at the storage system prompt:

      key_manager setup

   b. Complete the steps in the wizard to configure Storage Encryption.

Example

The following example shows how to configure Storage Encryption:

```
storage-system*> key_manager setup
Found client certificate file client.pem.
Registration successful for client.pem.
Found client private key file client_private.pem.
Is this file protected by a passphrase? [no]:
Registration successful for client_private.pem.
Enter the IP address for a key server, 'q' to quit:  172.22.192.192
Enter the IP address for a key server, 'q' to quit:  q
Enter the TCP port number for kmip server [6001] :

You will now be prompted to enter a key tag name. The
key tag name is used to identify all keys belonging to this
Data ONTAP system. The default key tag name is based on the
system's hostname.

Would you like to use <storage-system> as the default key tag name? [yes]:

Registering 1 key servers...
Found client CA certificate file 172.22.192.192_CA.pem.
Registration successful for 172.22.192.192_CA.pem.
Registration complete.

You will now be prompted for a subset of your network configuration
setup.  These parameters will define a pre-boot network environment
allowing secure connections to the registered key server(s).

Enter network interface:  e0a
Enter IP address:  172.16.132.165
Enter netmask:   255.255.252.0
Enter gateway:  172.16.132.1

Do you wish to enter or generate a passphrase for the system's
encrypting drives at this time? [yes]:  yes

Would you like the system to autogenerate a passphrase? [yes]:  yes

Key ID: 080CDCB2000000000100000000000003FE505B0C5E3E76061EE48E02A29822C

Make sure that you keep a copy of your passphrase, key ID, and key tag
name in a secure location in case it is ever needed for recovery purposes.

Should the system lock all encrypting drives at this time? yes
Completed rekey on 4 disks: 4 successes, 0 failures, including 0 unknown key and 0
authentication failures.
Completed lock on 4 disks: 4 successes, 0 failures, including 0 unknown key and 0
authentication failures.
```

4. On the *existing* controller module, enter the applicable command to restore authentication keys either from all linked key management servers or from a specific one:

   - **key_manager restore -all**

- **`key_manager restore -key_server` *`key_server_ip_address`***

5. On the *existing* controller module, rekey all of the disks by entering the following command at the prompt:

   `key_manager rekey -keytag` *`key_tag`*

   `key_tag` is the key tag name specified in the setup wizard in step 3.

Verifying the configuration with the Config Advisor

The Config Advisor utility verifies that the controller modules are properly configured for failover. This utility checks licenses, network configurations, options, and so on, and provides output that shows when error conditions occur.

Steps

1. Go to the Config Advisor page on the NetApp Support Site at *support.netapp.com/NOW/download/tools/config_advisor/*.
2. Use the links and information on the page to download and run the tool.

Copyright information

Trademark information

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

doccomments@netapp.com

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277