Stage 6. Complete the upgrade

Upgrade controllers

NetApp
February 22, 2024

Table of Contents

Stage 6. Complete the upgrade

Overview

During Stage 6, you confirm that the new nodes are set up correctly and, if the new nodes are encryption-enabled, you configure and set up Storage Encryption or NetApp Volume Encryption. You should also decommission the old nodes and resume the SnapMirror operations.

1. Manage authentication using KMIP servers
2. Confirm that the new controllers are set up correctly
3. Set up Storage Encryption on the new controller module
4. Set up NetApp Volume or Aggregate Encryption on the new controller module
5. Decommission the old system
6. Resume SnapMirror operations

Manage authentication using KMIP servers

With ONTAP 9.5 and later, you can use Key Management Interoperability Protocol (KMIP) servers to manage authentication keys.

Steps

1. Add a new controller:

   ```
   security key-manager setup -node new_controller_name
   ```

2. Add the key manager:

   ```
   security key-manager -add key_management_server_ip_address
   ```

3. Verify that the key management servers are configured and available to all nodes in the cluster:

   ```
   security key-manager show -status
   ```

4. Restore the authentication keys from all linked key management servers to the new node:

   ```
   security key-manager restore -node new_controller_name
   ```

Confirm that the new controllers are set up correctly

To confirm correct setup, you enable the HA pair. You also verify that node3 and node4 can access each other's storage and that neither owns data LIFs belonging to other nodes on the cluster. In addition, you confirm that node3 owns node1's aggregates and that node4 owns node2's aggregates, and that the volumes for both nodes are online.

Steps

1. Enable storage failover by entering the following command on one of the nodes:

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

2. Verify that storage failover is enabled:

   ```
   storage failover show
   ```

   The following example shows the output of the command when storage failover is enabled:

   ```
   cluster::> storage failover show
                                   Takeover
   Node            Partner         Possible State Description
   -------------- -------------- -------- ------------------
   node3           node4           true     Connected to node4
   node4           node3           true     Connected to node3
   ```

3. Take one of the following actions:

If the cluster is a…	Description
Two-node cluster	Enable cluster high availability by entering the following command on either node: `cluster ha modify -configured true`
Cluster with more than two nodes	Go to Step 4.

4. Verify that node3 and node4 belong to the same cluster by entering the following command and examining the output:

   ```
   cluster show
   ```

5. Verify that node3 and node4 can access each other's storage by entering the following command and examining the output:

   ```
   storage failover show -fields local-missing-disks,partner-missing-disks
   ```

6. Verify that neither node3 nor node4 owns data LIFs home-owned by other nodes in the cluster by entering the following command and examining the output:

   ```
   network interface show
   ```

   If either node3 or node4 owns data LIFs home-owned by other nodes in the cluster, use the `network interface revert` command to revert the data LIFs to their home-owner.

7. Verify that node3 owns the aggregates from node1 and that node4 owns the aggregates from node2:

   ```
   storage aggregate show -owner-name node3
   storage aggregate show -owner-name node4
   ```

8. Determine whether any volumes are offline:

```
volume show -node node3 -state offline
volume show -node node4 -state offline
```

9. If any volumes are offline, compare them with the list of offline volumes that you captured in Step 19 (d) in *Prepare the nodes for upgrade*, and bring online any of the offline volumes, as required, by entering the following command, once for each volume:

```
volume online -vserver vserver_name -volume volume_name
```

10. Install new licenses for the new nodes by entering the following command for each node:

```
system license add -license-code license_code,license_code,license_code…
```

The license-code parameter accepts a list of 28 upper-case alphabetic character keys. You can add one license at a time, or you can add multiple licenses at once, each license key separated by a comma.

11. If self-encrypting drives are being used in the configuration and you have set the `kmip.init.maxwait` variable to `off` (for example, in Step 16 of *Install and boot node3*), you must unset the variable:

```
set diag; systemshell -node node_name -command sudo kenv -u -p
kmip.init.maxwait
```

12. To remove all of the old licenses from the original nodes, enter one of the following commands:

```
system license clean-up -unused -expired
system license delete -serial-number node_serial_number -package
licensable_package
```

   ◦ To delete all expired licenses, enter:

   ```
   system license clean-up -expired
   ```

   ◦ To delete all unused licenses, enter:

   ```
   system license clean-up -unused
   ```

   ◦ To delete a specific license from a cluster, enter the following commands on the nodes:

   ```
   system license delete -serial-number node1_serial_number -package *
   system license delete -serial-number node2_serial_number -package *
   ```

   The following output is displayed:

   ```
   Warning: The following licenses will be removed:
   <list of each installed package>
   Do you want to continue? {y|n}: y
   ```

   Enter `y` to remove all of the packages.

13. Verify that the licenses are correctly installed by entering the following command and examining its output:

```
system license show
```

You can compare the output with the output that you captured in of *Prepare the nodes for upgrade*.

14. Configure the SPs by performing the following command on both nodes:

    ```
    system service-processor network modify -node node_name
    ```

    Go to References to link to the *System Administration Reference* for information about the SPs and the *ONTAP 9 Commands: Manual Page Reference* for detailed information about the `system service-processor network modify` command.

15. If you want to set up a switchless cluster on the new nodes, go to References to link to the *Network Support Site* and follow the instructions in *Transitioning to a two-node switchless cluster*.

After you finish

If Storage Encryption is enabled on node3 and node4, complete the steps in Set up Storage Encryption on the new controller module. Otherwise, complete the steps in Decommission the old system.

Set up Storage Encryption on the new controller module

If the replaced controller or the HA partner of the new controller uses Storage Encryption, you must configure the new controller module for Storage Encryption, including installing SSL certificates and setting up key management servers.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

    ```
    security key-manager show -status
    ```

    ```
    security key-manager query
    ```

2. Add the key management servers listed in the previous step to the key management server list in the new controller.

    a. Add the key management server:

       ```
       security key-manager -add key_management_server_ip_address
       ```

    b. Repeat the previous step for each listed key management server.

       You can link up to four key management servers.

    c. Verify the that the key management servers were added successfully:

       ```
       security key-manager show
       ```

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

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

   a. Launch the key management server setup wizard on the new node:

```
security key-manager setup -node new_controller_name
```

   b. Complete the steps in the wizard to configure key management servers.

4. Restore authentication keys from all linked key management servers to the new node:

```
security key-manager restore -node new_controller_name
```

Set up NetApp Volume or Aggregate Encryption on the new controller module

If the replaced controller or the high-availability (HA) partner of the new controller uses NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE), you must configure the new controller module for NVE or NAE.

About this task

This procedure includes steps that are performed on the new controller module. You must enter the command on the correct node.

Steps

1. Verify that the key management servers are still available, their status, and their authentication key information:

For this ONTAP version…	Use this command…
ONTAP 9.6 or 9.7	`security key-manager key query -node node`
ONTAP 9.5 or earlier	`security key-manager key show`

2. Add the key management servers listed in the previous step to the key management server list in the new controller:

   a. Add the key management server by using the following command:

```
security key-manager -add key_management_server_ip_address
```

   b. Repeat the previous step for each listed key management server. You can link up to four key management servers.

   c. Verify the that the key management servers were added successfully by using the following command:

```
security key-manager show
```

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

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

   a. Launch the key management server setup wizard on the new node by using the following command:

```
security key-manager setup -node new_controller_name
```

    b. Complete the steps in the wizard to configure key management servers.

4. Restore authentication keys from all linked key management servers to the new node.

   ◦ Restore authentication for external key manager:

   ```
   security key-manager external restore
   ```

   This command needs the Onboard Key Manager (OKM) passphrase

   For more information, see the Knowledge Base article How to restore external key manager server configuration from the ONTAP boot menu.

   ◦ Restore authentication for the OKM:

For this ONTAP version…	Use this command…
All other ONTAP versions	`security key-manager onboard sync`
ONTAP 9.5	`security key-manager setup -node node_name`

After you finish

Check if any volumes were taken offline because authentication keys were not available or external key management servers could not be reached. Bring those volumes back online by using the `volume online` command.

Decommission the old system

After upgrading, you can decommission the old system through the NetApp Support Site. Decommissioning the system tells NetApp that the system is no longer in operation and removes it from support databases.

Steps

1. Refer to References to link to the *NetApp Support Site* and log in.

2. Select **Products > My Products** from the menu.

3. On the **View Installed Systems** page, choose which **Selection Criteria** you want to use to display information about your system.

   You can choose one of the following to locate your system:

   ◦ Serial Number (located on the back of the unit)

   ◦ Serial Numbers for My Location

4. Select **Go!**

   A table displays cluster information, including the serial numbers.

5. Locate the cluster in the table and select **Decommission this system** from the Product Tool Set drop-down menu.

Resume SnapMirror operations

You can resume SnapMirror transfers that were quiesced before upgrade and resume the SnapMirror relationships. The updates are on schedule after the upgrade is completed.

Steps

1. Verify the SnapMirror status on the destination:

   ```
   snapmirror show
   ```

2. Resume the SnapMirror relationship:

   ```
   snapmirror resume -destination-vserver vserver_name
   ```