Stage 6. Boot node2 with the replacement system modules

Upgrade controllers

NetApp
February 22, 2024

Table of Contents

Stage 6. Boot node2 with the replacement system modules

Overview

During Stage 6, you boot node2 with upgraded system modules and verify the upgraded node2 installation. If you are using NetApp Volume Encryption (NVE), you restore key-manager configuration. You also relocate node1 non-root aggregates and NAS data LIFs from node1 to the upgraded node2 and verify that the SAN LIFs exist on node2.

1. Boot node2 with the replacement system modules
2. Verify the node2 installation
3. Restore key-manager configuration on node2
4. Move non-root aggregates and NAS data LIFs back to node2

Boot node2 with the replacement system modules

Node2 with the replacement modules is now ready to boot. Upgrading by swapping the system modules involves moving only the console and management connections. This section provides the steps required to boot node2 with the replacement modules for the following upgrade configurations:

Old node2 controller	Replacement node2 system modules
AFF A220 configured as an ASA	ASA A150 controller module
AFF A220 AFF A200 AFF C190	AFF A150 controller module
FAS2620 FAS2720	FAS2820 controller module
AFF A700 configured as an ASA	ASA A900 controller and NVRAM modules
AFF A700	AFF A900 controller and NVRAM modules
FAS9000	FAS9500 controller and NVRAM modules

Steps

1. If you have NetApp Storage Encryption (NSE) drives installed, perform the following steps.

> If you have not already done so earlier in the procedure, see the Knowledge Base article How to tell if a drive is FIPS certified to determine the type of self-encrypting drives that are in use.

   a. Set `bootarg.storageencryption.support` to `true` or `false`:

If the following drives are in use…	Then…
NSE drives that conform to FIPS 140-2 Level 2 self-encryption requirements	`setenv bootarg.storageencryption.support` **`true`**
NetApp non-FIPS SEDs	`setenv bootarg.storageencryption.support` **`false`**

> You cannot mix FIPS drives with other types of drives on the same node or HA pair. You can mix SEDs with non-encrypting drives on the same node or HA pair.

b. Go to the special boot menu and select option `(10) Set Onboard Key Manager recovery secrets`.

   Enter the passphrase and the backup information that you recorded earlier procedure. See Manage storage encryption using the Onboard Key Manager.

2. Boot the node into the boot menu:

   ```
   boot_ontap menu
   ```

3. Reassign the old node2 disks to the replacement node2 by entering "22/7" and selecting the hidden option `boot_after_controller_replacement` when the node stops at the boot menu.

   After a short delay, you are prompted to enter the name of the node that is being replaced. If there are shared disks (also called Advanced Disk Partitioning (ADP) or partitioned disks), you are prompted to enter the node name of the HA partner.

   These prompts might get buried in the console messages. If you do not enter a node name or enter an incorrect name, you are prompted to enter the name again.

   > If `[localhost:disk.encryptNoSupport:ALERT]: Detected FIPS-certified encrypting drive` and, or, `[localhost:diskown.errorDuringIO:error]: error 3 (disk failed) on disk` errors occur, perform the following steps:
   >
   > 1. Halt the node at the LOADER prompt.
   > 2. Check and reset the storage encryption bootargs mentioned in Step 1.
   > 3. At the loader prompt, boot up:
   >
   >  ```
   >  boot_ontap
   >  ```

   You can use the following example as a reference:

Expand the console output example

```
LOADER-A> boot_ontap menu
.
.
<output truncated>
.
All rights reserved.
******************************
*                            *
* Press Ctrl-C for Boot Menu. *
*                            *
******************************
.
<output truncated>
.
Please choose one of the following:

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? 22/7

(22/7)                          Print this secret List
(25/6)                          Force boot with multiple filesystem
disks missing.
(25/7)                          Boot w/ disk labels forced to clean.
(29/7)                          Bypass media errors.
(44/4a)                         Zero disks if needed and create new
flexible root volume.
(44/7)                          Assign all disks, Initialize all
disks as SPARE, write DDR labels
.
.
<output truncated>
.
.
(wipeconfig)                        Clean all configuration on boot
```

```
device
(boot_after_controller_replacement) Boot after controller upgrade
(boot_after_mcc_transition)         Boot after MCC transition
(9a)                                Unpartition all disks and remove
their ownership information.
(9b)                                Clean configuration and
initialize node with partitioned disks.
(9c)                                Clean configuration and
initialize node with whole disks.
(9d)                                Reboot the node.
(9e)                                Return to main boot menu.



The boot device has changed. System configuration information could
be lost. Use option (6) to restore the system configuration, or
option (4) to initialize all disks and setup a new system.
Normal Boot is prohibited.

Please choose one of the following:

(1)  Normal Boot.
(2)  Boot without /etc/rc.
(3)  Change password.
(4)  Clean configuration and initialize all disks.
(5)  Maintenance mode boot.
(6)  Update flash from backup config.
(7)  Install new software first.
(8)  Reboot node.
(9)  Configure Advanced Drive Partitioning.
(10) Set Onboard Key Manager recovery secrets.
(11) Configure node for external key management.
Selection (1-11)? boot_after_controller_replacement

This will replace all flash-based configuration with the last backup
to disks. Are you sure you want to continue?: yes

.
.
<output truncated>
.
.
Controller Replacement: Provide name of the node you would like to
replace:<nodename of the node being replaced>
Changing sysid of node node1 disks.
Fetched sanown old_owner_sysid = 536940063 and calculated old sys id
```

```
= 536940063
Partner sysid = 4294967295, owner sysid = 536940063
.
.
<output truncated>
.
.
varfs_backup_restore: restore using /mroot/etc/varfs.tgz
varfs_backup_restore: attempting to restore /var/kmip to the boot
device
varfs_backup_restore: failed to restore /var/kmip to the boot device
varfs_backup_restore: attempting to restore env file to the boot
device
varfs_backup_restore: successfully restored env file to the boot
device wrote key file "/tmp/rndc.key"
varfs_backup_restore: timeout waiting for login
varfs_backup_restore: Rebooting to load the new varfs
Terminated
<node reboots>

System rebooting...

.
.
Restoring env file from boot media...
copy_env_file:scenario = head upgrade
Successfully restored env file from boot media...
Rebooting to load the restored env file...
.
System rebooting...
.
.
.
<output truncated>
.
.
.
.
WARNING: System ID mismatch. This usually occurs when replacing a
boot device or NVRAM cards!
Override system ID? {y|n} y
.
.
.
.
Login:
```

> The system IDs shown in the preceding example are example IDs. The actual system IDs of the nodes that you are upgrading will be different.
>
> Between entering node names at the prompt and the login prompt, the node reboots a few times to restore the environment variables, update firmware on the cards in the system, and for other ONTAP updates.

Verify the node2 installation

You must verify the node2 installation with the replacement system modules. Because there is no change to physical ports, you are not required to map the physical ports from the old node2 to the replacement node2.

About this task

After you boot node1 with the replacement system module, you verify that it is installed correctly. You must wait for node2 to join quorum and then resume the controller replacement operation.

At this point in the procedure, the operation pauses while node2 joins quorum.

Steps

1. Verify that node2 has joined quorum:

   ```
   cluster show -node node2 -fields health
   ```

   The output of the `health` field should be `true`.

2. Verify that node2 is part of the same cluster as node1 and that it is healthy:

   ```
   cluster show
   ```

3. Switch to advanced privilege mode:

   ```
   set advanced
   ```

4. Check the status of the controller replacement operation and verify that it is in a paused state and in the same state that it was in before node2 was halted to perform the physical tasks of installing new controllers and moving cables:

   ```
   system controller replace show
   ```

   ```
   system controller replace show-details
   ```

5. Resume the controller replacement operation:

   ```
   system controller replace resume
   ```

6. The controller replacement operation pauses for intervention with the following message:

```
Cluster::*> system controller replace show
Node          Status                        Error-Action
------------  ------------------------
-----------------------------------
Node2         Paused-for-intervention       Follow the instructions given
in
                                            Step Details
Node1         None

Step Details:
---------------------------------------------
To complete the Network Reachability task, the ONTAP network
configuration must be manually adjusted to match the new physical
network configuration of the hardware. This includes:


1. Re-create the interface group, if needed, before restoring VLANs. For
detailed commands and instructions, refer to the "Re-creating VLANs,
ifgrps, and broadcast domains" section of the upgrade controller
hardware guide for the ONTAP version running on the new controllers.
2. Run the command "cluster controller-replacement network displaced-
vlans show" to check if any VLAN is displaced.
3. If any VLAN is displaced, run the command "cluster controller-
replacement network displaced-vlans restore" to restore the VLAN on the
desired port.
2 entries were displayed.
```

In this procedure, section *Re-creating VLANs, ifgrps, and broadcast domains* has been renamed *Restore network configuration on node2*.

7. With the controller replacement in a paused state, proceed to Restore network configuration on node2.

Restore network configuration on node2

After you confirm that node2 is in quorum and can communicate with node1, verify that node1's VLANs, interface groups, and broadcast domains are seen on node2. Also, verify that all node2 network ports are configured in their correct broadcast domains.

About this task

For more information on creating and re-creating VLANs, interface groups, and broadcast domains, refer to References to link to the *Network Management* content.

Steps

1. List all the physical ports that are on upgraded node2:

```
network port show -node node2
```

All physical network ports, VLAN ports, and interface group ports on the node are displayed. From this output, you can see any physical ports that have been moved into the `Cluster` broadcast domain by ONTAP. You can use this output to aid in deciding which ports should be used as interface group member ports, VLAN base ports, or standalone physical ports for hosting LIFs.

2. List the broadcast domains on the cluster:

   ```
   network port broadcast-domain show
   ```

3. List network port reachability of all ports on node2:

   ```
   network port reachability show -node node2
   ```

   You should see output similar to the following example. The port and broadcast names vary.

   ```
   Cluster::*> network port reachability show -node local
   Node      Port     Expected Reachability                Reachability
   Status
   --------- -------- -----------------------------------
   ---------------------
   Node2
             e0M      Default:Mgmt                         no-reachability
             e10a     Default:Default-3                    ok
             e10b     Default:Default-4                    ok
             e11a     Cluster:Cluster                      no-reachability
             e11b     Cluster:Cluster                      no-reachability
             e11c     -                                    no-reachability
             e11d     -                                    no-reachability
             e2a      Default:Default-1                    ok
             e2b      Default:Default-2                    ok
             e9a      Default:Default                      no-reachability
             e9b      Default:Default                      no-reachability
             e9c      Default:Default                      no-reachability
             e9d      Default:Default                      no-reachability
   13 entries were displayed.
   ```

   In the preceding example, node2 has booted and joined quorum after controller replacement. It has several ports that have no reachability and are pending a reachability scan.

4. Repair the reachability for each of the ports on node2 with a reachability status other than `ok` by using the following command, in the following order:

   ```
   network port reachability repair -node node_name -port port_name
   ```

   a. Physical ports
   b. VLAN ports

   You should see output like the following example:

```
Cluster ::> reachability repair -node node2 -port e9d
```

```
Warning: Repairing port "node2:e9d" may cause it to move into a
different broadcast domain, which can cause LIFs to be re-homed away
from the port. Are you sure you want to continue? {y|n}:
```

A warning message, as shown in the preceding example, is expected for ports with a reachability status that might be different from the reachability status of the broadcast domain where it is currently located. Review the connectivity of the port and answer `y` or `n` as appropriate.

Verify that all physical ports have their expected reachability:

```
network port reachability show
```

As the reachability repair is performed, ONTAP attempts to place the ports in the correct broadcast domains. However, if a port's reachability cannot be determined and does not belong to any of the existing broadcast domains, ONTAP will create new broadcast domains for these ports.

5. Verify port reachability:

   ```
   network port reachability show
   ```

   When all ports are correctly configured and added to the correct broadcast domains, the `network port reachability show` command should report the reachability status as `ok` for all connected ports, and the status as `no-reachability` for ports with no physical connectivity. If any port reports a status other than these two, perform the reachability repair and add or remove ports from their broadcast domains as instructed in Step 4.

6. Verify that all ports have been placed into broadcast domains:

   ```
   network port show
   ```

7. Verify that all ports in the broadcast domains have the correct maximum transmission unit (MTU) configured:

   ```
   network port broadcast-domain show
   ```

8. Restore LIF home ports, specifying the Vserver and LIF home ports, if any, that need to be restored by using the following steps:

   a. List any LIFs that are displaced:

      ```
      displaced-interface show
      ```

   b. Restore LIF home nodes and home ports:

      ```
      displaced-interface restore-home-node -node node_name -vserver vserver_name
      -lif-name LIF_name
      ```

9. Verify that all LIFs have a home port and are administratively up:

```
network interface show -fields home-port,status-admin
```

Restore key-manager configuration on node2

If you are using NetApp Aggregate Encryption (NAE) or NetApp Volume Encryption (NVE) to encrypt volumes on the system that you are upgrading, the encryption configuration must be synchronized to the new nodes. If you do not resynchronize the key-manager, when you relocate the node2 aggregates from the upgraded node1 to the upgraded node2 by using ARL, failures might occur because node2 does not have the required encryption keys to bring encrypted volumes and aggregates online.

About this task

Synchronize the encryption configuration to the new nodes by performing the following steps:

Steps

1. Run the following command from node2:

   ```
   security key-manager onboard sync
   ```

2. Verify that the SVM-KEK key is restored to "true" on node2 before you relocate the data aggregates:

   ```
   ::> security key-manager key query -node node2 -fields restored -key
   -type SVM-KEK
   ```

   Example

   ```
   ::> security key-manager key query -node node2 -fields restored -key
   -type SVM-KEK

   node     vserver   key-server   key-id
   restored
   -------- --------- -----------  ----------------------------------------
   --------
   node2    svm1      ""           00000000000000000200000000000a008a81976
   true
                                   2190178f9350e071fbb90f00000000000000000
   ```

Move non-root aggregates and NAS data LIFs back to node2

After you verify network configuration on node2 and before you relocate aggregates from node1 to node2, you verify that the NAS data LIFs belonging to node2 that are currently on node1 are relocated from node1 to node2. You must also verify that the SAN LIFs exist on node2.

About this task

Remote LIFs handle traffic to SAN LUNs during the upgrade procedure. Moving SAN LIFs is not necessary for cluster or service health during the upgrade. SAN LIFs are not moved unless they need to be mapped to new ports. After you bring node2 online, you must verify that the LIFs are healthy and located on the appropriate ports.

Steps

1. Resume the relocation operation:

   ```
   system controller replace resume
   ```

   The system performs the following tasks:

   - Cluster quorum check
   - System ID check
   - Image version check
   - Target platform check
   - Network reachability check

   The operation pauses at this stage in the network reachability check.

2. Resume the relocation operation:

   ```
   system controller replace resume
   ```

   The system performs the following checks:

   - Cluster health check
   - Cluster LIF status check

   After performing these checks, the system relocates the non-root aggregates and NAS data LIFs back to node2, which is now running on the replacement controller.

   The controller replacement operation pauses after the resource relocation is complete.

3. Check the status of the aggregate relocation and NAS data LIF move operations:

   ```
   system controller replace show-details
   ```

   If the controller replacement procedure is paused, check and correct the error, if any, and then issue `resume` to continue the operation.

4. If necessary, restore and revert any displaced LIFs. List any displaced LIFs:

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

   If any LIFs are displaced, restore the home node back to node2:

   ```
   cluster controller-replacement network displaced-interface restore-home-node
   ```

5. Resume the operation to prompt the system to perform the required post-checks:

```
system controller replace resume
```

The system performs the following post-checks:

- Cluster quorum check
- Cluster health check
- Aggregates reconstruction check
- Aggregate status check
- Disk status check
- Cluster LIF status check
- Volume check