Installing NX-OS software and RCFs on Cisco® Nexus 5010, 5020, and 5596 cluster switches

Installing the NX-OS software and reference configuration files (RCFs) on Cisco Nexus 5010, 5020, 5596UP, and 5596T cluster switches is a nondisruptive procedure (NDU).

Before you begin

- The cluster must be a fully functioning cluster.

- You might want to check or set your desired boot variables in the RCF to reflect the desired boot images if you are installing only the NX-OS software and keeping your current RCF version.
  You might need to change the boot variables to reflect the current boot images. You must perform this operation before reapplying the RCF so that the correct version will be instantiated on future reboots.

- The ONTAP, NX-OS, and RCF versions that are supported in this procedure are listed in the switch compatibility matrix on the "Cisco Ethernet Switches" page.
  There can be command dependencies between command syntax in the RCF and NX-OS versions.
  Cisco Ethernet Switch

- You can see the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.
  Install and Upgrade Guides

About this task

The procedure in the "Installing the NX-OS software" section includes steps to install the RCF. If you only need to update or reapply the RCF on your Cisco cluster switch, then you should use the procedure provided in the "Installing the reference configuration file" section.

The examples in both procedures use Nexus 5596UP or 5596T switches, and use the following switch and node nomenclature:

- The names of the two Cisco switches are cs1 and cs2.

- clus1 and clus2 are the cluster logical interfaces (LIFs) that are typically associated with the switches cs1 and cs2, respectively.

- The SVMs are node1 and node2.

- The `cluster::*>` prompt indicates the name of the cluster.

- The cluster ports on each node are named e1a and e2a and are connected to switches cs1 and cs2, respectively.
  The *Hardware Universe* lists the current cluster ports supported on your platform.

- Cisco cluster switches use the following inter-switch link (ISL) ports:

  ◦ Nexus 5596UP or 5596T: ports 41 through 48

  ◦ Nexus 5010: ports 13 through 20

  ◦ Nexus 5020: ports 33 through 40

- Cisco cluster switches support the following node connections:

- ◦ Nexus 5596UP or 5596T: ports 1/1 through 1/40
- ◦ Nexus 5010: ports 1/1 through e1/12 (and ports 2/1 through 2/6 if an expansion module is installed)
- ◦ Nexus 5020: ports 1/1 through 1/32

Related tasks

Installing the NX-OS software

The NX-OS software installation includes steps for installing the reference configuration file (RCF).

About this task

The procedure begins with preparing the cluster interconnect connections to Cisco switch cs2.

Steps

1. Change the privilege level to advanced, entering **y** when prompted to continue: `set -privilege advanced`

   The advanced prompt (`*>`) appears.

2. Migrate clus2 to port e1a on the console of each node: `network interface migrate`

   Example

   The following command migrates clus2 to port e1a on node1 and node2:

   ```
   cluster::*> network interface migrate -vserver node1 -lif clus2 -source-node node1 -dest-
   node node1 -dest-port e1a
   ```

   ```
   cluster::*> network interface migrate -vserver node2 -lif clus2 -source-node node1 -dest-
   node node2 -dest-port e1a
   ```

3. Verify that the migration took place on the console of each node: `network interface show`

   Example

   The following example shows that clus2 was migrated to port e1a on both nodes:

   ```
   cluster::*> network interface show -role cluster

             Logical    Status     Network           Current       Current Is
   Vserver   Interface  Admin/Open Address/Mask      Node          Port    Home
   --------- ---------- ---------- ----------------- ------------- ------- ----
   node1
             clus1      up/up      10.10.10.1/16     node1         e1a     true
             clus2      up/up      10.10.10.2/16     node1         e1a     false
   node2
             clus1      up/up      10.10.10.1/16     node2         e1a     true
             clus2      up/up      10.10.10.2/16     node2         e1a     false
   ```

4. Shut down cluster port e2a on both nodes: `network port modify`

Example

The following command shuts down port e2a on both nodes:

```
cluster::*> network port modify -node node1 -port e2a -up-admin false

cluster::*> network port modify -node node2 -port e2a -up-admin false
```

5. Verify the status of cluster port e2a: `network port show`

Example

The following example shows that cluster port e2a was shut down on node1 and node2:

```
cluster::*> network port show -role cluster

                                    Auto-Negot  Duplex     Speed (Mbps)
Node   Port   Role         Link MTU  Admin/Oper  Admin/Oper Admin/Oper
------ ------ ------------ ---- -----  -----------  ---------- ------------
node1
       e1a    cluster      up   9000  true/true   full/full  auto/10000
       e2a    cluster      down 9000  true/true   full/full  auto/10000
node2
       e1a    cluster      up   9000  true/true   full/full  auto/10000
       e2a    cluster      down 9000  true/true   full/full  auto/10000
```

6. Shut down the ISL ports on the other Cisco switch, cs1.

Example

The following example shuts down ISL ports 41 through 48 on cs1 of a Nexus 5596UP or Nexus 5596T switch. The "About this task" section contains the correct ISL ports for your switch model.

```
cs1# configure
cs1(config)# interface ethernet 1/41-48
cs1(config-if-range)# shutdown
cs1(config-if-range)# exit
cs1(config)# exit
cs1(config)#
```

Note: For Nexus 5010 switches, you need to shut down ISL ports 1/13-20. For Nexus 5020 switches, you need to shut down ISL ports 1/33-40.

7. In switch cs2, copy the kickstart and system images to the switch bootflash using a transfer protocol, such as FTP, TFTP, SFTP, or SCP.

Example

The following command uses SFTP to copy the kickstart image to the switch bootflash on a Nexus 5596UP or Nexus 5596T switch:

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /tftpboot/n5000-uk9-kickstart.5.2.1.N1.8b.bin
Enter hostname for the sftp server: 10.55.27.160
Enter username: root

The authenticity of host '10.55.27.160 (10.55.27.160)' can't be established.
RSA key fingerprint is 31:a0:eg2:9d:c4:5a:1b:6c:74:27:6b:40:01:g5:48:5d.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '10.55.27.160' (RSA) to the list of known hosts.
root@10.55.27.160's password:
Connected to 10.55.27.160.
```

```
sftp> get /tftpboot/n5000-uk9-kickstart.5.2.1.N1.8b.bin /bootflash/n5000-uk9-kickstart.
5.2.1.N1.8b.bin
Fetching /tftpboot/n5000-uk9-kickstart.5.2.1.N1.8b.bin to /bootflash/n5000-uk9-kickstart.
5.2.1.N1.8b.bin
/tftpboot/n5000-uk9-kickstart.5.2.1.N1.8b.bin  100%   30MB   3.4MB/s   00:09
sftp> exit
```

The following command uses SFTP to copy the system image to the switch bootflash on a Nexus 5596UP or Nexus 5596T switch:

```
cs2#  copy sftp: bootflash: vrf management
Enter source filename: /tftpboot/n5000-uk9.5.2.1.N1.8b.bin
Enter hostname for the sftp server: 10.55.27.160
Enter username: root

root@10.55.27.160's password:
Connected to 10.55.27.160.
sftp> get /tftpboot/n5000-uk9.5.2.1.N1.8b.bin /bootflash/n5000-uk9.5.2.1.N1.8b.bin
Fetching /tftpboot/n5000-uk9.5.2.1.N1.8b.bin to /bootflash/n5000-uk9.5.2.1.N1.8b.bin
/tftpboot/n5000-uk9.5.2.1.N1.8b.bin              100% 165MB   3.2MB/s   00:52
sftp> exit
```

8. Verify that the bootflash has enough space for the new image.

    • If necessary, create space by removing any unnecessary files: **delete bootflash:*filename***

       The Cisco documentation contains more information about this command.

    • You can look at the size of the image to verify that you have enough space for the new image that is about to be downloaded: **dir bootflash:**

       If the desired image is already on the switch, then go to the next step.

    Example

    The following example shows the content of the bootflash file system, the amount of space used for the image, and the remaining space.

```
cs2# dir bootflash:
    31646720    Sep 16 18:18:51 2012  n5000-uk9-kickstart.5.2.1.N1.8b.bin
   173087826    Sep 16 19:00:17 2012  n5000-uk9.5.2.1.N1.8b.bin
    .
    .
    .
  Usage for bootflash://sup-local712589312 bytes used
  936034304 bytes free
 1648623616 bytes total
```

9. Verify that there are no incompatibility issues with the new image: **show incompatibility**

    Example

    The following example shows that there are no compatibility issues with the new `n5000-uk9.5.2.1.N1.1.bin` image:

```
cs2# show incompatibility system bootflash:/n5000-uk9.5.2.1.N1.8b.bin
No incompatible configurations
```

10. Install the kickstart and system images that you copied to the bootflash, responding **y** when prompted: **install all**

     This operation reboots the switch with the image that was just installed.

     Example

     The following command installs the kickstart and system images on switch cs2:

```
cs2# install all kickstart bootflash:n5000-uk9-kickstart.5.2.1.N1.8b.bin system
bootflash:n5000-uk9.5.2.1.N1.8b.bin
```

11. Verify that the kickstart and image files were installed correctly: **show version**

    It might take four to five minutes for the switch to reboot.

    Example

    The following example shows that the correct kickstart and image files were installed on switch cs2:

    ```
    cs2# show version
    Software
      BIOS:      version 3.6.0
      loader:    version N/A
      kickstart: version 5.2(1)N1(8b)
      system:    version 5.2(1)N1(8b)
      power-seq: Module 1: version v5.0
      uC:        version v1.0.0.2
      SFP uC:    Module 1: v1.0.0.0
      BIOS compile time:       05/09/2012
      kickstart image file is: bootflash:///n5000-uk9-kickstart.5.2.1.N1.8b.bin
      kickstart compile time:  7/12/2012 19:00:00 [07/12/2012 19:08:13]
      system image file is:    bootflash:///n5000-uk9.5.2.1.N1.8b.bin
      system compile time:     7/12/2012 19:00:00 [07/12/2012 20:23:16]
    ```

12. Verify that the bootflash has enough space for the new RCF.

    • If necessary, create space by removing any unnecessary files: **delete bootflash:*filename***
      The Cisco documentation contains more information about this command.

    • You can look at the size of the RCF to verify that you have enough space for the new RCF that is about to be
      downloaded: **dir bootflash:**
      If the desired RCF is already on the switch, then go to step *14* on page 6.

    Example

    The following example shows the content of the bootflash file system, the amount of space used for the RCF, and the
    remaining space.

    ```
    cs2# dir bootflash:
             0     Nov 13 01:14:17 2012  20121113_011417_poap_3498_init.log
           248     Apr 10 02:26:16 2012  convert_pfm1.log
           581     Apr 10 02:26:16 2012  fcoe_mgr_cnv.log
          4096     Nov 13 01:12:25 2012  lost+found/
          5027     Feb 28 18:50:37 2013  mts.log
      25140224     Mar 23 01:17:56 2012  n5000-uk9-kickstart.5.0.3.N2.1.bin
      31646720     Jul 17 06:19:54 2012  n5000-uk9-kickstart.5.2.1.N1.8b.bin
     147256572     Mar 23 01:18:34 2012  n5000-uk9.5.0.3.N2.1.bin
     172662417     Apr 10 01:41:08 2012  n5000-uk9.5.1.3.N2.1.bin
     173087826     Jul 17 02:19:16 2012  n5000-uk9.5.2.1.N1.1.bin
          3779     Apr 10 02:25:38 2012  stp.log.1
          4096     Jan 01 01:21:25 2009  vdc_2/
          4096     Jan 01 01:21:25 2009  vdc_3/
          4096     Jan 01 01:21:25 2009  vdc_4/
           348     Apr 10 02:26:16 2012  vfc_cnv.log

    Usage for bootflash://sup-local
      667828224 bytes used
      980795392 bytes free
     1648623616 bytes total
    ```

13. Copy the RCF to the switch bootflash using a transfer protocol such as FTP, TFTP, SFTP, or SCP.

Example

The following command copies the RCF to the switch bootflash using TFTP on a Nexus 5596UP or Nexus 5596T switch:

```
cs2# copy tftp: bootflash: vrf management
Enter source filename: NX5596T_RCF_v1.3.txt
Enter hostname for the tftp server: 10.10.10.120
Trying to connect to tftp server......
Connection to Server Established.

TFTP get operation was successful
Copy complete, now saving to disk (please wait)...
```

14. Apply the previously downloaded RCF to the bootflash: `copy bootflash`

Example

The following command installs the NX5596T_RCF_v1.3.txt RCF on a Nexus 5596T switch:

```
cs2# copy bootflash:NX5596_RCF_v1.3.txt running-config
```

15. Verify that the RCF version is the new one, that the node and port settings are correct, and that your site customizations are implemented: `show running-config`

Example

```
cs2# show running-config
```

Your output varies depending on your site configuration. You can check the port settings and then refer to the release notes as a reference for any changes specific to the RCF that you have installed.

16. To save the latest configuration, you need to copy the running-config file to the startup-config file.

Example

The following example copies the running-configuration file to the startup-configuration file:

```
cs2# copy running-config startup-config
[###################################] 100%
```

17. Bring up the ISL ports on the active switch cs1.

Example

The following command brings up ISL ports 41 through 48 on cs1 of a Nexus 5596UP or Nexus 5596T switch. The "About this task" section contains the correct ISL ports for your switch model.

```
cs1# configure
cs1(config)# interface ethernet 1/41-48
cs1(config-if-range)# no shutdown
cs1(config-if-range)#exit
cs1(config)#exit
cs1#
```

Note: For Nexus 5010 switches, you need to bring up ISL ports 1/13-20. For Nexus 5020 switches, you need to bring up ISL ports 1/33-40.

18. Verify that the ISLs are operational: **show port-channel summary**

    There should be a "(P)" after the Ethernet ports in the "Member Ports" column.

    Example

    The following example shows that the port-channel members 41 through 48 are up (**P**) on switch cs2:

    ```
    cs2# show port-channel summary
    Flags:  D - Down        P - Up in port-channel (members)
            I - Individual  H - Hot-standby (LACP only)
            s - Suspended   r - Module-removed
            S - Switched    R - Routed
            U - Up (port-channel)
            M - Not in use. Min-links not met
    --------------------------------------------------------------------------------
    Group Port-        Type      Protocol  Member Ports
          Channel
    --------------------------------------------------------------------------------
    1     Po1(SU)      Eth       LACP      Eth1/41(P)   Eth1/42(P)   Eth1/43(P)
                                           Eth1/44(P)   Eth1/45(P)   Eth1/46(P)
                                           Eth1/47(P)   Eth1/48(P)
    ```

19. Bring up cluster port e2a on both nodes: **network port modify**

    Example

    The following command brings up cluster port e2a on node1 and node2:

    ```
    cluster::*> network port modify -node node1 -port e2a -up-admin true

    cluster::*> network port modify -node node2 -port e2a -up-admin true
    ```

20. Verify that cluster port e2a is **up** on both nodes: **network port show -role cluster**

    Example

    The following example shows that the port e2a is **up** on both node1 and node2:

    ```
    cluster::*> network port show -role cluster

                                          Auto-Negot  Duplex     Speed (Mbps)
    Node   Port   Role          Link MTU  Admin/Oper  Admin/Oper Admin/Oper
    ------ ------ ------------- ---- ----- ----------- ---------- ------------
    node1
           e1a    cluster       up   9000  true/true   full/full  auto/10000
           e2a    cluster       up   9000  true/true   full/full  auto/10000
    node2
           e1a    cluster       up   9000  true/true   full/full  auto/10000
           e2a    cluster       up   9000  true/true   full/full  auto/10000
    ```

21. On both nodes, revert clus2 that is associated with port e2a: **network interface revert**

    The LIF might revert automatically, depending on your version of ONTAP.

    Example

    ```
    cluster::*> network interface revert -vserver node1 -lif clus2

    cluster::*> network interface revert -vserver node2 -lif clus2
    ```

22. Verify that the LIF is now home on both nodes: **network interface show -role cluster**

The LIF is home if the ports in the "Current Port" column have a status of **true** in the "Is Home" column, as shown in the following example:

```
cluster::*> network interface show -role cluster

            Logical     Status     Network            Current     Current Is
Vserver     Interface   Admin/Oper Address/Mask       Node        Port    Home
----------- ----------  ---------- ------------------ ----------- ------- ----
node1
            clus1       up/up      10.10.10.1/24      node1       e1a     true
            clus2       up/up      10.10.10.2/24      node1       e2a     true
node2
            clus1       up/up      10.10.10.1/24      node2       e1a     true
            clus2       up/up      10.10.10.2/24      node2       e2a     true
```

23. Verify the health and eligibility status of the nodes in the cluster: **cluster show**

Example

The following example shows that node1 and node2 are healthy and eligible:

```
cluster::*> cluster show
Node                    Health  Eligibility  Epsilon
--------------------- ------- ------------ ------------
node1                   true    true         false
node2                   true    true         false
```

24. Change the privilege level back to admin: **set -privilege admin**

25. Repeat step *1* on page 2 through step *24* on page 8 to upgrade the NX-OS software on the other Cisco switch, cs1.

Installing the reference configuration file

You can use "Installing the reference configuration file" (RCF) procedure to upgrade or reapply the RCF on the Cisco cluster switch.

Before you begin

You must have saved the configuration that is currently running on your switch.

Steps

1. Change the privilege level to advanced, entering **y** when prompted to continue: **set -privilege advanced**

The advanced prompt (***>**) appears.

2. Save your current switch configuration information: **copy running-config startup-config**

Example

The following example shows how to save your current switch configuration to the startup configuration file on the switch cs2:

```
cs2# copy running-config startup-config
[########################################] 100%
```

3. Migrate clus2 to port e1a on the console of each node: **network interface migrate**

Example

The following command migrates clus2 to port e1a on node1 and node2:

```
cluster::*> network interface migrate -vserver node1 -lif clus2 -source-node node1 -dest-
node node1 -dest-port e1a
```

```
cluster::*> network interface migrate -vserver node2 -lif clus2 -source-node node1 -dest-
node node2 -dest-port e1a
```

4. Verify that the migration took place on the console of each node: `network interface show`

Example

The following example shows that clus2 was migrated to port e1a on both nodes:

```
cluster::*> network interface show -role cluster

            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Open Address/Mask       Node          Port    Home
----------- ---------- ---------- ------------------ ------------- ------- ----
node1
            clus1      up/up      10.10.10.1/16      node1         e1a     true
            clus2      up/up      10.10.10.2/16      node1         e1a     false
node2
            clus1      up/up      10.10.10.1/16      node2         e1a     true
            clus2      up/up      10.10.10.2/16      node2         e1a     false
```

5. Shut down the cluster port e2a on both nodes: `network port modify`

Example

The following command shuts down port e2a on both nodes:

```
cluster::*> network port modify -node node1 -port e2a -up-admin false

cluster::*> network port modify -node node2 -port e2a -up-admin false
```

6. Verify the status of cluster port e2a: `network port show`

Example

The following example shows that cluster port e2a was shut down on node1 and node2:

```
cluster::*> network port show -role cluster

                                    Auto-Negot  Duplex      Speed (Mbps)
Node   Port   Role         Link MTU Admin/Oper  Admin/Oper  Admin/Oper
------ ------ ------------ ---- ----- ----------- ---------- ------------
node1
       e1a    cluster      up   9000  true/true   full/full   auto/10000
       e2a    cluster      down 9000  true/true   full/full   auto/10000
node2
       e1a    cluster      up   9000  true/true   full/full   auto/10000
       e2a    cluster      down 9000  true/true   full/full   auto/10000
```

7. Shut down the ISL ports on the other Cisco switch, cs1.

Example

The following example shuts down ISL ports 41 through 48 on cs1 of a Nexus 5596UP or Nexus 5596T switch. The "About this task" section contains the correct ISL ports for your switch model.

```
cs1# configure
cs1(config)# interface ethernet 1/41-48
cs1(config-if-range)# shutdown
cs1(config-if-range)# exit
cs1(config)# exit
cs1(config)#
```

Note: For Nexus 5010 switches, you need to shut down ISL ports 1/13-20. For Nexus 5020 switches, you need to shut down ISL ports 1/33-40.

8. Before applying the RCF, use the `no channel-group` command to temporarily disable the port channel to avoid syntax errors.

The RCF contains commands to create a port channel that cannot be executed if a port channel is already active.

Example

The following example shows ISL ports 41 through 48 on cs1 of a Nexus 5596UP or 5596T switch and the `no channel-group` command removes the configuration from the interface:

```
cs2 # configure
cs2(config)# interface ethernet 1/41-48
cs2(config-if-range)#no channel-group
cs2(config-if-range)#exit
cs2(config)#exit
cs2#
```

Note: For Nexus 5010 switches, you need to shut down ISL ports 1/13-20. For Nexus 5020 switches, you need to shut down ISL ports 1/33-40.

9. Verify that the bootflash has enough space for the new RCF.

   - If necessary, create space by removing any unnecessary files: **delete bootflash:*filename***
     The Cisco documentation contains more information about this command.

   - You can look at the size of the RCF to verify that you have enough space for the new RCF that is about to be downloaded: **dir bootflash:**
     If the desired RCF is already on the switch, then go to step *11* on page 11.

Example

The following example shows the content of the bootflash file system, the amount of space used for the RCF, and the remaining space.

```
cs2# dir bootflash:
         0     Nov 13 01:14:17 2012   20121113_011417_poap_3498_init.log
       248     Apr 10 02:26:16 2012   convert_pfm1.log
       581     Apr 10 02:26:16 2012   fcoe_mgr_cnv.log
      4096     Nov 13 01:12:25 2012   lost+found/
      5027     Feb 28 18:50:37 2013   mts.log
  25140224     Mar 23 01:17:56 2012   n5000-uk9-kickstart.5.0.3.N2.1.bin
  31646720     Jul 17 06:19:54 2012   n5000-uk9-kickstart.5.2.1.N1.8b.bin
 147256572     Mar 23 01:18:34 2012   n5000-uk9.5.0.3.N2.1.bin
 172662417     Apr 10 01:41:08 2012   n5000-uk9.5.1.3.N2.1.bin
 173087826     Jul 17 02:19:16 2012   n5000-uk9.5.2.1.N1.1.bin
      3779     Apr 10 02:25:38 2012   stp.log.1
      4096     Jan 01 01:21:25 2009   vdc_2/
```

```
        4096     Jan 01 01:21:25 2009  vdc_3/
        4096     Jan 01 01:21:25 2009  vdc_4/
         348     Apr 10 02:26:16 2012  vfc_cnv.log

 Usage for bootflash://sup-local
   667828224 bytes used
   980795392 bytes free
  1648623616 bytes total
```

10. Copy the RCF to the switch bootflash using a transfer protocol such as FTP, TFTP, SFTP, or SCP.

    Example

    The following command copies the RCF to the switch bootflash using TFTP on a Nexus 5596UP or Nexus 5596T switch:

    ```
    cs2# copy tftp: bootflash: vrf management
    Enter source filename: NX5596T_RCF_v1.3.txt
    Enter hostname for the tftp server: 10.10.10.120
    Trying to connect to tftp server......
    Connection to Server Established.

    TFTP get operation was successful
    Copy complete, now saving to disk (please wait)...
    ```

11. Apply the previously downloaded RCF to the bootflash: `copy bootflash`

    Example

    The following command installs the NX5596T_RCF_v1.3.txt RCF on a Nexus 5596T switch

    ```
    cs2# copy bootflash:NX5596_RCF_v1.3.txt running-config
    ```

12. Verify that the RCF version is the new one, that the node and port settings are correct, and that your site customizations are implemented: `show running-config`

    Example

    ```
    cs2# show running-config
    ```

    Your output varies depending on your site configuration. You can check the port settings and then refer to the release notes as a reference for any changes specific to the RCF that you have installed.

13. Bring up the ISL ports on the active switch cs1.

    Example

    The following command brings up ISL ports 41 through 48 on cs1 of a Nexus 5596UP or 5596T switch. The "About this task" section contains the correct ISL ports for your switch model.

    ```
    cs1# configure
    cs1(config)# interface ethernet 1/41-48
    cs1(config-if-range)# no shutdown
    cs1(config-if-range)#exit
    cs1(config)#exit
    cs1#
    ```

    Note: For Nexus 5010 switches, you need to shut down ISL ports 1/13-20. For Nexus 5020 switches, you need to shut down ISL ports 1/33-40.

14. Verify that the ISLs are operational: `show port-channel summary`

There should be a "(P)" after the Ethernet ports in the "Member Ports" column.

Example

The following example shows that the port-channel members 41 through 48 are up (**P**) on switch cs2:

```
cs2# show port-channel summary
Flags:  D - Down         P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
--------------------------------------------------------------------------------
Group Port-       Type        Protocol  Member Ports
      Channel
--------------------------------------------------------------------------------
1     Po1(SU)     Eth         LACP      Eth1/41(P)   Eth1/42(P)   Eth1/43(P)
                                        Eth1/44(P)   Eth1/45(P)   Eth1/46(P)
                                        Eth1/47(P)   Eth1/48(P)
```

15. Bring up cluster port e2a on both nodes: **network port modify**

    Example

    The following command brings up cluster port e2a on node1 and node2:

    ```
    cluster::*> network port modify -node node1 -port e2a -up-admin true

    cluster::*> network port modify -node node2 -port e2a -up-admin true
    ```

16. Verify that cluster port e2a is **up** on both nodes: **network port show -role cluster**

    Example

    The following example shows that the port e2a is **up** on both node1 and node2:

    ```
    cluster::*> network port show -role cluster

                                           Auto-Negot  Duplex     Speed (Mbps)
    Node   Port   Role          Link MTU   Admin/Oper  Admin/Oper Admin/Oper
    ------ ------ ------------- ---- ----- ----------- ---------- ------------
    node1
           e1a    cluster       up   9000  true/true   full/full  auto/10000
           e2a    cluster       up   9000  true/true   full/full  auto/10000
    node2
           e1a    cluster       up   9000  true/true   full/full  auto/10000
           e2a    cluster       up   9000  true/true   full/full  auto/10000
    ```

17. On both nodes, revert clus2 that is associated with port e2a: **network interface revert**

    The LIF might revert automatically, depending on your version of ONTAP.

    Example

    ```
    cluster::*> network interface revert -vserver node1 -lif clus2

    cluster::*> network interface revert -vserver node2 -lif clus2
    ```

18. Verify that the LIF is now home on both nodes: **network interface show -role cluster**

Example

The LIF is home if the ports in the "Current Port" column have a status of **true** in the "Is Home" column, as shown in the following example:

```
cluster::*> network interface show -role cluster

          Logical    Status     Network            Current     Current Is
Vserver   Interface  Admin/Oper Address/Mask       Node        Port    Home
--------- ---------- ---------- ------------------ ----------- ------- ----
node1
          clus1      up/up      10.10.10.1/24      node1       e1a     true
          clus2      up/up      10.10.10.2/24      node1       e2a     true
node2
          clus1      up/up      10.10.10.1/24      node2       e1a     true
          clus2      up/up      10.10.10.2/24      node2       e2a     true
```

19. Verify the health and eligibility status of the nodes in the cluster: **cluster show**

Example

The following example shows that node1 and node2 are healthy and eligible:

```
cluster::*> cluster show
Node                 Health  Eligibility  Epsilon
-------------------- ------- ------------ ------------
node1                true    true         false
node2                true    true         false
```

20. To save the latest configuration, you need to copy the running-config file to the startup-config file.

Example

The following example copies the running-configuration file to the startup-configuration file:

```
cs2# copy running-config startup-config
[########################################] 100%
```

21. Change the privilege level back to admin: **set -privilege admin**

22. Repeat step *1* on page 8 through step *21* on page 13 to upgrade the reference configuration file on the other Cisco switch, cs1.

Related information

support.netapp.com

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

Trademark information

Active IQ, AltaVault, Arch Design, ASUP, AutoSupport, Campaign Express, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Element, Fitness, Flash Accel, Flash Cache, Flash Pool, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, Fueled by SolidFire, GetSuccessful, Helix Design, LockVault, Manage ONTAP, MetroCluster, MultiStore, NetApp, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANscreen, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, Snap Creator, SnapCenter, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, SolidFire Helix, StorageGRID, SyncMirror, Tech OnTap, Unbound Cloud, and WAFL and other names are trademarks or registered trademarks of NetApp, Inc., in the United States, and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web.

http://www.netapp.com/us/legal/netapptmlist.aspx