



Cisco Nexus 9336C-FX2

Cluster and storage switches

NetApp
April 05, 2024

This PDF was generated from <https://docs.netapp.com/us-en/ontap-systems-switches/switch-cisco-9336c-fx2/configure-switch-overview-9336c-cluster.html> on April 05, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Cisco Nexus 9336C-FX2 1
 - Overview 1
 - Install hardware 5
 - Configure software 16
 - Migrate switches 72
 - Replace switches 127

Cisco Nexus 9336C-FX2

Overview

Overview of installation and configuration for Cisco Nexus 9336C-FX2 cluster switches

The Cisco Nexus 9336C-FX2 cluster switch is part of the Cisco Nexus 9000 platform and can be installed in a NetApp system cabinet. Cluster switches allow you to build ONTAP clusters with more than two nodes.

Initial configuration overview

To initially configure a Cisco Nexus 9336C-FX2 switch on systems running ONTAP, follow these steps:

1. [Complete the Cisco Nexus 9336C-FX2 cabling worksheet](#). The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.
2. [Install the switch](#). Set up the switch hardware.
3. [Configure the 9336C-FX2 cluster switch](#). Set up the Cisco Nexus 9336C-FX2 switch.
4. [Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet](#). Depending on your configuration, you can install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet with the standard brackets that are included with the switch.
5. [Prepare to install NX-OS software and RCF](#). Follow preliminary procedures in preparation for installing the Cisco NX-OS software and reference configuration files (RCFs).
6. [Install the NX-OS software](#). Install the NX-OS software on the Nexus 9336C-FX2 cluster switch.
7. [Install the Reference Configuration File \(RCF\)](#). Install the RCF after setting up the Nexus 9336C-FX2 switch for the first time. You can also use this procedure to upgrade your RCF version.

Additional information

Before you begin installation or maintenance, be sure to review the following:

- [Configuration requirements](#)
- [Components and part numbers](#)
- [Required documentation](#)
- [Smart Call Home requirements](#)

Configuration requirements for Cisco Nexus 9336C-FX2 cluster switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review configuration and network requirements.

ONTAP support

From ONTAP 9.9.1, you can use Cisco Nexus 9336C-FX2 switches to combine storage and cluster functionality into a shared switch configuration.

If you want to build ONTAP clusters with more than two nodes, you need two supported network switches.

Configuration requirements

Make sure that:

- You have the appropriate number and type of cables and cable connectors for your switches. See the [Hardware Universe](#).
- Depending on the type of switch you are initially configuring, you need to connect to the switch console port with the included console cable.

Network requirements

You need the following network information for all switch configurations.

- IP subnet for management network traffic
- Host names and IP addresses for each of the storage system controllers and all applicable switches
- Most storage system controllers are managed through the e0M interface by connecting to the Ethernet service port (wrench icon). On AFF A800 and AFF A700s systems, the e0M interface uses a dedicated Ethernet port.
- Refer to the [Hardware Universe](#) for the latest information.

For more information about the initial configuration of your switch, see the following guide: [Cisco Nexus 9336C-FX2 Installation and Upgrade Guide](#).

Components and part numbers for Cisco Nexus 9336C-FX2 cluster switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review the list of components and part numbers.

The following table lists the part number and description for the 9336C-FX2 switch, fans, and power supplies:

Part number	Description
X190200-CS-PE	N9K-9336C-FX2, CS, PTSX, 36PT10/25/40/100GQSFP28
X190200-CS-PI	N9K-9336C-FX2, CS, PSIN, 36PT10/25/40/100GQSFP28
X190210-FE-PE	N9K-9336C, FTE, PTSX, 36PT10/25/40/100GQSFP28
X190210-FE-PI	N9K-9336C, FTE, PSIN, 36PT10/25/40/100GQSFP28
X190002	Accessory Kit X190001/X190003
X-NXA-PAC-1100W-PE2	N9K-9336C AC 1100W PSU - Port side exhaust airflow
X-NXA-PAC-1100W-PI2	N9K-9336C AC 1100W PSU - Port side Intake airflow

Part number	Description
X-NXA-FAN-65CFM-PE	N9K-9336C 65CFM, Port side exhaust airflow
X-NXA-FAN-65CFM-PI	N9K-9336C 65CFM, Port side intake airflow

Documentation requirements for Cisco Nexus 9336C-FX2 switches

For Cisco Nexus 9336C-FX2 switch installation and maintenance, be sure to review specific switch and controller documentation to set up your Cisco 9336-FX2 switches and ONTAP cluster.

Switch documentation

To set up the Cisco Nexus 9336C-FX2 switches, you need the following documentation from the [Cisco Nexus 9000 Series Switches Support](#) page:

Document title	Description
<i>Nexus 9000 Series Hardware Installation Guide</i>	Provides detailed information about site requirements, switch hardware details, and installation options.
<i>Cisco Nexus 9000 Series Switch Software Configuration Guides</i> (choose the guide for the NX-OS release installed on your switches)	Provides initial switch configuration information that you need before you can configure the switch for ONTAP operation.
<i>Cisco Nexus 9000 Series NX-OS Software Upgrade and Downgrade Guide</i> (choose the guide for the NX-OS release installed on your switches)	Provides information on how to downgrade the switch to ONTAP supported switch software, if necessary.
<i>Cisco Nexus 9000 Series NX-OS Command Reference Master Index</i>	Provides links to the various command references provided by Cisco.
<i>Cisco Nexus 9000 MIBs Reference</i>	Describes the Management Information Base (MIB) files for the Nexus 9000 switches.
<i>Nexus 9000 Series NX-OS System Message Reference</i>	Describes the system messages for Cisco Nexus 9000 series switches, those that are informational, and others that might help diagnose problems with links, internal hardware, or the system software.
<i>Cisco Nexus 9000 Series NX-OS Release Notes</i> (choose the notes for the NX-OS release installed on your switches)	Describes the features, bugs, and limitations for the Cisco Nexus 9000 Series.

Document title	Description
Regulatory Compliance and Safety Information for Cisco Nexus 9000 Series	Provides international agency compliance, safety, and statutory information for the Nexus 9000 series switches.

ONTAP systems documentation

To set up an ONTAP system, you need the following documents for your version of the operating system from the [ONTAP 9 Documentation Center](#).

Name	Description
Controller-specific <i>Installation and Setup Instructions</i>	Describes how to install NetApp hardware.
ONTAP documentation	Provides detailed information about all aspects of the ONTAP releases.
Hardware Universe	Provides NetApp hardware configuration and compatibility information.

Rail kit and cabinet documentation

To install a Cisco 9336-FX2 switch in a NetApp cabinet, see the following hardware documentation.

Name	Description
42U System Cabinet, Deep Guide	Describes the FRUs associated with the 42U system cabinet, and provides maintenance and FRU replacement instructions.
Install a Cisco 9336-FX2 switch in a NetApp Cabinet	Describes how to install a Cisco Nexus 9336C-FX2 switch in a four-post NetApp cabinet.

Smart Call Home requirements

To use the Smart Call Home feature, review the following guidelines.

Smart Call Home monitors the hardware and software components on your network. When a critical system configuration occurs, it generates an email-based notification and raises an alert to all the recipients that are configured in your destination profile. To use Smart Call Home, you must configure a cluster network switch to communicate using email with the Smart Call Home system. In addition, you can optionally set up your cluster network switch to take advantage of Cisco's embedded Smart Call Home support feature.

Before you can use Smart Call Home, be aware of the following considerations:

- An email server must be in place.
- The switch must have IP connectivity to the email server.
- The contact name (SNMP server contact), phone number, and street address information must be configured. This is required to determine the origin of messages received.

- A CCO ID must be associated with an appropriate Cisco SMARTnet Service contract for your company.
- Cisco SMARTnet Service must be in place for the device to be registered.

The [Cisco support site](#) contains information about the commands to configure Smart Call Home.

Install hardware

Complete the Cisco Nexus 9336C-FX2 cabling worksheet

If you want to document the supported platforms, download a PDF of this page and complete the cabling worksheet.

The sample cabling worksheet provides examples of recommended port assignments from the switches to the controllers. The blank worksheet provides a template that you can use in setting up your cluster.

Sample cabling worksheet

The sample port definition on each pair of switches is as follows:

Cluster switch A		Cluster switch B	
Switch port	Node and port usage	Switch port	Node and port usage
1	4x10GbE node 1	1	4x10GbE node 1
2	4x10GbE node 2	2	4x10GbE node 2
3	4x10GbE node 3	3	4x10GbE node 3
4	4x25GbE node 4	4	4x25GbE node 4
5	4x25GbE node 5	5	4x25GbE node 5
6	4x25GbE node 6	6	4x25GbE node 6
7	40/100GbE node 7	7	40/100GbE node 7
8	40/100GbE node 8	8	40/100GbE node 8
9	40/100GbE node 9	9	40/100GbE node 9
10	40/100GbE node 10	10	40/100GbE node 10
11	40/100GbE node 11	11	40/100GbE node 11
12	40/100GbE node 12	12	40/100GbE node 12

Cluster switch A		Cluster switch B	
13	40/100GbE node 13	13	40/100GbE node 13
14	40/100GbE node 14	14	40/100GbE node 14
15	40/100GbE node 15	15	40/100GbE node 15
16	40/100GbE node 16	16	40/100GbE node 16
17	40/100GbE node 17	17	40/100GbE node 17
18	40/100GbE node 18	18	40/100GbE node 18
19	40/100GbE node 19	19	40/100GbE node 19
20	40/100GbE node 20	20	40/100GbE node 20
21	40/100GbE node 21	21	40/100GbE node 21
22	40/100GbE node 22	22	40/100GbE node 22
23	40/100GbE node 23	23	40/100GbE node 23
24	40/100GbE node 24	24	40/100GbE node 24
25 through 34	Reserved	25 through 34	Reserved
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

Blank cabling worksheet

You can use the blank cabling worksheet to document the platforms that are supported as nodes in a cluster. The *Supported Cluster Connections* section of the [Hardware Universe](#) defines the cluster ports used by the platform.

Cluster switch A		Cluster switch B	
1		1	
2		2	

Cluster switch A		Cluster switch B	
3		3	
4		4	
5		5	
6		6	
7		7	
8		8	
9		9	
10		10	
11		11	
12		12	
13		13	
14		14	
15		15	
16		16	
17		17	
18		18	
19		19	
20		20	
21		21	
22		22	
23		23	
24		24	

Cluster switch A		Cluster switch B	
25 through 34	Reserved	25 through 34	Reserved
35	100GbE ISL to switch B port 35	35	100GbE ISL to switch A port 35
36	100GbE ISL to switch B port 36	36	100GbE ISL to switch A port 36

See the [Hardware Universe](#) for more information on switch ports.

Install the 9336C-FX2 cluster switch

Follow this procedure to set up and configure the Cisco Nexus 9336C-FX2 switch.

What you'll need

- Access to an HTTP, FTP, or TFTP server at the installation site to download the applicable NX-OS and Reference Configuration File (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco Software Download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at mysupport.netapp.com. All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- [Required switch and ONTAP documentation](#).

Steps

1. Rack the cluster network and management network switches and controllers.

If you are installing the...	Then...
Cisco Nexus 9336C-FX2 in a NetApp system cabinet	See the <i>Installing a Cisco Nexus 9336C-FX2 cluster switch and pass-through panel in a NetApp cabinet</i> guide for instructions to install the switch in a NetApp cabinet.
Equipment in a Telco rack	See the procedures provided in the switch hardware installation guides and the NetApp installation and setup instructions.

2. Cable the cluster network and management network switches to the controllers using the completed cabling worksheets.
3. Power on the cluster network and management network switches and controllers.

What's next?

Go to [Configure the Cisco Nexus 9336C-FX2 switch](#).

Configure the 9336C-FX2 cluster switch

Follow this procedure to configure the Cisco Nexus 9336C-FX2 switch.

What you'll need

- Access to an HTTP, FTP, or TFTP server at the installation site to download the applicable NX-OS and Reference Configuration File (RCF) releases.
- Applicable NX-OS version, downloaded from the [Cisco software download](#) page.
- Applicable licenses, network and configuration information, and cables.
- Completed [cabling worksheets](#).
- Applicable NetApp cluster network and management network RCFs downloaded from the NetApp Support Site at [mysupport.netapp.com](#). All Cisco cluster network and management network switches arrive with the standard Cisco factory-default configuration. These switches also have the current version of the NX-OS software but do not have the RCFs loaded.
- [Required switch and ONTAP documentation](#).


Steps

1. Perform an initial configuration of the cluster network switches.

Provide applicable responses to the following initial setup questions when you first boot the switch. Your site's security policy defines the responses and services to enable.

Prompt	Response
Abort Auto Provisioning and continue with normal setup? (yes/no)	Respond with yes . The default is no.
Do you want to enforce secure password standard? (yes/no)	Respond with yes . The default is yes.
Enter the password for admin.	The default password is "admin"; you must create a new, strong password. A weak password can be rejected.
Would you like to enter the basic configuration dialog? (yes/no)	Respond with yes at the initial configuration of the switch.
Create another login account? (yes/no)	Your answer depends on your site's policies on alternate administrators. The default is no .
Configure read-only SNMP community string? (yes/no)	Respond with no . The default is no.
Configure read-write SNMP community string? (yes/no)	Respond with no . The default is no.
Enter the switch name.	Enter the switch name, which is limited to 63 alphanumeric characters.

Prompt	Response
Continue with Out-of-band (mgmt0) management configuration? (yes/no)	Respond with yes (the default) at that prompt. At the mgmt0 IPv4 address: prompt, enter your IP address: ip_address.
Configure the default-gateway? (yes/no)	Respond with yes . At the IPv4 address of the default-gateway: prompt, enter your default_gateway.
Configure advanced IP options? (yes/no)	Respond with no . The default is no.
Enable the telnet service? (yes/no)	Respond with no . The default is no.
Enabled SSH service? (yes/no)	Respond with yes . The default is yes. <div>  <p>SSH is recommended when using Cluster Switch Health Monitor (CSHM) for its log collection features. SSHv2 is also recommended for enhanced security.</p> </div>
Enter the type of SSH key you want to generate (dsa/rsa/rsa1).	The default is rsa .
Enter the number of key bits (1024-2048).	Enter the number of key bits from 1024 to 2048.
Configure the NTP server? (yes/no)	Respond with no . The default is no.
Configure default interface layer (L3/L2)	Respond with L2 . The default is L2.
Configure default switch port interface state (shut/noshut)	Respond with noshut . The default is noshut.
Configure CoPP system profile (strict/moderate/lenient/dense)	Respond with strict . The default is strict.
Would you like to edit the configuration? (yes/no)	You should see the new configuration at this point. Review and make any necessary changes to the configuration you just entered. Respond with no at the prompt if you are satisfied with the configuration. Respond with yes if you want to edit your configuration settings.

Prompt	Response
Use this configuration and save it? (yes/no)	<p>Respond with yes to save the configuration. This automatically updates the kickstart and system images.</p> <div>  <p>If you do not save the configuration at this stage, none of the changes will be in effect the next time you reboot the switch.</p> </div>

2. Verify the configuration choices you made in the display that appears at the end of the setup, and make sure that you save the configuration.
3. Check the version on the cluster network switches, and if necessary, download the NetApp-supported version of the software to the switches from the [Cisco software download](#) page.

What's next?

Optionally, you can [install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet](#). Otherwise, go to [Prepare to install NX-OS and RCF](#).

Install a Cisco Nexus 9336C-FX2 switch in a NetApp cabinet

Depending on your configuration, you might need to install the Cisco Nexus 9336C-FX2 switch and pass-through panel in a NetApp cabinet. Standard brackets are included with the switch.

What you'll need

- The pass-through panel kit, which is available from NetApp (part number X8784-R6).

The NetApp pass-through panel kit contains the following hardware:

- One pass-through blanking panel
- Four 10-32 x .75 screws
- Four 10-32 clip nuts
- For each switch, eight 10-32 or 12-24 screws and clip nuts to mount the brackets and slider rails to the front and rear cabinet posts.
- The Cisco standard rail kit to install the switch in a NetApp cabinet.



The jumper cords are not included with the pass-through kit and should be included with your switches. If they were not shipped with the switches, you can order them from NetApp (part number X1558A-R6).


- For initial preparation requirements, kit contents, and safety precautions, see [Cisco Nexus 9000 Series Hardware Installation Guide](#).

Steps

1. Install the pass-through blanking panel in the NetApp cabinet.
 - a. Determine the vertical location of the switches and blanking panel in the cabinet.


In this procedure, the blanking panel is installed in U40.

- b. Install two clip nuts on each side in the appropriate square holes for front cabinet rails.
- c. Center the panel vertically to prevent intrusion into adjacent rack space, and then tighten the screws.
- d. Insert the female connectors of both 48-inch jumper cords from the rear of the panel and through the brush assembly.




(1) Female connector of the jumper cord.

2. Install the rack-mount brackets on the Nexus 9336C-FX2 switch chassis.
 - a. Position a front rack-mount bracket on one side of the switch chassis so that the mounting ear is aligned with the chassis faceplate (on the PSU or fan side), and then use four M4 screws to attach the bracket to the chassis.




- b. Repeat step [2a](#) with the other front rack-mount bracket on the other side of the switch.
 - c. Install the rear rack-mount bracket on the switch chassis.
 - d. Repeat step [2c](#) with the other rear rack-mount bracket on the other side of the switch.
3. Install the clip nuts in the square hole locations for all four IEA posts.



The two 9336C-FX2 switches are always mounted in the top 2U of the cabinet RU41 and 42.

4. Install the slider rails in the cabinet.

- a. Position the first slider rail at the RU42 mark on the back side of the rear left post, insert screws with the matching thread type, and then tighten the screws with your fingers.



(1) As you gently slide the slider rail, align it to the screw holes in the rack.

(2) Tighten the screws of the slider rails to the cabinet posts.

- b. Repeat step [4a](#) for the right-side rear post.


c. Repeat steps 4a and 4b at the RU41 locations on the cabinet.

5. Install the switch in the cabinet.



This step requires two people: one person to support the switch from the front and another to guide the switch into the rear slider rails.


a. Position the back of the switch at RU41.



(1) As the chassis is pushed toward the rear posts, align the two rear rack-mount guides with the slider rails.

(2) Gently slide the switch until the front rack-mount brackets are flush with the front posts.

b. Attach the switch to the cabinet.



(1) With one person holding the front of the chassis level, the other person should fully tighten the four rear screws to the cabinet posts.

c. With the chassis now supported without assistance, fully tighten the front screws to the posts.

- d. Repeat steps 5a through 5c for the second switch at the RU42 location.



By using the fully installed switch as a support, it is not necessary to hold the front of the second switch during the installation process.

6. When the switches are installed, connect the jumper cords to the switch power inlets.
7. Connect the male plugs of both jumper cords to the closest available PDU outlets.



To maintain redundancy, the two cords must be connected to different PDUs.

8. Connect the management port on each 9336C-FX2 switch to either of the management switches (if ordered) or connect them directly to your management network.

The management port is the upper-right port located on the PSU side of the switch. The CAT6 cable for each switch needs to be routed through the pass-through panel after the switches are installed to connect to the management switches or management network.

What's next?

[Configure the Cisco Nexus 9336C-FX2 switch.](#)

Review cabling and configuration considerations

Before configuring your Cisco 9336C-FX2 switch, review the following considerations.

Support for NVIDIA CX6, CX6-DX, and CX7 Ethernet ports

If connecting a switch port to an ONTAP controller using NVIDIA ConnectX-6 (CX6), ConnectX-6 Dx (CX6-DX), or ConnectX-7 (CX7) NIC ports, you must hard-code the switch port speed.

```
(cs1)(config)# interface Ethernet1/19
For 100GbE speed:
(cs1)(config-if)# speed 100000
For 40GbE speed:
(cs1)(config-if)# speed 40000
(cs1)(config-if)# no negotiate auto
(cs1)(config-if)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

See the [Hardware Universe](#) for more information on switch ports.

25GbE FEC requirements

FAS2820 e0a/e0b ports

FAS2820 e0a and e0b ports require FEC configuration changes to link up with 9336C-FX2 switch ports. For switch ports e0a and e0b, the fec setting is set to `rs-cons16`.

```
(cs1)(config)# interface Ethernet1/8-9
(cs1)(config-if-range)# fec rs-cons16
(cs1)(config-if-range)# exit
(cs1)(config)# exit
Save the changes:
(cs1)# copy running-config startup-config
```

Configure software

Software install workflow for Cisco Nexus 9336C-FX2 cluster switches

To install and configure the software for a Cisco Nexus 9336C-FX2 switch, follow these steps:

1. [Prepare to install NX-OS software and RCF.](#)
2. [Install the NX-OS software.](#)
3. [Install the Reference Configuration File \(RCF\).](#)

Install the RCF after setting up the Nexus 9336C-FX2 switch for the first time. You can also use this procedure to upgrade your RCF version.

Available RCF configurations

The following table describes the RCFs available for different configurations. Choose the RCF applicable to your configuration.

For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

RCF name	Description
2-Cluster-HA-Breakout	Supports two ONTAP clusters with at least eight nodes, including nodes that use shared Cluster+HA ports.
4-Cluster-HA-Breakout	Supports four ONTAP clusters with at least four nodes, including nodes that use shared Cluster+HA ports.
1-Cluster-HA	All ports are configured for 40/100GbE. Supports shared cluster/HA traffic on ports. Required for AFF A320, AFF A250, and FAS500f systems. Additionally, all ports can be used as dedicated cluster ports.
1-Cluster-HA-Breakout	Ports are configured for 4x10GbE breakout, 4x25GbE breakout (RCF 1.6+ on 100GbE switches), and 40/100GbE. Supports shared cluster/HA traffic on ports for nodes that use shared cluster/HA ports: AFF A320, AFF A250, and FAS500f systems. Additionally, all ports can be used as dedicated cluster ports.

RCF name	Description
Cluster-HA-Storage	Ports are configured for 40/100GbE for Cluster+HA, 4x10GbE Breakout for Cluster and 4x25GbE Breakout for Cluster+HA, and 100GbE for each Storage HA Pair.
Cluster	Two flavors of RCF with different allocations of 4x10GbE ports (breakout) and 40/100GbE ports. All FAS/AFF nodes are supported, except for AFF A320, AFF A250, and FAS500f systems.
Storage	All ports are configured for 100GbE NVMe storage connections.

Prepare to install NX-OS software and RCF

Before you install the NX-OS software and the Reference Configuration File (RCF), follow this procedure.

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the two Cisco switches are cs1 and cs2.
- The node names are cluster1-01 and cluster1-02.
- The cluster LIF names are cluster1-01_clus1 and cluster1-01_clus2 for cluster1-01 and cluster1-02_clus1 and cluster1-02_clus2 for cluster1-02.
- The `cluster1::*>` prompt indicates the name of the cluster.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=x h`

where x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

3. Display how many cluster interconnect interfaces are configured in each node for each cluster interconnect switch:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
cluster1-02/cdp	e0a	cs1	Eth1/2	N9K-
C9336C	e0b	cs2	Eth1/2	N9K-
C9336C				
cluster1-01/cdp	e0a	cs1	Eth1/1	N9K-
C9336C	e0b	cs2	Eth1/1	N9K-
C9336C				

4 entries were displayed.

4. Check the administrative or operational status of each cluster interface.
 - a. Display the network port attributes:

```
`network port show -ipspace Cluster`
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: cluster1-02
```

Health					Speed (Mbps)	
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status						
-----	-----	-----	-----	----	----	-----

e0a	Cluster	Cluster		up	9000	auto/10000
healthy						
e0b	Cluster	Cluster		up	9000	auto/10000
healthy						

```
Node: cluster1-01
```

Health					Speed (Mbps)	
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status						
-----	-----	-----	-----	----	----	-----

e0a	Cluster	Cluster		up	9000	auto/10000
healthy						
e0b	Cluster	Cluster		up	9000	auto/10000
healthy						

```
4 entries were displayed.
```

b. Display information about the LIFs:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Vserver Port	Home	Logical Current Is Interface	Status Admin/Oper	Network Address/Mask	Node

Cluster					
		cluster1-01_clus1	up/up	169.254.209.69/16	
cluster1-01		e0a true			
		cluster1-01_clus2	up/up	169.254.49.125/16	
cluster1-01		e0b true			
		cluster1-02_clus1	up/up	169.254.47.194/16	
cluster1-02		e0a true			
		cluster1-02_clus2	up/up	169.254.19.183/16	
cluster1-02		e0b true			

4 entries were displayed.

5. Ping the remote cluster LIFs:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::*> cluster ping-cluster -node cluster1-02
Host is cluster1-02
Getting addresses from network interface table...
Cluster cluster1-01_clus1 169.254.209.69 cluster1-01      e0a
Cluster cluster1-01_clus2 169.254.49.125 cluster1-01      e0b
Cluster cluster1-02_clus1 169.254.47.194 cluster1-02      e0a
Cluster cluster1-02_clus2 169.254.19.183 cluster1-02      e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

6. Verify that the auto-revert command is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert
```

Vserver	Logical Interface	Auto-revert
Cluster	cluster1-01_clus1	true
	cluster1-01_clus2	true
	cluster1-02_clus1	true
	cluster1-02_clus2	true

4 entries were displayed.

7. For ONTAP 9.8 and later, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

```
system switch ethernet log setup-password and system switch ethernet log enable-collection
```


Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

8. For ONTAP releases 9.5P16, 9.6P12, and 9.7P10 and later patch releases, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

`system cluster-switch log setup-password` and `system cluster-switch log enable-`

collection

Show example

```
cluster1::*> system cluster-switch log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

What's next?

[Install the NX-OS software.](#)

Install the NX-OS software

Follow this procedure to install the NX-OS software on the Nexus 9336C-FX2 cluster switch.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- [Cisco Ethernet switch page](#). Consult the switch compatibility table for the supported ONTAP and NX-OS versions.
- Appropriate software and upgrade guides available on the Cisco web site for the Cisco switch upgrade and downgrade procedures. See [Cisco Nexus 9000 Series Switches](#).

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the two Cisco switches are cs1 and cs2.
- The node names are cluster1-01, cluster1-02, cluster1-03, and cluster1-04.
- The cluster LIF names are cluster1-01_clus1, cluster1-01_clus2, cluster1-02_clus1, cluster1-02_clus2 , cluster1-03_clus1, cluster1-03_clus2, cluster1-04_clus1, and cluster1-04_clus2.
- The `cluster1::*>` prompt indicates the name of the cluster.

Install the software

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

Steps

1. Connect the cluster switch to the management network.
2. Use the ping command to verify connectivity to the server hosting the NX-OS software and the RCF.

Show example

This example verifies that the switch can reach the server at IP address 172.19.2.1:

```
cs2# ping 172.19.2.1
Pinging 172.19.2.1 with 0 bytes of data:

Reply From 172.19.2.1: icmp_seq = 0. time= 5910 usec.
```

3. Copy the NX-OS software and EPLD images to the Nexus 9336C-FX2 switch.

Show example

```
cs2# copy sftp: bootflash: vrf management
Enter source filename: /code/nxos.9.3.5.bin
Enter hostname for the sftp server: 172.19.2.1
Enter username: user1

Outbound-ReKey for 172.19.2.1:22
Inbound-ReKey for 172.19.2.1:22
user1@172.19.2.1's password:
sftp> progress
Progress meter enabled
sftp> get /code/nxos.9.3.5.bin /bootflash/nxos.9.3.5.bin
/code/nxos.9.3.5.bin 100% 1261MB 9.3MB/s 02:15
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.

cs2# copy sftp: bootflash: vrf management

Enter source filename: /code/n9000-epld.9.3.5.img
Enter hostname for the sftp server: 172.19.2.1
Enter username: user1

Outbound-ReKey for 172.19.2.1:22
Inbound-ReKey for 172.19.2.1:22
user1@172.19.2.1's password:
sftp> progress
Progress meter enabled
sftp> get /code/n9000-epld.9.3.5.img /bootflash/n9000-
epld.9.3.5.img
/code/n9000-epld.9.3.5.img 100% 161MB 9.5MB/s 00:16
sftp> exit
Copy complete, now saving to disk (please wait)...
Copy complete.
```

4. Verify the running version of the NX-OS software:

```
show version
```

Show example

```
cs2# show version
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2020, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their
own
licenses, such as open source. This software is provided "as is,"
and unless
otherwise stated, there is no warranty, express or implied,
including but not
limited to warranties of merchantability and fitness for a
particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0 or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
```

Software

```
BIOS: version 08.38
NXOS: version 9.3(4)
BIOS compile time: 05/29/2020
NXOS image file is: bootflash:///nxos.9.3.4.bin
NXOS compile time: 4/28/2020 21:00:00 [04/29/2020 02:28:31]
```

Hardware

```
cisco Nexus9000 C9336C-FX2 Chassis
Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
Processor Board ID FOC20291J6K

Device name: cs2
bootflash: 53298520 kB
Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 157524 usecs after Mon Nov  2 18:32:06 2020
Reason: Reset Requested by CLI command reload
System version: 9.3(4)
Service:
```

```
plugin
  Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

```
cs2#
```

5. Install the NX-OS image.

Installing the image file causes it to be loaded every time the switch is rebooted.

Show example

```
cs2# install all nxos bootflash:nxos.9.3.5.bin
```

Installer will perform compatibility check first. Please wait.

Installer is forced disruptive

Verifying image bootflash:/nxos.9.3.5.bin for boot variable "nxos".

[#####] 100% -- SUCCESS

Verifying image type.

[#####] 100% -- SUCCESS

Preparing "nxos" version info using image bootflash:/nxos.9.3.5.bin.

[#####] 100% -- SUCCESS

Preparing "bios" version info using image bootflash:/nxos.9.3.5.bin.

[#####] 100% -- SUCCESS

Performing module support checks.

[#####] 100% -- SUCCESS

Notifying services about system upgrade.

[#####] 100% -- SUCCESS

Compatibility check is done:

Module	bootable	Impact	Install-type	Reason
1	yes	disruptive	reset	default upgrade is not hitless

Images will be upgraded according to following table:

Module	Image	Running-Version(pri:alt	New-
Version		Upg-Required	
1	nxos	9.3(4)	9.3(5)
yes			
1	bios	v08.37(01/28/2020):v08.23(09/23/2015)	
v08.38(05/29/2020)		yes	

```
Switch will be reloaded for disruptive upgrade.

Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Performing runtime checks.
[#####] 100% -- SUCCESS

Setting boot variables.
[#####] 100% -- SUCCESS

Performing configuration copy.
[#####] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading
bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[#####] 100% -- SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

6. Verify the new version of NX-OS software after the switch has rebooted:

```
show version
```


Show example

```
cs2# show version
```

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2020, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their
own
licenses, such as open source.  This software is provided "as is,"
and unless
otherwise stated, there is no warranty, express or implied,
including but not
limited to warranties of merchantability and fitness for a
particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0  or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
```

Software

```
  BIOS: version 05.33
  NXOS: version 9.3(5)
  BIOS compile time:  09/08/2018
  NXOS image file is: bootflash:///nxos.9.3.5.bin
  NXOS compile time:  11/4/2018 21:00:00 [11/05/2018 06:11:06]
```

Hardware

```
  cisco Nexus9000 C9336C-FX2 Chassis
  Intel(R) Xeon(R) CPU E5-2403 v2 @ 1.80GHz with 8154432 kB of
memory.
  Processor Board ID FOC20291J6K

  Device name: cs2
  bootflash:  53298520 kB
  Kernel uptime is 0 day(s), 0 hour(s), 3 minute(s), 42 second(s)
```

```
Last reset at 277524 usecs after Mon Nov  2 22:45:12 2020
```

```
Reason: Reset due to upgrade
```

```
System version: 9.3(4)
```

```
Service:
```

```
plugin
```

```
Core Plugin, Ethernet Plugin
```

```
Active Package(s):
```

7. Upgrade the EPLD image and reboot the switch.

Show example



```
cs2# show version module 1 epld
```

EPLD Device	Version
MI FPGA	0x7
IO FPGA	0x17
MI FPGA2	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2
GEM FPGA	0x2

```
cs2# install epld bootflash:n9000-epld.9.3.5.img module 1
```

Compatibility check:

Module	Type	Upgradable	Impact	Reason
1	SUP	Yes	disruptive	Module Upgradable

Retrieving EPLD versions.... Please wait.

Images will be upgraded according to following table:

Module	Type	EPLD	Running-Version	New-Version	Upg-Required
1	SUP	MI FPGA	0x07	0x07	No
1	SUP	IO FPGA	0x17	0x19	Yes
1	SUP	MI FPGA2	0x02	0x02	No

The above modules require upgrade.

The switch will be reloaded at the end of the upgrade

Do you want to continue (y/n) ? [n] y

Proceeding to upgrade Modules.

Starting Module 1 EPLD Upgrade

Module 1 : IO FPGA [Programming] : 100.00% (64 of 64 sectors)

Module 1 EPLD upgrade is successful.

Module	Type	Upgrade-Result
1	SUP	Success

EPLDs upgraded.

Module 1 EPLD upgrade is successful.

8. After the switch reboot, log in again and verify that the new version of EPLD loaded successfully.

Show example

```
cs2# show version module 1 epld
```

EPLD	Device	Version
MI	FPGA	0x7
IO	FPGA	0x19
MI	FPGA2	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2
GEM	FPGA	0x2

9. Repeat steps 1 to 8 to install the NX-OS software on switch cs1.

What's next?

[Install the Reference Configuration File \(RCF\).](#)

Install the Reference Configuration File (RCF)

You can install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 switch for the first time. You can also use this procedure to upgrade your RCF version.

Before you begin, complete the procedure in [Prepare to install NX-OS and RCF](#).

For details of the available RCF configurations, see [Software install workflow](#).

Review requirements

What you'll need

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF file.
- A console connection to the switch, required when installing the RCF.

Suggested documentation

- [Cisco Ethernet switch page](#) Consult the switch compatibility table for the supported ONTAP and RCF versions. Note that there can be command dependencies between the command syntax in the RCF and that found in versions of NX-OS.
- [Cisco Nexus 3000 Series Switches](#). Refer to the appropriate software and upgrade guides available on the Cisco web site for complete documentation on the Cisco switch upgrade and downgrade procedures.

Install the RCF

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the two Cisco switches are cs1 and cs2.
- The node names are cluster1-01, cluster1-02, cluster1-03, and cluster1-04.
- The cluster LIF names are cluster1-01_clus1, cluster1-01_clus2, cluster1-02_clus1, cluster1-02_clus2 , cluster1-03_clus1, cluster1-03_clus2, cluster1-04_clus1, and cluster1-04_clus2.
- The `cluster1::*>` prompt indicates the name of the cluster.

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.



The command outputs might vary depending on different releases of ONTAP.

About this task

The procedure requires the use of both ONTAP commands and Cisco Nexus 9000 Series Switches commands; ONTAP commands are used unless otherwise indicated.

No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.



Before installing a new switch software version and RCFs, you must erase the switch settings and perform basic configuration. You must be connected to the switch using the serial console. This task resets the configuration of the management network.

Step 1: Prepare for the installation

1. Display the cluster ports on each node that are connected to the cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
-----
cluster1-01/cdp
           e0a    cs1                      Ethernet1/7      N9K-
C9336C
           e0d    cs2                      Ethernet1/7      N9K-
C9336C
cluster1-02/cdp
           e0a    cs1                      Ethernet1/8      N9K-
C9336C
           e0d    cs2                      Ethernet1/8      N9K-
C9336C
cluster1-03/cdp
           e0a    cs1                      Ethernet1/1/1    N9K-
C9336C
           e0b    cs2                      Ethernet1/1/1    N9K-
C9336C
cluster1-04/cdp
           e0a    cs1                      Ethernet1/1/2    N9K-
C9336C
           e0b    cs2                      Ethernet1/1/2    N9K-
C9336C
cluster1::*>
```

2. Check the administrative and operational status of each cluster port.

a. Verify that all the cluster ports are **up** with a healthy status:

```
network port show -role cluster
```

Show example

```
cluster1::*> network port show -role cluster
```

```
Node: cluster1-01
```

```
Ignore
```

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e0d	Cluster	Cluster		up	9000	auto/100000
healthy	false					

```
Node: cluster1-02
```

```
Ignore
```

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e0d	Cluster	Cluster		up	9000	auto/100000
healthy	false					

8 entries were displayed.

```
Node: cluster1-03
```

```
Ignore
```

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: cluster1-04

Ignore

Health	Health					Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----		----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

cluster1::*>

b. Verify that all the cluster interfaces (LIFs) are on the home port:

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

Show example

```
cluster1::*> network interface show -role cluster
```

	Logical	Status	Network	
Current	Current Is			
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	cluster1-01_clus1	up/up	169.254.3.4/23	
cluster1-01	e0a true			
	cluster1-01_clus2	up/up	169.254.3.5/23	
cluster1-01	e0d true			
	cluster1-02_clus1	up/up	169.254.3.8/23	
cluster1-02	e0a true			
	cluster1-02_clus2	up/up	169.254.3.9/23	
cluster1-02	e0d true			
	cluster1-03_clus1	up/up	169.254.1.3/23	
cluster1-03	e0a true			
	cluster1-03_clus2	up/up	169.254.1.1/23	
cluster1-03	e0b true			
	cluster1-04_clus1	up/up	169.254.1.6/23	
cluster1-04	e0a true			
	cluster1-04_clus2	up/up	169.254.1.7/23	
cluster1-04	e0b true			
8 entries were displayed.				
cluster1::*>				

c. Verify that the cluster displays information for both cluster switches:

```
system cluster-switch show -is-monitoring-enabled-operational true
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                Address
Model
-----
cs1                                     cluster-network     10.233.205.90
N9K-C9336C
    Serial Number: FOCXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
    9.3(5)
    Version Source: CDP

cs2                                     cluster-network     10.233.205.91
N9K-C9336C
    Serial Number: FOCXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
    9.3(5)
    Version Source: CDP
cluster1::*>
```

3. Disable auto-revert on the cluster LIFs.

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert false
```

Step 2: Configure ports

1. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes.

Show example

```
cs2(config)# interface eth1/1/1-2,eth1/7-8
cs2(config-if-range)# shutdown
```

2. Verify that the cluster LIFs have migrated to the ports hosted on cluster switch cs1. This might take a few seconds.

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

Show example

```
cluster1::*> network interface show -role cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	cluster1-01_clus1	up/up	169.254.3.4/23	
cluster1-01	e0a true			
	cluster1-01_clus2	up/up	169.254.3.5/23	
cluster1-01	e0a false			
	cluster1-02_clus1	up/up	169.254.3.8/23	
cluster1-02	e0a true			
	cluster1-02_clus2	up/up	169.254.3.9/23	
cluster1-02	e0a false			
	cluster1-03_clus1	up/up	169.254.1.3/23	
cluster1-03	e0a true			
	cluster1-03_clus2	up/up	169.254.1.1/23	
cluster1-03	e0a false			
	cluster1-04_clus1	up/up	169.254.1.6/23	
cluster1-04	e0a true			
	cluster1-04_clus2	up/up	169.254.1.7/23	
cluster1-04	e0a false			
8 entries were displayed.				
cluster1::*>				

3. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node           Health Eligibility  Epsilon
-----
cluster1-01    true   true        false
cluster1-02    true   true        false
cluster1-03    true   true         true
cluster1-04    true   true        false
4 entries were displayed.
cluster1::*>
```

4. If you have not already done so, save a copy of the current switch configuration by copying the output of the following command to a text file:

```
show running-config
```

5. Clean the configuration on switch cs2 and perform a basic setup.



When updating or applying a new RCF, you must erase the switch settings and perform basic configuration. You must be connected to the switch serial console port to set up the switch again.

- a. Clean the configuration:

Show example

```
(cs2)# write erase

Warning: This command will erase the startup-configuration.

Do you wish to proceed anyway? (y/n)  [n]  y
```

- b. Perform a reboot of the switch:

Show example

```
(cs2)# reload

Are you sure you would like to reset the system? (y/n) y
```

6. Copy the RCF to the bootflash of switch cs2 using one of the following transfer protocols: FTP, TFTP, SFTP, or SCP. For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows TFTP being used to copy an RCF to the bootflash on switch cs2:

```
cs2# copy tftp: bootflash: vrf management
Enter source filename: Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
Enter hostname for the tftp server: 172.22.201.50
Trying to connect to tftp server.....Connection to Server
Established.
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...
```

7. Apply the RCF previously downloaded to the bootflash.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

This example shows the RCF file `Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt` being installed on switch cs2:

```
cs2# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

8. Examine the banner output from the `show banner motd` command. You must read and follow these instructions to ensure the proper configuration and operation of the switch.

Show example

```
cs2# show banner motd

*****
*****
* NetApp Reference Configuration File (RCF)
*
* Switch      : Nexus N9K-C9336C-FX2
* Filename    : Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt
* Date       : 10-23-2020
* Version    : v1.6
*
* Port Usage:
* Ports 1- 3: Breakout mode (4x10G) Intra-Cluster Ports, int
e1/1/1-4, e1/2/1-4
, e1/3/1-4
* Ports 4- 6: Breakout mode (4x25G) Intra-Cluster/HA Ports, int
e1/4/1-4, e1/5/
1-4, e1/6/1-4
* Ports 7-34: 40/100GbE Intra-Cluster/HA Ports, int e1/7-34
* Ports 35-36: Intra-Cluster ISL Ports, int e1/35-36
*
* Dynamic breakout commands:
* 10G: interface breakout module 1 port <range> map 10g-4x
* 25G: interface breakout module 1 port <range> map 25g-4x
*
* Undo breakout commands and return interfaces to 40/100G
configuration in confi
g mode:
* no interface breakout module 1 port <range> map 10g-4x
* no interface breakout module 1 port <range> map 25g-4x
* interface Ethernet <interfaces taken out of breakout mode>
* inherit port-profile 40-100G
* priority-flow-control mode auto
* service-policy input HA
* exit
*
*****
*****
```

9. Verify that the RCF file is the correct newer version:

```
show running-config
```

When you check the output to verify you have the correct RCF, make sure that the following information is correct:

- The RCF banner
- The node and port settings
- Customizations

The output varies according to your site configuration. Check the port settings and refer to the release notes for any changes specific to the RCF that you have installed.

10. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file.

For more information on Cisco commands, see the appropriate guide in the [Cisco Nexus 9000 Series NX-OS Command Reference](#) guides.

Show example

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

11. Reboot switch cs2. You can ignore the “cluster ports down” events reported on the nodes while the switch reboots.

Show example

```
cs2# reload
This command will reboot the system. (y/n)? [n] y
```

12. Verify the health of cluster ports on the cluster.
 - a. Verify that e0d ports are up and healthy across all nodes in the cluster:

```
network port show -role cluster
```


Show example

```
cluster1::*> network port show -role cluster
```

Node: cluster1-01

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: cluster1-02

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

Node: cluster1-03

Ignore

						Speed (Mbps)
Health	Health					
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e0d	Cluster	Cluster		up	9000	auto/100000
healthy	false					

Node: cluster1-04

Ignore

Health	Health				Speed (Mbps)	
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status	Status					
-----	-----	-----	-----	-----	-----	-----
e0a	Cluster	Cluster		up	9000	auto/100000
healthy	false					
e0d	Cluster	Cluster		up	9000	auto/100000
healthy	false					

8 entries were displayed.

- b. Verify the switch health from the cluster (this might not show switch cs2, since LIFs are not homed on e0d).

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered	
Protocol	Port	Device (LLDP: ChassisID)	Interface
Platform			

cluster1-01/cdp			
	e0a	cs1	Ethernet1/7
N9K-C9336C			
	e0d	cs2	Ethernet1/7
N9K-C9336C			
cluster01-2/cdp			
	e0a	cs1	Ethernet1/8
N9K-C9336C			
	e0d	cs2	Ethernet1/8
N9K-C9336C			
cluster01-3/cdp			
	e0a	cs1	Ethernet1/1/1
N9K-C9336C			
	e0b	cs2	Ethernet1/1/1
N9K-C9336C			
cluster1-04/cdp			
	e0a	cs1	Ethernet1/1/2
N9K-C9336C			
	e0b	cs2	Ethernet1/1/2
N9K-C9336C			


```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
```

Switch	Type	Address
Model		

cs1	cluster-network	10.233.205.90
NX9-C9336C		
Serial Number: FOCXXXXXXGD		
Is Monitored: true		
Reason: None		
Software Version: Cisco Nexus Operating System (NX-OS)		
Software, Version		
9.3(5)		
Version Source: CDP		
cs2	cluster-network	10.233.205.91

```
NX9-C9336C
  Serial Number: FOCXXXXXXGS
    Is Monitored: true
      Reason: None
  Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
                9.3(5)
  Version Source: CDP

2 entries were displayed.
```

You might observe the following output on the cs1 switch console depending on the RCF version previously loaded on the switch:

```
2020 Nov 17 16:07:18 cs1 %$ VDC-1 %$ %STP-2-UNBLOCK_CONSIST_PORT:
Unblocking port port-channel1 on VLAN0092. Port consistency
restored.
2020 Nov 17 16:07:23 cs1 %$ VDC-1 %$ %STP-2-BLOCK_PVID_PEER:
Blocking port-channel1 on VLAN0001. Inconsistent peer vlan.
2020 Nov 17 16:07:23 cs1 %$ VDC-1 %$ %STP-2-BLOCK_PVID_LOCAL:
Blocking port-channel1 on VLAN0092. Inconsistent local vlan.
```

13. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes.

Show example

The following example uses the interface example output:

```
cs1(config)# interface eth1/1/1-2,eth1/7-8
cs1(config-if-range)# shutdown
```

14. Verify that the cluster LIFs have migrated to the ports hosted on switch cs2. This might take a few seconds.

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

Show example

```
cluster1::*> network interface show -role cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	cluster1-01_clus1	up/up	169.254.3.4/23	
cluster1-01	e0d	false		
	cluster1-01_clus2	up/up	169.254.3.5/23	
cluster1-01	e0d	true		
	cluster1-02_clus1	up/up	169.254.3.8/23	
cluster1-02	e0d	false		
	cluster1-02_clus2	up/up	169.254.3.9/23	
cluster1-02	e0d	true		
	cluster1-03_clus1	up/up	169.254.1.3/23	
cluster1-03	e0b	false		
	cluster1-03_clus2	up/up	169.254.1.1/23	
cluster1-03	e0b	true		
	cluster1-04_clus1	up/up	169.254.1.6/23	
cluster1-04	e0b	false		
	cluster1-04_clus2	up/up	169.254.1.7/23	
cluster1-04	e0b	true		
8 entries were displayed.				
cluster1::*>				

15. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node                Health    Eligibility    Epsilon
-----
cluster1-01         true     true           false
cluster1-02         true     true           false
cluster1-03         true     true           true
cluster1-04         true     true           false
4 entries were displayed.
cluster1::*>
```

16. Repeat steps 4 to 11 on switch cs1.
17. Enable auto-revert on the cluster LIFs.

Show example

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto
-revert True
```

18. Reboot switch cs1. You do this to trigger the cluster LIFs to revert to their home ports. You can ignore the “cluster ports down” events reported on the nodes while the switch reboots.

Show example

```
cs1# reload
This command will reboot the system. (y/n)? [n] y
```

Step 3: Verify the configuration

1. Verify that the switch ports connected to the cluster ports are **up**.

```
show interface brief
```

Show example

```
cs1# show interface brief | grep up
.
.
Eth1/1/1      1      eth  access up      none
10G(D)  --
Eth1/1/2      1      eth  access up      none
10G(D)  --
Eth1/7        1      eth  trunk  up      none
100G(D)  --
Eth1/8        1      eth  trunk  up      none
100G(D)  --
.
.
```

2. Verify that the expected nodes are still connected:

```
show cdp neighbors
```

Show example

```
cs1# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID          Local Intrfce  Hldtme Capability  Platform
Port ID
node1              Eth1/1        133      H                FAS2980
e0a
node2              Eth1/2        133      H                FAS2980
e0a
cs2                Eth1/35       175      R S I s          N9K-C9336C
Eth1/35
cs2                Eth1/36       175      R S I s          N9K-C9336C
Eth1/36

Total entries displayed: 4
```

3. Verify that the cluster nodes are in their correct cluster VLANs using the following commands:

```
show vlan brief
```

```
show interface trunk
```


Show example

```
cs1# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Pol, Eth1/1, Eth1/2, Eth1/3, Eth1/4, Eth1/5, Eth1/6, Eth1/7, Eth1/8, Eth1/35, Eth1/36, Eth1/9/1, Eth1/9/2, Eth1/9/3, Eth1/9/4, Eth1/10/1, Eth1/10/2, Eth1/10/3, Eth1/10/4
17	VLAN0017	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4, Eth1/5, Eth1/6, Eth1/7, Eth1/8, Eth1/9/1, Eth1/9/2, Eth1/9/3, Eth1/9/4, Eth1/10/1, Eth1/10/2, Eth1/10/3, Eth1/10/4
18	VLAN0018	active	Eth1/1, Eth1/2, Eth1/3, Eth1/4, Eth1/5, Eth1/6, Eth1/7, Eth1/8, Eth1/9/1, Eth1/9/2, Eth1/9/3, Eth1/9/4, Eth1/10/1, Eth1/10/2, Eth1/10/3, Eth1/10/4
31	VLAN0031	active	Eth1/11, Eth1/12, Eth1/13, Eth1/14, Eth1/15, Eth1/16, Eth1/17, Eth1/18, Eth1/19, Eth1/20, Eth1/21, Eth1/22
32	VLAN0032	active	Eth1/23, Eth1/24, Eth1/25

```

Eth1/28
Eth1/31
Eth1/34
33    VLAN0033          active  Eth1/11, Eth1/12,
Eth1/13
Eth1/16
Eth1/19
Eth1/22
34    VLAN0034          active  Eth1/23, Eth1/24,
Eth1/25
Eth1/28
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

Port	Native Vlan	Status	Port Channel
Eth1/1	1	trunking	--
Eth1/2	1	trunking	--
Eth1/3	1	trunking	--
Eth1/4	1	trunking	--
Eth1/5	1	trunking	--
Eth1/6	1	trunking	--
Eth1/7	1	trunking	--
Eth1/8	1	trunking	--
Eth1/9/1	1	trunking	--
Eth1/9/2	1	trunking	--
Eth1/9/3	1	trunking	--
Eth1/9/4	1	trunking	--
Eth1/10/1	1	trunking	--
Eth1/10/2	1	trunking	--
Eth1/10/3	1	trunking	--
Eth1/10/4	1	trunking	--
Eth1/11	33	trunking	--

Eth1/12	33	trunking	--
Eth1/13	33	trunking	--
Eth1/14	33	trunking	--
Eth1/15	33	trunking	--
Eth1/16	33	trunking	--
Eth1/17	33	trunking	--
Eth1/18	33	trunking	--
Eth1/19	33	trunking	--
Eth1/20	33	trunking	--
Eth1/21	33	trunking	--
Eth1/22	33	trunking	--
Eth1/23	34	trunking	--
Eth1/24	34	trunking	--
Eth1/25	34	trunking	--
Eth1/26	34	trunking	--
Eth1/27	34	trunking	--
Eth1/28	34	trunking	--
Eth1/29	34	trunking	--
Eth1/30	34	trunking	--
Eth1/31	34	trunking	--
Eth1/32	34	trunking	--
Eth1/33	34	trunking	--
Eth1/34	34	trunking	--
Eth1/35	1	trnk-bndl	Pol
Eth1/36	1	trnk-bndl	Pol
Pol	1	trunking	--

```

-----
Port                Vlans Allowed on Trunk
-----
Eth1/1              1,17-18
Eth1/2              1,17-18
Eth1/3              1,17-18
Eth1/4              1,17-18
Eth1/5              1,17-18
Eth1/6              1,17-18
Eth1/7              1,17-18
Eth1/8              1,17-18
Eth1/9/1            1,17-18
Eth1/9/2            1,17-18
Eth1/9/3            1,17-18
Eth1/9/4            1,17-18
Eth1/10/1           1,17-18
Eth1/10/2           1,17-18
Eth1/10/3           1,17-18
Eth1/10/4           1,17-18

```

Eth1/11	31, 33
Eth1/12	31, 33
Eth1/13	31, 33
Eth1/14	31, 33
Eth1/15	31, 33
Eth1/16	31, 33
Eth1/17	31, 33
Eth1/18	31, 33
Eth1/19	31, 33
Eth1/20	31, 33
Eth1/21	31, 33
Eth1/22	31, 33
Eth1/23	32, 34
Eth1/24	32, 34
Eth1/25	32, 34
Eth1/26	32, 34
Eth1/27	32, 34
Eth1/28	32, 34
Eth1/29	32, 34
Eth1/30	32, 34
Eth1/31	32, 34
Eth1/32	32, 34
Eth1/33	32, 34
Eth1/34	32, 34
Eth1/35	1
Eth1/36	1
Pol	1
..	
..	
..	
..	
..	



For specific port and VLAN usage details, refer to the banner and important notes section in your RCF.

4. Verify that the ISL between cs1 and cs2 is functional:

```
show port-channel summary
```

Show example

```
cs1# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met

-----
-----
Group Port-          Type      Protocol  Member Ports      Channel
-----
-----
1      Pol (SU)      Eth      LACP      Eth1/35 (P)      Eth1/36 (P)
cs1#
```

5. Verify that the cluster LIFs have reverted to their home port:

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

Show example

```
cluster1::*> network interface show -role cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	cluster1-01_clus1	up/up	169.254.3.4/23	
cluster1-01	e0d	true		
	cluster1-01_clus2	up/up	169.254.3.5/23	
cluster1-01	e0d	true		
	cluster1-02_clus1	up/up	169.254.3.8/23	
cluster1-02	e0d	true		
	cluster1-02_clus2	up/up	169.254.3.9/23	
cluster1-02	e0d	true		
	cluster1-03_clus1	up/up	169.254.1.3/23	
cluster1-03	e0b	true		
	cluster1-03_clus2	up/up	169.254.1.1/23	
cluster1-03	e0b	true		
	cluster1-04_clus1	up/up	169.254.1.6/23	
cluster1-04	e0b	true		
	cluster1-04_clus2	up/up	169.254.1.7/23	
cluster1-04	e0b	true		
8 entries were displayed.				
cluster1::*>				

6. Verify that the cluster is healthy:

```
cluster show
```

Show example

```
cluster1::*> cluster show
Node           Health Eligibility Epsilon
-----
cluster1-01    true   true      false
cluster1-02    true   true      false
cluster1-03    true   true      true
cluster1-04    true   true      false
4 entries were displayed.
cluster1::*>
```

7. Ping the remote cluster interfaces to verify connectivity:

```
cluster ping-cluster -node local
```

Show example

```
cluster1::*> cluster ping-cluster -node local
Host is cluster1-03
Getting addresses from network interface table...
Cluster cluster1-03_clus1 169.254.1.3 cluster1-03 e0a
Cluster cluster1-03_clus2 169.254.1.1 cluster1-03 e0b
Cluster cluster1-04_clus1 169.254.1.6 cluster1-04 e0a
Cluster cluster1-04_clus2 169.254.1.7 cluster1-04 e0b
Cluster cluster1-01_clus1 169.254.3.4 cluster1-01 e0a
Cluster cluster1-01_clus2 169.254.3.5 cluster1-01 e0d
Cluster cluster1-02_clus1 169.254.3.8 cluster1-02 e0a
Cluster cluster1-02_clus2 169.254.3.9 cluster1-02 e0d
Local = 169.254.1.3 169.254.1.1
Remote = 169.254.1.6 169.254.1.7 169.254.3.4 169.254.3.5 169.254.3.8
169.254.3.9
Cluster Vserver Id = 4294967293
Ping status:
.....
Basic connectivity succeeds on 12 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 12 path(s):
    Local 169.254.1.3 to Remote 169.254.1.6
    Local 169.254.1.3 to Remote 169.254.1.7
    Local 169.254.1.3 to Remote 169.254.3.4
    Local 169.254.1.3 to Remote 169.254.3.5
    Local 169.254.1.3 to Remote 169.254.3.8
    Local 169.254.1.3 to Remote 169.254.3.9
    Local 169.254.1.1 to Remote 169.254.1.6
    Local 169.254.1.1 to Remote 169.254.1.7
    Local 169.254.1.1 to Remote 169.254.3.4
    Local 169.254.1.1 to Remote 169.254.3.5
    Local 169.254.1.1 to Remote 169.254.3.8
    Local 169.254.1.1 to Remote 169.254.3.9
Larger than PMTU communication succeeds on 12 path(s)
RPC status:
6 paths up, 0 paths down (tcp check)
6 paths up, 0 paths down (udp check)
```

Enable SSH on Cisco 9336C-FX2 cluster switches

If you are using the Cluster Switch Health Monitor (CSHM) and log collection features, you must generate the SSH keys and then enable SSH on the cluster switches.

Steps

1. Verify that SSH is disabled:

```
show ip ssh
```

Show example

```
(switch)# show ip ssh
```

```
SSH Configuration
```

```
Administrative Mode: ..... Disabled
SSH Port: ..... 22
Protocol Level: ..... Version 2
SSH Sessions Currently Active: ..... 0
Max SSH Sessions Allowed: ..... 5
SSH Timeout (mins): ..... 5
Keys Present: ..... DSA(1024) RSA(1024)
ECDSA(521)
Key Generation In Progress: ..... None
SSH Public Key Authentication Mode: ..... Disabled
SCP server Administrative Mode: ..... Disabled
```

2. Generate the SSH keys:

```
crypto key generate
```

Show example

```
(switch)# config

(switch) (Config)# crypto key generate rsa

Do you want to overwrite the existing RSA keys? (y/n): y

(switch) (Config)# crypto key generate dsa

Do you want to overwrite the existing DSA keys? (y/n): y

(switch) (Config)# crypto key generate ecdsa 521

Do you want to overwrite the existing ECDSA keys? (y/n): y

(switch) (Config)# aaa authorization commands "noCmdAuthList" none
(switch) (Config)# exit
(switch)# ip ssh server enable
(switch)# ip scp server enable
(switch)# ip ssh pubkey-auth
(switch)# write mem

This operation may take a few minutes.
Management interfaces will not be available during this time.
Are you sure you want to save? (y/n) y

Config file 'startup-config' created successfully.

Configuration Saved!
```

3. Reboot the switch:

```
reload
```

4. Verify that SSH is enabled:

```
show ip ssh
```

Show example

```
(switch)# show ip ssh
```

SSH Configuration

```
Administrative Mode: ..... Enabled
SSH Port: ..... 22
Protocol Level: ..... Version 2
SSH Sessions Currently Active: ..... 0
Max SSH Sessions Allowed: ..... 5
SSH Timeout (mins): ..... 5
Keys Present: ..... DSA(1024) RSA(1024)
ECDSA(521)
Key Generation In Progress: ..... None
SSH Public Key Authentication Mode: ..... Enabled
SCP server Administrative Mode: ..... Enabled
```

What's next?

[Enable log collection.](#)

Ethernet Switch Health Monitoring log collection

You can use the log collection feature to collect switch-related log files in ONTAP. The Ethernet switch health monitor (CSHM) is responsible for ensuring the operational health of Cluster and Storage network switches and collecting switch logs for debugging purposes. This procedure guides you through the process of setting up and starting the collection of detailed **Support** logs from the switch and starts an hourly collection of **Periodic** data that is collected by AutoSupport.

Before you begin

- Verify that you have set up your environment using the 9336C-FX2 cluster switch **CLI**.
- Switch health monitoring must be enabled for the switch. Verify this by ensuring the `Is Monitored:` field is set to **true** in the output of the `system switch ethernet show` command.

Steps

1. Create a password for the Ethernet switch health monitor log collection feature:

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
Would you like to specify a user other than admin for log
collection? {y|n}: n

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

2. To start log collection, run the following command, replacing **DEVICE** with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

Wait for 10 minutes and then check that the log collection completes:

```
system switch ethernet log show
```



If any of these commands return an error or if the log collection does not complete, contact NetApp support.

Troubleshooting

If you encounter any of the following error statuses reported by the log collection feature (visible in the output of `system switch ethernet log show`), try the corresponding debug steps:

Log collection error status	Resolution
RSA keys not present	Regenerate ONTAP SSH keys. Contact NetApp support.
switch password error	Verify credentials, test SSH connectivity, and regenerate ONTAP SSH keys. Review the switch documentation or contact NetApp support for instructions.
ECDSA keys not present for FIPS	If FIPS mode is enabled, ECDSA keys need to be generated on the switch before retrying.
pre-existing log found	Remove the previous log collection file on the switch.

switch dump log error	Ensure the switch user has log collection permissions. Refer to the prerequisites above.
------------------------------	--

Configure SNMPv3

Follow this procedure to configure SNMPv3, which supports Ethernet switch health monitoring (CSHM).

About this task

The following commands configure an SNMPv3 username on Cisco 9336C-FX2 switches:

- For **no authentication**:

```
snmp-server user SNMPv3_USER NoAuth
```
- For **MD5/SHA authentication**:

```
snmp-server user SNMPv3_USER auth [md5|sha] AUTH-PASSWORD
```
- For **MD5/SHA authentication with AES/DES encryption**:

```
snmp-server user SNMPv3_USER AuthEncrypt auth [md5|sha] AUTH-PASSWORD priv  
aes-128 PRIV-PASSWORD
```

The following command configures an SNMPv3 username on the ONTAP side:

```
cluster1::*> security login create -user-or-group-name SNMPv3_USER -application  
snmp -authentication-method usm -remote-switch-ipaddress ADDRESS
```

The following command establishes the SNMPv3 username with CSHM:

```
cluster1::*> system switch ethernet modify -device DEVICE -snmp-version SNMPv3  
-community-or-username SNMPv3_USER
```

Steps

1. Set up the SNMPv3 user on the switch to use authentication and encryption:

```
show snmp user
```

Show example

```
(sw1) (Config)# snmp-server user SNMPv3User auth md5 <auth_password>
priv aes-128 <priv_password>

(sw1) (Config)# show snmp user

-----
-----
                        SNMP USERS
-----
-----

User                Auth                Priv(enforce)    Groups
acl_filter
-----
-----
admin               md5                des(no)          network-admin
SNMPv3User          md5                aes-128(no)      network-operator
-----
-----

      NOTIFICATION TARGET USERS (configured for sending V3 Inform)
-----
-----

User                Auth                Priv
-----
-----

(sw1) (Config)#
```

2. Set up the SNMPv3 user on the ONTAP side:

```
security login create -user-or-group-name <username> -application snmp
-authentication-method usm -remote-switch-ipaddress 10.231.80.212
```

Show example

```
cluster1::*> system switch ethernet modify -device "sw1
(b8:59:9f:09:7c:22)" -is-monitoring-enabled-admin true

cluster1::*> security login create -user-or-group-name <username>
-application snmp -authentication-method usm -remote-switch
-ipaddress 10.231.80.212

Enter the authoritative entity's EngineID [remote EngineID]:

Which authentication protocol do you want to choose (none, md5, sha,
sha2-256)
[none]: md5

Enter the authentication protocol password (minimum 8 characters
long):

Enter the authentication protocol password again:

Which privacy protocol do you want to choose (none, des, aes128)
[none]: aes128

Enter privacy protocol password (minimum 8 characters long):
Enter privacy protocol password again:
```

3. Configure CSHM to monitor with the new SNMPv3 user:

```
system switch ethernet show-all -device "sw1" -instance
```


Show example

```
cluster1::*> system switch ethernet show-all -device "sw1" -instance

Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv2c
Is Discovered: true
SNMPv2c Community String or SNMPv3 Username: cshml!
Model Number: N9K-C9336C-FX2
Switch Network: cluster-network
Software Version: Cisco Nexus
Operating System (NX-OS) Software, Version 9.3(7)
Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored ?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
cluster1::*> system switch ethernet modify -device "sw1" -snmp
-version SNMPv3 -community-or-username <username>
cluster1::*>
```

4. Verify that the serial number to be queried with the newly created SNMPv3 user is the same as detailed in the previous step after the CSHM polling period has completed.

```
system switch ethernet polling-interval show
```

Show example

```
cluster1::*> system switch ethernet polling-interval show
Polling Interval (in minutes): 5

cluster1::*> system switch ethernet show-all -device "sw1" -instance

Device Name: sw1
IP Address: 10.231.80.212
SNMP Version: SNMPv3
Is Discovered: true
SNMPv2c Community String or SNMPv3 Username: SNMPv3User
Model Number: N9K-C9336C-FX2
Switch Network: cluster-network
Software Version: Cisco Nexus
Operating System (NX-OS) Software, Version 9.3(7)
Reason For Not Monitoring: None <---- displays
when SNMP settings are valid
Source Of Switch Version: CDP/ISDP
Is Monitored?: true
Serial Number of the Device: QTFCU3826001C
RCF Version: v1.8X2 for

Cluster/HA/RDMA

cluster1::*>
```

Migrate switches

Migrate from a NetApp CN1610 cluster switch to a Cisco 9336C-FX2 cluster switch

You can migrate NetApp CN1610 cluster switches for an ONTAP cluster to Cisco 9336C-FX2 cluster switches. This is a nondisruptive procedure.

Review requirements

You must be aware of certain configuration information, port connections and cabling requirements when you are replacing NetApp CN1610 cluster switches with Cisco 9336C-FX2 cluster switches.

Supported switches

The following cluster switches are supported:

- NetApp CN1610
- Cisco 9336C-FX2

For details of supported ports and their configurations, see the [Hardware Universe](#).

What you'll need

Verify that your configuration meets the following requirements:

- The existing cluster is correctly set up and functioning.
- All cluster ports are in the **up** state to ensure nondisruptive operations.
- The Cisco 9336C-FX2 cluster switches are configured and operating under the correct version of NX-OS installed with the reference configuration file (RCF) applied.
- The existing cluster network configuration has the following:
 - A redundant and fully functional NetApp cluster using NetApp CN1610 switches.
 - Management connectivity and console access to both the NetApp CN1610 switches and the new switches.
 - All cluster LIFs in the up state with the cluster LIFs are on their home ports.
- Some of the ports are configured on Cisco 9336C-FX2 switches to run at 40GbE or 100GbE.
- You have planned, migrated, and documented 40GbE and 100GbE connectivity from nodes to Cisco 9336C-FX2 cluster switches.

Migrate the switches

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The existing CN1610 cluster switches are *C1* and *C2*.
- The new 9336C-FX2 cluster switches are *cs1* and *cs2*.
- The nodes are *node1* and *node2*.
- The cluster LIFs are *node1_clus1* and *node1_clus2* on node 1, and *node2_clus1* and *node2_clus2* on node 2 respectively.
- The `cluster1: :*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are *e3a* and *e3b*.

About this task

This procedure covers the following scenario:

- Switch C2 is replaced by switch cs2 first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - The cabling between the nodes and C2 is then disconnected from C2 and reconnected to cs2.
- Switch C1 is replaced by switch cs1.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - The cabling between the nodes and C1 is then disconnected from C1 and reconnected to cs1.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where *x* is the duration of the maintenance window in hours.

2. Change the privilege level to advanced, entering *y* when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (*>) appears.

3. Disable auto-revert on the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

Step 2: Configure ports and cabling

1. Determine the administrative or operational status for each cluster interface.

Each port should display up for `Link` and `healthy` for `Health Status`.

- a. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
e3a	Cluster	Cluster		up	9000
healthy	false				auto/100000
e3b	Cluster	Cluster		up	9000
healthy	false				auto/100000

Node: node2

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
-----	-----				
e3a	Cluster	Cluster		up	9000
healthy	false				auto/100000
e3b	Cluster	Cluster		up	9000
healthy	false				auto/100000

b. Display information about the LIFs and their designated home nodes:

```
network interface show -vserver Cluster
```

Each LIF should display up/up for Status Admin/Oper and true for Is Home.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e3a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e3b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e3a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e3b	true			

2. The cluster ports on each node are connected to existing cluster switches in the following way (from the nodes' perspective) using the command:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered		
Protocol	Port	Device (LLDP: ChassisID)	Interface	
Platform				

node1	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/1	-
	e3b	C2 (6a:ad:4f:98:4c:a4)	0/1	-
node2	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/2	-
	e3b	C2 (6a:ad:4f:98:4c:a4)	0/2	-

3. The cluster ports and switches are connected in the following way (from the switches' perspective) using the command:

```
show cdp neighbors
```

Show example

C1# **show cdp neighbors**

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3a	Eth1/1	124	H	AFF-A400
node2 e3a	Eth1/2	124	H	AFF-A400
C2 0/13	0/13	179	S I s	CN1610
C2 0/14	0/14	175	S I s	CN1610
C2 0/15	0/15	179	S I s	CN1610
C2 0/16	0/16	175	S I s	CN1610

C2# **show cdp neighbors**

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e3b	Eth1/1	124	H	AFF-A400
node2 e3b	Eth1/2	124	H	AFF-A400
C1 0/13	0/13	175	S I s	CN1610
C1 0/14	0/14	175	S I s	CN1610
C1 0/15	0/15	175	S I s	CN1610
C1 0/16	0/16	175	S I s	CN1610

4. Verify that the cluster network has full connectivity using the command:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::*> cluster ping-cluster -node node2

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1      e3a
Cluster node1_clus2 169.254.49.125 node1      e3b
Cluster node2_clus1 169.254.47.194 node2      e3a
Cluster node2_clus2 169.254.19.183 node2      e3b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

5. On switch C2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(C2) # configure
(C2) (Config) # interface 0/1-0/12
(C2) (Interface 0/1-0/12) # shutdown
(C2) (Interface 0/1-0/12) # exit
(C2) (Config) # exit
```

6. Move the node cluster ports from the old switch C2 to the new switch cs2, using appropriate cabling supported by Cisco 9336C-FX2.
7. Display the network port attributes:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	-----
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

8. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node1	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/1	
CN1610				
	e3b	cs2 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
C9336C-FX2				
node2	/cdp			
	e3a	C1 (6a:ad:4f:98:3b:3f)	0/2	
CN1610				
	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-
C9336C-FX2				

9. On switch cs2, verify that all node cluster ports are up:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interfac	Admin/Oper	Address/Mask	Node
Port	Home			
Cluster				
	node1_clus1	up/up	169.254.3.4/16	node1
e0b	false			
	node1_clus2	up/up	169.254.3.5/16	node1
e0b	true			
	node2_clus1	up/up	169.254.3.8/16	node2
e0b	false			
	node2_clus2	up/up	169.254.3.9/16	node2
e0b	true			

10. On switch C1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs.

```
(C1) # configure
(C1) (Config) # interface 0/1-0/12
(C1) (Interface 0/1-0/12) # shutdown
(C1) (Interface 0/1-0/12) # exit
(C1) (Config) # exit
```

11. Move the node cluster ports from the old switch C1 to the new switch cs1, using appropriate cabling supported by Cisco 9336C-FX2.
12. Verify the final configuration of the cluster:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

Node: node2

Ignore

						Speed (Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	-----	-----	-----	
-----	-----						
e3a	Cluster	Cluster		up	9000	auto/100000	
healthy	false						
e3b	Cluster	Cluster		up	9000	auto/100000	
healthy	false						

13. The cluster ports on each node are now connected to cluster switches in the following way, from the nodes' perspective:

```
network device-discovery show -protocol
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	

node1	/cdp			
	e3a	cs1 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
C9336C-FX2				
	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-
C9336C-FX2				
node2	/cdp			
	e3a	cs1 (b8:ce:f6:19:1a:7e)	Ethernet1/1/1	N9K-
C9336C-FX2				
	e3b	cs2 (b8:ce:f6:19:1b:96)	Ethernet1/1/2	N9K-
C9336C-FX2				

14. On switches cs1 and cs2, verify that all node cluster ports are up:

```
network port show -ipSpace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	----	----	-----	
-----	-----						
e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

Node: node2

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	----	----	-----	
-----	-----						
e0a	Cluster	Cluster		up	9000	auto/10000	
healthy	false						
e0b	Cluster	Cluster		up	9000	auto/10000	
healthy	false						

15. Verify that both nodes each have one connection to each switch:

```
network device-discovery show -protocol
```

Show example

The following example shows the appropriate results for both switches:

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered		
Protocol	Port	Device	(LLDP: ChassisID)	Interface
Platform				

node1	/cdp			
	e0a	cs1	(b8:ce:f6:19:1b:42)	Ethernet1/1/1 N9K-
C9336C-FX2				
	e0b	cs2	(b8:ce:f6:19:1b:96)	Ethernet1/1/2 N9K-
C9336C-FX2				
node2	/cdp			
	e0a	cs1	(b8:ce:f6:19:1b:42)	Ethernet1/1/1 N9K-
C9336C-FX2				
	e0b	cs2	(b8:ce:f6:19:1b:96)	Ethernet1/1/2 N9K-
C9336C-FX2				

Step 3: Complete the procedure

1. Enable auto-revert on the cluster LIFs:

```
cluster1::*> network interface modify -vserver Cluster -lif * -auto-revert true
```

2. Verify that all cluster network LIFs are back on their home ports:

```
network interface show
```


Show example

```
cluster1::*> network interface show -vserver Cluster
```

		Logical	Status	Network	Current
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask		Node
Port	Home				

Cluster					
		node1_clus1	up/up	169.254.209.69/16	node1
e3a	true				
		node1_clus2	up/up	169.254.49.125/16	node1
e3b	true				
		node2_clus1	up/up	169.254.47.194/16	node2
e3a	true				
		node2_clus2	up/up	169.254.19.183/16	node2
e3b	true				

3. To set up log collection, run the following command for each switch. You are prompted to enter the switch name, username, and password for log collection.

```
system switch ethernet log setup-password
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>
```

4. To start log collection, run the following command, replacing **DEVICE** with the switch used in the previous command. This starts both types of log collection: the detailed **Support** logs and an hourly collection of **Periodic** data.

```
system switch ethernet log modify -device <switch-name> -log-request true
```

Show example

```
cluster1::*> system switch ethernet log modify -device cs1 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*> system switch ethernet log modify -device cs2 -log  
-request true
```

```
Do you want to modify the cluster switch log collection  
configuration? {y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*>
```

Wait for 10 minutes and then check that the log collection was successful using the command:

```
system switch ethernet log show
```



If any of these commands return an error, contact NetApp support.

5. Change the privilege level back to admin:

```
set -privilege admin
```

6. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Migrate from an older Cisco switch to a Cisco Nexus 9336C-FX2 cluster switch

You can perform a nondisruptive migration from an older Cisco cluster switch to a Cisco Nexus 9336C-FX2 cluster network switch.

Review requirements

Ensure that:

- Some of the ports on Nexus 9336C-FX2 switches are configured to run at 10GbE or 40GbE.
- The 10GbE and 40GbE connectivity from nodes to Nexus 9336C-FX2 cluster switches have been planned, migrated, and documented.
- The cluster is fully functioning (there should be no errors in the logs or similar issues).

- Initial customization of the Cisco Nexus 9336C-FX2 switches is complete, so that:
 - 9336C-FX2 switches are running the latest recommended version of software.
 - Reference Configuration Files (RCFs) have been applied to the switches.
 - Any site customization, such as DNS, NTP, SMTP, SNMP, and SSH, are configured on the new switches.
- You have access to the switch compatibility table on the [Cisco Ethernet Switches](#) page for the supported ONTAP, NX-OS, and RCF versions.
- You have reviewed the appropriate software and upgrade guides available on the Cisco web site for the Cisco switch upgrade and downgrade procedures at [Cisco Nexus 9000 Series Switches Support](#) page.



If you are changing the port speed of the e0a and e1a cluster ports on AFF A800 or AFF C800 systems, you might observe malformed packets being received after the speed conversion. See [Bug 1570339](#) and the Knowledge Base article [CRC errors on T6 ports after converting from 40GbE to 100GbE](#) for guidance.

Migrate the switches

About the examples

The examples in this procedure use two nodes. These nodes use two 10GbE cluster interconnect ports e0a and e0b. See the [Hardware Universe](#) to verify the correct cluster ports on your platforms.




The command outputs might vary depending on the different releases of ONTAP.

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing two Cisco switches are **cs1** and **cs2**
- The new Nexus 9336C-FX2 cluster switches are **cs1-new** and **cs2-new**.
- The node names are **node1** and **node2**.
- The cluster LIF names are **node1_clus1** and **node1_clus2** for node 1, and **node2_clus1** and **node2_clus2** for node 2.
- The **cluster1::>*** prompt indicates the name of the cluster.

During this procedure, refer to the following example:



About this task

The procedure requires the use of both ONTAP commands and [Nexus 9000 Series Switches](#) commands;

ONTAP commands are used, unless otherwise indicated.

This procedure covers the following scenario:

- Switch cs2 is replaced by switch cs2-new first.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and cs2 are then disconnected from cs2 and reconnected to cs2-new.
- Switch cs1 is replaced by switch cs1-new.
 - Shut down the ports to the cluster nodes. All ports must be shut down simultaneously to avoid cluster instability.
 - Cabling between the nodes and cs1 are then disconnected from cs1 and reconnected to cs1-new.



No operational inter-switch link (ISL) is needed during this procedure. This is by design because RCF version changes can affect ISL connectivity temporarily. To ensure non-disruptive cluster operations, the following procedure migrates all of the cluster LIFs to the operational partner switch while performing the steps on the target switch.

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:
`message: system node autosupport invoke -node * -type all -message MAINT=xh`

where *x* is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Change the privilege level to advanced, entering **y** when prompted to continue:

```
set -privilege advanced
```

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

Step 2: Configure ports and cabling

1. On the new switches, confirm that the ISL is cabled and healthy between the switches cs1-new and cs2-new:

```
show port-channel summary
```

Show example

```
cs1-new# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met

-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)        Eth       LACP      Eth1/35(P)  Eth1/36(P)

cs2-new# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met

-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)        Eth       LACP      Eth1/35(P)  Eth1/36(P)
```

2. Display the cluster ports on each node that are connected to the existing cluster switches:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/	Local	Discovered	
Protocol	Port	Device (LLDP: ChassisID)	Interface
Platform			

node1	/cdp		
	e0a	cs1	Ethernet1/1 N5K-
C5596UP			
	e0b	cs2	Ethernet1/2 N5K-
C5596UP			
node2	/cdp		
	e0a	cs1	Ethernet1/1 N5K-
C5596UP			
	e0b	cs2	Ethernet1/2 N5K-
C5596UP			

3. Determine the administrative or operational status for each cluster port.

a. Verify that all the cluster ports are up with a healthy status:

```
network port show -ipSPACE Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
e0a	Cluster	Cluster		up	9000
healthy	false				auto/10000
e0b	Cluster	Cluster		up	9000
healthy	false				auto/10000

Node: node2

Ignore

Health	Health				Speed (Mbps)
Port	IPspace	Broadcast	Domain	Link	MTU
Status	Status				Admin/Oper
-----	-----	-----	----	----	-----
-----	-----				
e0a	Cluster	Cluster		up	9000
healthy	false				auto/10000
e0b	Cluster	Cluster		up	9000
healthy	false				auto/10000

- b. Verify that all the cluster interfaces (LIFs) are on their home ports:

```
network interface show -vserver Cluster
```


Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
e0a	node1_clus1	up/up	169.254.209.69/16	node1
e0b	true			
e0a	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
e0a	node2_clus1	up/up	169.254.47.194/16	node2
e0b	true			
e0a	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

- c. Verify that the cluster displays information for both cluster switches:

```
system cluster-switch show -is-monitoring-enabled-operational true
```

Show example

```
cluster1::*> system cluster-switch show -is-monitoring-enabled
-operational true
Switch                                     Type                               Address
Model
-----
cs1                                       cluster-network                   10.233.205.92
N5K-C5596UP
    Serial Number: FOXXXXXXXXGS
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
    9.3(4)
    Version Source: CDP

cs2                                       cluster-network                   10.233.205.93
N5K-C5596UP
    Serial Number: FOXXXXXXXXGD
    Is Monitored: true
    Reason: None
    Software Version: Cisco Nexus Operating System (NX-OS)
Software, Version
    9.3(4)
    Version Source: CDP
```

4. Disable auto-revert on the cluster LIFs.

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. On cluster switch cs2, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs:

```
cs2(config)# interface eth1/1-1/2
cs2(config-if-range)# shutdown
```

6. Verify that the cluster LIFs have failed over to the ports hosted on cluster switch cs1. This might take a few seconds.

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.3.4/16	node1
e0a	true			
	node1_clus2	up/up	169.254.3.5/16	node1
e0a	false			
	node2_clus1	up/up	169.254.3.8/16	node2
e0a	true			
	node2_clus2	up/up	169.254.3.9/16	node2
e0a	false			

7. Verify that the cluster is healthy:

```
cluster show
```


Show example

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
-----	-----	-----	-----
node1	true	true	false
node2	true	true	false

8. Move all cluster node connection cables from the old cs2 switch to the new cs2-new switch.

Cluster node connection cables moved to the cs2-new switch



9. Confirm the health of the network connections moved to cs2-new:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status	Speed(Mbps)	Health
e0a	Cluster	Cluster		up	9000	auto/10000			
healthy	false								
e0b	Cluster	Cluster		up	9000	auto/10000			
healthy	false								

```
Node: node2
```

```
Ignore
```

Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status	Speed(Mbps)	Health
e0a	Cluster	Cluster		up	9000	auto/10000			
healthy	false								
e0b	Cluster	Cluster		up	9000	auto/10000			
healthy	false								

All cluster ports that were moved should be up.

10. Check neighbor information on the cluster ports:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
node1	/cdp			
	e0a	cs1	Ethernet1/1	N5K-
C5596UP				
	e0b	cs2-new	Ethernet1/1/1	N9K-
C9336C-FX2				
node2	/cdp			
	e0a	cs1	Ethernet1/2	N5K-
C5596UP				
	e0b	cs2-new	Ethernet1/1/2	N9K-
C9336C-FX2				

Verify that the moved cluster ports see the cs2-new switch as the neighbor.

11. Confirm the switch port connections from switch cs2-new's perspective:

```
cs2-new# show interface brief
cs2-new# show cdp neighbors
```

12. On cluster switch cs1, shut down the ports connected to the cluster ports of the nodes in order to fail over the cluster LIFs. The following example uses the interface example output from step 7.

```
cs1(config)# interface eth1/1-1/2
cs1(config-if-range)# shutdown
```

All cluster LIFs will move to the cs2-new switch.

13. Verify that the cluster LIFs have failed over to the ports hosted on switch cs2-new. This might take a few seconds:

```
network interface show -vserver Cluster
```

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interfac	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.3.4/16	node1
e0b	false			
	node1_clus2	up/up	169.254.3.5/16	node1
e0b	true			
	node2_clus1	up/up	169.254.3.8/16	node2
e0b	false			
	node2_clus2	up/up	169.254.3.9/16	node2
e0b	true			

14. Verify that the cluster is healthy:

```
cluster show
```


Show example

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
-----	-----	-----	-----
node1	true	true	false
node2	true	true	false

15. Move the cluster node connection cables from cs1 to the new cs1-new switch.

Cluster node connection cables moved to the cs1-new switch



16. Confirm the health of the network connections moved to cs1-new:

```
network port show -ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status	Speed(Mbps)	Health
e0a	Cluster	Cluster		up	9000	auto/10000			
healthy	false								
e0b	Cluster	Cluster		up	9000	auto/10000			
healthy	false								

```
Node: node2
```

```
Ignore
```

Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status	Speed(Mbps)	Health
e0a	Cluster	Cluster		up	9000	auto/10000			
healthy	false								
e0b	Cluster	Cluster		up	9000	auto/10000			
healthy	false								

All cluster ports that were moved should be up.

17. Check neighbor information on the cluster ports:

```
network device-discovery show
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface
Platform
-----
node1      /cdp
           e0a    cs1-new                  Ethernet1/1/1  N9K-
C9336C-FX2
           e0b    cs2-new                  Ethernet1/1/2  N9K-
C9336C-FX2

node2      /cdp
           e0a    cs1-new                  Ethernet1/1/1  N9K-
C9336C-FX2
           e0b    cs2-new                  Ethernet1/1/2  N9K-
C9336C-FX2
```

Verify that the moved cluster ports see the cs1-new switch as the neighbor.

18. Confirm the switch port connections from switch cs1-new's perspective:

```
cs1-new# show interface brief
cs1-new# show cdp neighbors
```

19. Verify that the ISL between cs1-new and cs2-new is still operational:

```
show port-channel summary
```


Show example

```
cs1-new# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met
```

```
-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)       Eth      LACP      Eth1/35(P)  Eth1/36(P)
```

```
cs2-new# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met
```

```
-----
-----
Group Port-          Type      Protocol  Member Ports
Channel
-----
-----
1      Po1(SU)       Eth      LACP      Eth1/35(P)  Eth1/36(P)
```

Step 3: Verify the configuration

1. Enable auto-revert on the cluster LIFs.

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

2. Verify that the cluster LIFs have reverted to their home ports (this might take a minute):

```
network interface show -vserver Cluster
```

If the cluster LIFs have not reverted to their home port, manually revert them:

```
network interface revert -vserver Cluster -lif *
```

3. Verify that the cluster is healthy:

```
cluster show
```

4. Verify the connectivity of the remote cluster interfaces:

ONTAP 9.9.1 and later

You can use the `network interface check cluster-connectivity` command to start an accessibility check for cluster connectivity and then display the details:

```
network interface check cluster-connectivity start and network interface check cluster-connectivity show
```

```
cluster1::*> network interface check cluster-connectivity start
```

NOTE: Wait for a number of seconds before running the show command to display the details.

```
cluster1::*> network interface check cluster-connectivity show
```

				Source	Destination
Packet					
Node	Date			LIF	LIF
Loss					

node1					
	3/5/2022	19:21:18	-06:00	node1_clus2	node2_clus1
none					
	3/5/2022	19:21:20	-06:00	node1_clus2	node2_clus2
none					
node2					
	3/5/2022	19:21:18	-06:00	node2_clus2	node1_clus1
none					
	3/5/2022	19:21:20	-06:00	node2_clus2	node1_clus2
none					

All ONTAP releases

For all ONTAP releases, you can also use the `cluster ping-cluster -node <name>` command to check the connectivity:

```
cluster ping-cluster -node <name>
```

```

cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1      e0a
Cluster node1_clus2 169.254.49.125 node1      e0b
Cluster node2_clus1 169.254.47.194 node2      e0a
Cluster node2_clus2 169.254.19.183 node2      e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
    Local 169.254.19.183 to Remote 169.254.209.69
    Local 169.254.19.183 to Remote 169.254.49.125
    Local 169.254.47.194 to Remote 169.254.209.69
    Local 169.254.47.194 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)

```

5. Enable the Ethernet switch health monitor log collection feature for collecting switch-related log files.

ONTAP 9.8 and later

Enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the following two commands: `system switch ethernet log setup-password` and `system switch ethernet log enable-collection`

NOTE: You will need the password for the **admin** user on the switches.

Enter: `system switch ethernet log setup-password`

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: <return>
```

```
The switch name entered is not recognized.
```

```
Choose from the following list:
```

```
cs1-new
```

```
cs2-new
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs1-new
```

```
RSA key fingerprint is e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
```

```
Do you want to continue? {y|n}::[n] y
```

```
Enter the password: <password of switch's admin user>
```

```
Enter the password again: <password of switch's admin user>
```

```
cluster1::*> system switch ethernet log setup-password
```

```
Enter the switch name: cs2-new
```

```
RSA key fingerprint is 57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
```

```
Do you want to continue? {y|n}:: [n] y
```

```
Enter the password: <password of switch's admin user>
```

```
Enter the password again: <password of switch's admin user>
```

Followed by: `system switch ethernet log enable-collection`

```
cluster1::*> system switch ethernet log enable-collection
```

Do you want to enable cluster log collection for all nodes in the cluster?

```
{y|n}: [n] y
```

Enabling cluster switch log collection.

```
cluster1::*>
```

NOTE: If any of these commands return an error, contact NetApp support.

ONTAP releases 9.5P16, 9.6P12, and 9.7P10 and later patch releases

Enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands: `system cluster-switch log setup-password` and `system cluster-switch log enable-collection`

NOTE: You will need the password for the **admin** user on the switches.

Enter: `system cluster-switch log setup-password`

```
cluster1::*> system cluster-switch log setup-password
```

Enter the switch name: <return>

The switch name entered is not recognized.

Choose from the following list:

cs1-new

cs2-new

```
cluster1::*> system cluster-switch log setup-password
```

Enter the switch name: **cs1-new**

RSA key fingerprint is e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc

Do you want to continue? {y|n}::[n] **y**

Enter the password: <password of switch's admin user>

Enter the password again: <password of switch's admin user>

```
cluster1::*> system cluster-switch log setup-password
```

Enter the switch name: **cs2-new**

RSA key fingerprint is 57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1

Do you want to continue? {y|n}:: [n] **y**

Enter the password: <password of switch's admin user>

Enter the password again: <password of switch's admin user>

Followed by: `system cluster-switch log enable-collection`

```
cluster1::*> system cluster-switch log enable-collection
```

```
Do you want to enable cluster log collection for all nodes in the
cluster?
```

```
{y|n}: [n] y
```

```
Enabling cluster switch log collection.
```

```
cluster1::*>
```

NOTE: If any of these commands return an error, contact NetApp support.

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message: `system node autosupport invoke -node * -type all -message MAINT=END`

Migrate to two-node switched cluster

If you have an existing two-node *switchless* cluster environment, you can migrate to a two-node *switched* cluster environment using Cisco Nexus 9336C-FX2 switches.

The migration process works for all nodes using optical or Twinax ports, but is not supported on this switch if nodes are using onboard 10Gb BASE-T RJ45 ports for the cluster-network ports.

Review requirements

What you'll need

- For the two-node switchless configuration:
 - The two-node switchless configuration is properly set up and functioning.
 - All cluster ports are in the **up** state.
 - All cluster logical interfaces (LIFs) are in the **up** state and on their home ports.
 - See [Hardware Universe](#) for all supported ONTAP versions.
- For the Cisco Nexus 9336C-FX2 switch configuration:
 - Both switches have management network connectivity.
 - There is console access to the cluster switches.
 - Nexus 9336C-FX2 node-to-node switch and switch-to-switch connections use Twinax or fiber cables.

See [Hardware Universe](#) for more information about cabling.

- Inter-Switch Link (ISL) cables are connected to ports 1/35 and 1/36 on both 9336C-FX2 switches.
- Initial customization of both the 9336C-FX2 switches are completed, so that:
 - 9336C-FX2 switches are running the latest version of software.
 - Reference Configuration Files (RCFs) are applied to the switches.

Any site customization, such as SMTP, SNMP, and SSH, is configured on the new switches.

About the examples

The examples in this procedure use the following cluster switch and node nomenclature:

- The names of the 9336C-FX2 switches are cs1 and cs2.
- The names of the cluster SVMs are node1 and node2.
- The names of the LIFs are node1_clus1 and node1_clus2 on node 1, and node2_clus1 and node2_clus2 on node 2 respectively.
- The `cluster1::*>` prompt indicates the name of the cluster.
- The cluster ports used in this procedure are e0a and e0b.

See [Hardware Universe](#) for information about the cluster ports for your platforms.

Migrate the switches

Step 1: Prepare for migration

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Change the privilege level to advanced, entering `y` when prompted to continue:

```
set -privilege advanced
```

The advanced prompt (`*>`) appears.

Step 2: Configure ports and cabling

1. Disable all node-facing ports (not ISL ports) on both the new cluster switches cs1 and cs2.

Do not disable the ISL ports.

Show example

The following example shows that node-facing ports 1 through 34 are disabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs1(config-if-range)# shutdown
```

2. Verify that the ISL and the physical ports on the ISL between the two 9336C-FX2 switches cs1 and cs2 are up on ports 1/35 and 1/36:

```
show port-channel summary
```

Show example

The following example shows that the ISL ports are up on switch cs1:

```
cs1# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lACP mode (member)
        M - Not in use. Min-links not met
```

```
-----
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
-----
1      Po1 (SU)   Eth       LACP      Eth1/35 (P)  Eth1/36 (P)
```

The following example shows that the ISL ports are up 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lACP mode (member)
        M - Not in use. Min-links not met
```

```
-----
-----
Group Port-      Type      Protocol  Member Ports
Channel
-----
-----
1      Po1 (SU)   Eth       LACP      Eth1/35 (P)  Eth1/36 (P)
```

3. Display the list of neighboring devices:

```
show cdp neighbors
```

This command provides information about the devices that are connected to the system.

Show example

The following example lists the neighboring devices on switch cs1:

```
cs1# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID         Local Intrfce  Hldtme Capability  Platform
Port ID
cs2               Eth1/35       175    R S I s         N9K-C9336C
Eth1/35
cs2               Eth1/36       175    R S I s         N9K-C9336C
Eth1/36

Total entries displayed: 2
```

The following example lists the neighboring devices on switch cs2:

```
cs2# show cdp neighbors

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-
Bridge
                  S - Switch, H - Host, I - IGMP, r - Repeater,
                  V - VoIP-Phone, D - Remotely-Managed-Device,
                  s - Supports-STP-Dispute

Device-ID         Local Intrfce  Hldtme Capability  Platform
Port ID
cs1               Eth1/35       177    R S I s         N9K-C9336C
Eth1/35
cs1               Eth1/36       177    R S I s         N9K-C9336C
Eth1/36

Total entries displayed: 2
```

4. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```

Each port should display up for Link and healthy for Health Status.

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

Port	IPspace	Broadcast Domain	Link	MTU	Speed(Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster	up	9000	auto/10000	healthy
e0b	Cluster	Cluster	up	9000	auto/10000	healthy

```
Node: node2
```

Port	IPspace	Broadcast Domain	Link	MTU	Speed(Mbps) Admin/Oper	Health Status
e0a	Cluster	Cluster	up	9000	auto/10000	healthy
e0b	Cluster	Cluster	up	9000	auto/10000	healthy

```
4 entries were displayed.
```

5. Verify that all cluster LIFs are up and operational:

```
network interface show -vserver Cluster
```

Each cluster LIF should display true for Is Home and have a Status Admin/Oper of up/up.

Show example

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e0a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e0b	true			

4 entries were displayed.

6. Verify that auto-revert is enabled on all cluster LIFs:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::*> network interface show -vserver Cluster -fields auto-revert
```

	Logical	
Vserver	Interface	Auto-revert

Cluster		
	node1_clus1	true
	node1_clus2	true
	node2_clus1	true
	node2_clus2	true

4 entries were displayed.

7. Disconnect the cable from cluster port e0a on node1, and then connect e0a to port 1 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.

The [Hardware Universe - Switches](#) contains more information about cabling.

[Hardware Universe - Switches](#)

8. Disconnect the cable from cluster port e0a on node2, and then connect e0a to port 2 on cluster switch cs1, using the appropriate cabling supported by the 9336C-FX2 switches.
9. Enable all node-facing ports on cluster switch cs1.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs1:

```
cs1# config
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs1(config-if-range)# no shutdown
```

10. Verify that all cluster LIFs are up, operational, and display as `true` for `Is Home`:

```
network interface show -vserver Cluster
```

Show example

The following example shows that all of the LIFs are up on node1 and node2 and that Is Home results are true:

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	
-----	----				
Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true					
	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true					
	node2_clus2	up/up	169.254.19.183/16	node2	e0b
true					

4 entries were displayed.

11. Display information about the status of the nodes in the cluster:

```
cluster show
```

Show example

The following example displays information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
-----	-----	-----	-----
node1	true	true	false
node2	true	true	false

2 entries were displayed.

12. Disconnect the cable from cluster port e0b on node1, and then connect e0b to port 1 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.

13. Disconnect the cable from cluster port e0b on node2, and then connect e0b to port 2 on cluster switch cs2, using the appropriate cabling supported by the 9336C-FX2 switches.
14. Enable all node-facing ports on cluster switch cs2.

Show example

The following example shows that ports 1/1 through 1/34 are enabled on switch cs2:

```
cs2# config
Enter configuration commands, one per line. End with CNTL/Z.
cs2(config)# interface e1/1/1-4, e1/2/1-4, e1/3/1-4, e1/4/1-4,
e1/5/1-4, e1/6/1-4, e1/7-34
cs2(config-if-range)# no shutdown
```

15. Verify that all cluster ports are up:

```
network port show -ipspace Cluster
```


Show example

The following example shows that all of the cluster ports are up on node1 and node2:

```
cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster      up    9000  auto/10000
healthy  false

Node: node2

Ignore

Health
Port      IPspace      Broadcast Domain Link MTU  Admin/Oper  Status
Status
-----
e0a      Cluster      Cluster      up    9000  auto/10000
healthy  false
e0b      Cluster      Cluster      up    9000  auto/10000
healthy  false

4 entries were displayed.
```

Step 3: Verify the configuration

1. Verify that all interfaces display true for Is Home:

```
network interface show -vserver Cluster
```



This might take several minutes to complete.

Show example

The following example shows that all LIFs are up on node1 and node2 and that Is Home results are true:

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current	
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	
-----	----				
Cluster					
true	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true	node1_clus2	up/up	169.254.49.125/16	node1	e0b
true	node2_clus1	up/up	169.254.47.194/16	node2	e0a
true	node2_clus2	up/up	169.254.19.183/16	node2	e0b
true					

4 entries were displayed.

2. Verify that both nodes each have one connection to each switch:

```
show cdp neighbors
```

Show example

The following example shows the appropriate results for both switches:

```
(cs1)# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	133	H	FAS2980
node2 e0a	Eth1/2	133	H	FAS2980
cs2 Eth1/35	Eth1/35	175	R S I s	N9K-C9336C
cs2 Eth1/36	Eth1/36	175	R S I s	N9K-C9336C

Total entries displayed: 4

```
(cs2)# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	133	H	FAS2980
node2 e0b	Eth1/2	133	H	FAS2980
cs1 Eth1/35	Eth1/35	175	R S I s	N9K-C9336C
cs1 Eth1/36	Eth1/36	175	R S I s	N9K-C9336C

Total entries displayed: 4

3. Display information about the discovered network devices in your cluster:

```
network device-discovery show -protocol cdp
```

Show example

```
cluster1::*> network device-discovery show -protocol cdp
Node/      Local   Discovered
Protocol   Port    Device (LLDP: ChassisID)  Interface
Platform
-----
node2      /cdp
           e0a    cs1                      0/2      N9K-
C9336C
           e0b    cs2                      0/2      N9K-
C9336C
node1      /cdp
           e0a    cs1                      0/1      N9K-
C9336C
           e0b    cs2                      0/1      N9K-
C9336C

4 entries were displayed.
```

4. Verify that the settings are disabled:

```
network options switchless-cluster show
```



It might take several minutes for the command to complete. Wait for the '3 minute lifetime to expire' announcement.

Show example

The false output in the following example shows that the configuration settings are disabled:

```
cluster1::*> network options switchless-cluster show
Enable Switchless Cluster: false
```

5. Verify the status of the node members in the cluster:

```
cluster show
```

Show example

The following example shows information about the health and eligibility of the nodes in the cluster:

```
cluster1::*> cluster show
```

Node	Health	Eligibility	Epsilon
node1	true	true	false
node2	true	true	false

6. Verify that the cluster network has full connectivity:

```
cluster ping-cluster -node node-name
```

Show example

```
cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:

Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)

Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

7. Change the privilege level back to admin:

```
set -privilege admin
```

8. For ONTAP 9.8 and later, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

```
system switch ethernet log setup-password and system switch ethernet log enable-  
collection
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

9. For ONTAP releases 9.5P16, 9.6P12, and 9.7P10 and later patch releases, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

`system cluster-switch log setup-password` and `system cluster-switch log enable-`

collection

Show example

```
cluster1::*> system cluster-switch log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

10. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:


```
system node autosupport invoke -node * -type all -message MAINT=END
```

Replace switches

Replace a Cisco Nexus 9336C-FX2 cluster switch

Follow these steps to replace a defective Nexus 9336C-FX2 switch in a cluster network. This is a nondisruptive procedure (NDU).

Review requirements

Before performing the switch replacement, make sure that:

- On the existing cluster and network infrastructure:
 - The existing cluster is verified as completely functional, with at least one fully connected cluster switch.
 - All cluster ports are **up**.
 - All cluster logical interfaces (LIFs) are **up** and on their home ports.
 - The ONTAP `cluster ping-cluster -node node1` command must indicate that basic connectivity and larger than PMTU communication are successful on all paths.
- On the Nexus 9336C-FX2 replacement switch:
 - Management network connectivity on the replacement switch is functional.
 - Console access to the replacement switch is in place.
 - The node connections are ports 1/1 through 1/34.
 - All Inter-Switch Link (ISL) ports is disabled on ports 1/35 and 1/36.
 - The desired reference configuration file (RCF) and NX-OS operating system image switch is loaded onto the switch.
 - Initial customization of the switch is complete, as detailed in [Configure the 9336C-FX2 cluster switch](#).

Any previous site customizations, such as STP, SNMP, and SSH, are copied to the new switch.

- You have executed the command for migrating a cluster LIF from the node where the cluster LIF is hosted.

Replace the switch

About the examples

The examples in this procedure use the following switch and node nomenclature:

- The names of the existing Nexus 9336C-FX2 switches are cs1 and cs2.
- The name of the new Nexus 9336C-FX2 switch is newcs2.
- The node names are node1 and node2.
- The cluster ports on each node are named e0a and e0b.
- The cluster LIF names are node1_clus1 and node1_clus2 for node1, and node2_clus1 and node2_clus2 for node2.
- The prompt for changes to all cluster nodes is cluster1::*>

About this task

The following procedure is based on the following cluster network topology:

Show example

```
cluster1::*> network port show -ipspace Cluster
```

Node: node1

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	----	----	-----	-----

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

Node: node2

Ignore

						Speed(Mbps)	Health
Health							
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper	Status
Status							
-----	-----	-----	-----	----	----	-----	-----

e0a	Cluster	Cluster		up	9000	auto/10000	healthy
false							
e0b	Cluster	Cluster		up	9000	auto/10000	healthy
false							

4 entries were displayed.

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current	
Current Is					
Vserver	Interface	Admin/Oper	Address/Mask	Node	Port
Home					
-----	-----	-----	-----	-----	-----

Cluster					
	node1_clus1	up/up	169.254.209.69/16	node1	e0a
true					
	node1_clus2	up/up	169.254.49.125/16	node1	e0b

```

true
node2_clus1 up/up 169.254.47.194/16 node2 e0a
true
node2_clus2 up/up 169.254.19.183/16 node2 e0b
true
4 entries were displayed.

```

```
cluster1::*> network device-discovery show -protocol cdp
```

Node/ Protocol	Local Port	Discovered Device (LLDP: ChassisID)	Interface	Platform
node2	/cdp			
	e0a	cs1	Eth1/2	N9K-
C9336C				
	e0b	cs2	Eth1/2	N9K-
C9336C				
node1	/cdp			
	e0a	cs1	Eth1/1	N9K-
C9336C				
	e0b	cs2	Eth1/1	N9K-
C9336C				

4 entries were displayed.

```
cs1# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge
S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID ID	Local Intrfce	Hldtme	Capability	Platform	Port
node1	Eth1/1	144	H	FAS2980	e0a
node2	Eth1/2	145	H	FAS2980	e0a
cs2	Eth1/35	176	R S I s	N9K-C9336C	
Eth1/35					
cs2 (FD0220329V5)	Eth1/36	176	R S I s	N9K-C9336C	
Eth1/36					

Total entries displayed: 4

```
cs2# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge  
S - Switch, H - Host, I - IGMP, r - Repeater,  
V - VoIP-Phone, D - Remotely-Managed-Device,  
s - Supports-STP-Dispute
```

Device-ID	Local Intrfce	Hldtme	Capability	Platform	Port
ID					
node1	Eth1/1	139	H	FAS2980	e0b
node2	Eth1/2	124	H	FAS2980	e0b
cs1	Eth1/35	178	R S I s	N9K-C9336C	
Eth1/35					
cs1	Eth1/36	178	R S I s	N9K-C9336C	
Eth1/36					

```
Total entries displayed: 4
```

Step 1: Prepare for replacement

1. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=xh
```

where x is the duration of the maintenance window in hours.



The AutoSupport message notifies technical support of this maintenance task so that automatic case creation is suppressed during the maintenance window.

2. Install the appropriate RCF and image on the switch, newcs2, and make any necessary site preparations.

If necessary, verify, download, and install the appropriate versions of the RCF and NX-OS software for the new switch. If you have verified that the new switch is correctly set up and does not need updates to the RCF and NX-OS software, continue to step 2.

- a. Go to the *NetApp Cluster and Management Network Switches Reference Configuration File Description Page* on the NetApp Support Site.
 - b. Click the link for the *Cluster Network and Management Network Compatibility Matrix*, and then note the required switch software version.
 - c. Click your browser's back arrow to return to the Description page, click **CONTINUE**, accept the license agreement, and then go to the Download page.
 - d. Follow the steps on the Download page to download the correct RCF and NX-OS files for the version of ONTAP software you are installing.
3. On the new switch, log in as admin and shut down all of the ports that will be connected to the node cluster interfaces (ports 1/1 to 1/34).

If the switch that you are replacing is not functional and is powered down, go to Step 4. The LIFs on the

cluster nodes should have already failed over to the other cluster port for each node.

Show example

```
newcs2# config
Enter configuration commands, one per line. End with CNTL/Z.
newcs2(config)# interface e1/1-34
newcs2(config-if-range)# shutdown
```

4. Verify that all cluster LIFs have auto-revert enabled:

```
network interface show -vserver Cluster -fields auto-revert
```

Show example

```
cluster1::> network interface show -vserver Cluster -fields auto-
revert
```

Vserver	Logical Interface	Auto-revert
Cluster	node1_clus1	true
Cluster	node1_clus2	true
Cluster	node2_clus1	true
Cluster	node2_clus2	true

4 entries were displayed.

5. Verify that all the cluster LIFs can communicate:

```
cluster ping-cluster
```

Show example

```
cluster1::*> cluster ping-cluster node1

Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

Step 2: Configure cables and ports

1. Shut down the ISL ports 1/35 and 1/36 on the Nexus 9336C-FX2 switch cs1.

Show example

```
cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# interface e1/35-36
cs1(config-if-range)# shutdown
cs1(config-if-range)#
```

2. Remove all of the cables from the Nexus 9336C-FX2 cs2 switch, and then connect them to the same ports on the Nexus C9336C-FX2 newcs2 switch.

3. Bring up the ISLs ports 1/35 and 1/36 between the cs1 and newcs2 switches, and then verify the port channel operation status.

Port-Channel should indicate Po1(SU) and Member Ports should indicate Eth1/35(P) and Eth1/36(P).

Show example

This example enables ISL ports 1/35 and 1/36 and displays the port channel summary on switch cs1:

```
cs1# configure
Enter configuration commands, one per line. End with CNTL/Z.
cs1(config)# int e1/35-36
cs1(config-if-range)# no shutdown

cs1(config-if-range)# 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
        b - BFD Session Wait
        S - Switched      R - Routed
        U - Up (port-channel)
        p - Up in delay-lacp mode (member)
        M - Not in use. Min-links not met

-----
-----
Group Port-      Type      Protocol  Member      Ports
Channel
-----
-----
1      Po1 (SU)    Eth       LACP        Eth1/35 (P)  Eth1/36 (P)

cs1(config-if-range)#
```

4. Verify that port e0b is up on all nodes:

```
network port show ipspace Cluster
```


Show example

The output should be similar to the following:

```
cluster1::*> network port show -ipspace Cluster

Node: node1

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e0a         Cluster      Cluster      up    9000  auto/10000
healthy     false
e0b         Cluster      Cluster      up    9000  auto/10000
healthy     false

Node: node2

Ignore

Health      Health      Speed (Mbps)
Port        IPspace      Broadcast Domain Link MTU  Admin/Oper
Status      Status
-----
e0a         Cluster      Cluster      up    9000  auto/10000
healthy     false
e0b         Cluster      Cluster      up    9000  auto/auto  -
false

4 entries were displayed.
```

5. On the same node you used in the previous step, revert the cluster LIF associated with the port in the previous step by using the network interface revert command.

Show example

In this example, LIF node1_clus2 on node1 is successfully reverted if the Home value is true and the port is e0b.

The following commands return LIF node1_clus2 on node1 to home port e0a and displays information about the LIFs on both nodes. Bringing up the first node is successful if the Is Home column is true for both cluster interfaces and they show the correct port assignments, in this example e0a and e0b on node1.

```
cluster1::*> network interface show -vserver Cluster
```

Current Is	Logical	Status	Network	Current
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			

Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1
e0b	true			
	node2_clus1	up/up	169.254.47.194/16	node2
e0a	true			
	node2_clus2	up/up	169.254.19.183/16	node2
e0a	false			

4 entries were displayed.

6. Display information about the nodes in a cluster:

```
cluster show
```

Show example

This example shows that the node health for node1 and node2 in this cluster is true:

```
cluster1::*> cluster show
```

Node	Health	Eligibility
-----	-----	-----
node1	false	true
node2	true	true

7. Verify that all physical cluster ports are up:

```
network port show ipspace Cluster
```

Show example

```
cluster1::*> network port show -ipspace Cluster

Node node1
Ignore
Speed (Mbps)
Health  Health
Port    IPspace  Broadcast Domain  Link  MTU  Admin/Oper
Status  Status
-----
e0a      Cluster  Cluster          up    9000  auto/10000
healthy  false
e0b      Cluster  Cluster          up    9000  auto/10000
healthy  false

Node: node2

Ignore
Speed (Mbps)
Health  Health
Port    IPspace  Broadcast Domain  Link  MTU  Admin/Oper
Status  Status
-----
e0a      Cluster  Cluster          up    9000  auto/10000
healthy  false
e0b      Cluster  Cluster          up    9000  auto/10000
healthy  false

4 entries were displayed.
```

8. Verify that all the cluster LIFs can communicate:

```
cluster ping-cluster
```

Show example

```
cluster1::*> cluster ping-cluster -node node2
Host is node2
Getting addresses from network interface table...
Cluster node1_clus1 169.254.209.69 node1 e0a
Cluster node1_clus2 169.254.49.125 node1 e0b
Cluster node2_clus1 169.254.47.194 node2 e0a
Cluster node2_clus2 169.254.19.183 node2 e0b
Local = 169.254.47.194 169.254.19.183
Remote = 169.254.209.69 169.254.49.125
Cluster Vserver Id = 4294967293
Ping status:
....
Basic connectivity succeeds on 4 path(s)
Basic connectivity fails on 0 path(s)
.....
Detected 9000 byte MTU on 4 path(s):
Local 169.254.47.194 to Remote 169.254.209.69
Local 169.254.47.194 to Remote 169.254.49.125
Local 169.254.19.183 to Remote 169.254.209.69
Local 169.254.19.183 to Remote 169.254.49.125
Larger than PMTU communication succeeds on 4 path(s)
RPC status:
2 paths up, 0 paths down (tcp check)
2 paths up, 0 paths down (udp check)
```

9. Confirm the following cluster network configuration:

```
network port show
```

Show example

```
cluster1::*> network port show -ipspace Cluster
```

```
Node: node1
```

```
Ignore
```

				Speed (Mbps)		Health
Health						
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status						Status
-----	-----	-----	-----	----	-----	-----
-----	-----					
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

```
Node: node2
```

```
Ignore
```

				Speed (Mbps)		Health
Health						
Port	IPspace	Broadcast	Domain	Link	MTU	Admin/Oper
Status						Status
-----	-----	-----	-----	----	-----	-----
-----	-----					
e0a	Cluster	Cluster		up	9000	auto/10000
healthy	false					
e0b	Cluster	Cluster		up	9000	auto/10000
healthy	false					

```
4 entries were displayed.
```

```
cluster1::*> network interface show -vserver Cluster
```

	Logical	Status	Network	Current
Current Is				
Vserver	Interface	Admin/Oper	Address/Mask	Node
Port	Home			
-----	-----	-----	-----	-----
-----	-----			
Cluster				
	node1_clus1	up/up	169.254.209.69/16	node1
e0a	true			
	node1_clus2	up/up	169.254.49.125/16	node1

```

e0b      true
          node2_clus1  up/up    169.254.47.194/16  node2
e0a      true
          node2_clus2  up/up    169.254.19.183/16  node2
e0b      true

```

4 entries were displayed.

```
cluster1::> network device-discovery show -protocol cdp
```

Node/ Protocol Platform	Local Port	Discovered Device (LLDP: ChassisID)	Interface	
node2	/cdp			
	e0a	cs1	0/2	N9K-
C9336C				
	e0b	newcs2	0/2	N9K-
C9336C				
node1	/cdp			
	e0a	cs1	0/1	N9K-
C9336C				
	e0b	newcs2	0/1	N9K-
C9336C				

4 entries were displayed.

```
cs1# show cdp neighbors
```

Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-Bridge

S - Switch, H - Host, I - IGMP, r - Repeater,
V - VoIP-Phone, D - Remotely-Managed-Device,
s - Supports-STP-Dispute

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1	Eth1/1	144	H	FAS2980
e0a				
node2	Eth1/2	145	H	FAS2980
e0a				
newcs2	Eth1/35	176	R S I s	N9K-C9336C
Eth1/35				
newcs2	Eth1/36	176	R S I s	N9K-C9336C

```
Eth1/36
```

```
Total entries displayed: 4
```

```
cs2# show cdp neighbors
```

```
Capability Codes: R - Router, T - Trans-Bridge, B - Source-Route-  
Bridge
```

```
S - Switch, H - Host, I - IGMP, r - Repeater,  
V - VoIP-Phone, D - Remotely-Managed-Device,  
s - Supports-STP-Dispute
```

Device-ID Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0b	Eth1/1	139	H	FAS2980
node2 e0b	Eth1/2	124	H	FAS2980
cs1 Eth1/35	Eth1/35	178	R S I s	N9K-C9336C
cs1 Eth1/36	Eth1/36	178	R S I s	N9K-C9336C

```
Total entries displayed: 4
```

Step 3: Verify the configuration

1. For ONTAP 9.8 and later, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

```
system switch ethernet log setup-password and system switch ethernet log enable-  
collection
```

Show example

```
cluster1::*> system switch ethernet log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system switch ethernet log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

2. For ONTAP releases 9.5P16, 9.6P12, and 9.7P10 and later patch releases, enable the Ethernet switch health monitor log collection feature for collecting switch-related log files, using the commands:

`system cluster-switch log setup-password` and `system cluster-switch log enable-`

collection

Show example

```
cluster1::*> system cluster-switch log setup-password
Enter the switch name: <return>
The switch name entered is not recognized.
Choose from the following list:
cs1
cs2

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs1
RSA key fingerprint is
e5:8b:c6:dc:e2:18:18:09:36:63:d9:63:dd:03:d9:cc
Do you want to continue? {y|n}::[n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log setup-password

Enter the switch name: cs2
RSA key fingerprint is
57:49:86:a1:b9:80:6a:61:9a:86:8e:3c:e3:b7:1f:b1
Do you want to continue? {y|n}:: [n] y

Enter the password: <enter switch password>
Enter the password again: <enter switch password>

cluster1::*> system cluster-switch log enable-collection

Do you want to enable cluster log collection for all nodes in the
cluster?
{y|n}: [n] y

Enabling cluster switch log collection.

cluster1::*>
```



If any of these commands return an error, contact NetApp support.

3. If you suppressed automatic case creation, re-enable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

Replace Cisco Nexus 9336C-FX2 cluster switches with switchless connections

You can migrate from a cluster with a switched cluster network to one where two nodes are directly connected for ONTAP 9.3 and later.

Review requirements

Guidelines

Review the following guidelines:

- Migrating to a two-node switchless cluster configuration is a nondisruptive operation. Most systems have two dedicated cluster interconnect ports on each node, but you can also use this procedure for systems with a larger number of dedicated cluster interconnect ports on each node, such as four, six or eight.
- You cannot use the switchless cluster interconnect feature with more than two nodes.
- If you have an existing two-node cluster that uses cluster interconnect switches and is running ONTAP 9.3 or later, you can replace the switches with direct, back-to-back connections between the nodes.


What you'll need

- A healthy cluster that consists of two nodes connected by cluster switches. The nodes must be running the same ONTAP release.
- Each node with the required number of dedicated cluster ports, which provide redundant cluster interconnect connections to support your system configuration. For example, there are two redundant ports for a system with two dedicated cluster interconnect ports on each node.

Migrate the switches

About this task

The following procedure removes the cluster switches in a two-node cluster and replaces each connection to the switch with a direct connection to the partner node.



About the examples

The examples in the following procedure show nodes that are using "e0a" and "e0b" as cluster ports. Your

nodes might be using different cluster ports as they vary by system.

Step 1: Prepare for migration

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

```
set -privilege advanced
```

The advanced prompt `*>` appears.

2. ONTAP 9.3 and later supports automatic detection of switchless clusters, which is enabled by default.

You can verify that detection of switchless clusters is enabled by running the advanced privilege command:

```
network options detect-switchless-cluster show
```

Show example

The following example output shows if the option is enabled.

```
cluster::*> network options detect-switchless-cluster show
(network options detect-switchless-cluster show)
Enable Switchless Cluster Detection: true
```

If "Enable Switchless Cluster Detection" is `false`, contact NetApp support.

3. If AutoSupport is enabled on this cluster, suppress automatic case creation by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message
MAINT=<number_of_hours>h
```

where `h` is the duration of the maintenance window in hours. The message notifies technical support of this maintenance task so that they can suppress automatic case creation during the maintenance window.

In the following example, the command suppresses automatic case creation for two hours:

Show example

```
cluster::*> system node autosupport invoke -node * -type all
-message MAINT=2h
```


Step 2: Configure ports and cabling

1. Organize the cluster ports on each switch into groups so that the cluster ports in group1 go to cluster switch1 and the cluster ports in group2 go to cluster switch2. These groups are required later in the procedure.

2. Identify the cluster ports and verify link status and health:

```
network port show -ipspace Cluster
```

In the following example for nodes with cluster ports "e0a" and "e0b", one group is identified as "node1:e0a" and "node2:e0a" and the other group as "node1:e0b" and "node2:e0b". Your nodes might be using different cluster ports because they vary by system.



Verify that the ports have a value of `up` for the "Link" column and a value of `healthy` for the "Health Status" column.

Show example

```
cluster::> network port show -ipspace Cluster
Node: node1

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false

Node: node2

Ignore
Speed (Mbps) Health
Health
Port IPspace Broadcast Domain Link MTU Admin/Oper Status
Status
-----
-----
e0a Cluster Cluster up 9000 auto/10000 healthy
false
e0b Cluster Cluster up 9000 auto/10000 healthy
false
4 entries were displayed.
```

3. Confirm that all the cluster LIFs are on their home ports.

Verify that the “is-home” column is `true` for each of the cluster LIFs:

```
network interface show -vserver Cluster -fields is-home
```

Show example

```
cluster::*> net int show -vserver Cluster -fields is-home
(network interface show)
vserver  lif          is-home
-----  -
Cluster  node1_clus1  true
Cluster  node1_clus2  true
Cluster  node2_clus1  true
Cluster  node2_clus2  true
4 entries were displayed.
```

If there are cluster LIFs that are not on their home ports, revert those LIFs to their home ports:

```
network interface revert -vserver Cluster -lif *
```

4. Disable auto-revert for the cluster LIFs:

```
network interface modify -vserver Cluster -lif * -auto-revert false
```

5. Verify that all ports listed in the previous step are connected to a network switch:

```
network device-discovery show -port cluster_port
```

The “Discovered Device” column should be the name of the cluster switch that the port is connected to.

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to cluster switches "cs1" and "cs2".

```
cluster::> network device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol  Port   Device (LLDP: ChassisID)  Interface  Platform
-----  -
node1/cdp
          e0a    cs1                      0/11       BES-53248
          e0b    cs2                      0/12       BES-53248
node2/cdp
          e0a    cs1                      0/9        BES-53248
          e0b    cs2                      0/9        BES-53248
4 entries were displayed.
```

6. Verify the cluster connectivity:

```
cluster ping-cluster -node local
```

7. Verify that the cluster is healthy:

```
cluster ring show
```

All units must be either master or secondary.


8. Set up the switchless configuration for the ports in group 1.



To avoid potential networking issues, you must disconnect the ports from group1 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


a. Disconnect all the cables from the ports in group1 at the same time.

In the following example, the cables are disconnected from port "e0a" on each node, and cluster traffic continues through the switch and port "e0b" on each node:



b. Cable the ports in group1 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2:



9. The switchless cluster network option transitions from `false` to `true`. This might take up to 45 seconds. Confirm that the switchless option is set to `true`:

```
network options switchless-cluster show
```

The following example shows that the switchless cluster is enabled:

```
cluster::*> network options switchless-cluster show
Enable Switchless Cluster: true
```

10. Verify that the cluster network is not disrupted:

```
cluster ping-cluster -node local
```



Before proceeding to the next step, you must wait at least two minutes to confirm a working back-to-back connection on group 1.


11. Set up the switchless configuration for the ports in group 2.



To avoid potential networking issues, you must disconnect the ports from group2 and reconnect them back-to-back as quickly as possible, for example, **in less than 20 seconds**.


- a. Disconnect all the cables from the ports in group2 at the same time.

In the following example, the cables are disconnected from port "e0b" on each node, and cluster traffic continues through the direct connection between the "e0a" ports:



b. Cable the ports in group2 back-to-back.

In the following example, "e0a" on node1 is connected to "e0a" on node2 and "e0b" on node1 is connected to "e0b" on node2:



Step 3: Verify the configuration

1. Verify that the ports on both nodes are correctly connected:

```
network device-discovery show -port cluster_port
```

Show example

The following example shows that cluster ports "e0a" and "e0b" are correctly connected to the corresponding port on the cluster partner:

```
cluster::> net device-discovery show -port e0a|e0b
(network device-discovery show)
Node/      Local  Discovered
Protocol   Port   Device (LLDP: ChassisID)  Interface  Platform
-----
node1/cdp
          e0a    node2                      e0a        AFF-A300
          e0b    node2                      e0b        AFF-A300
node1/lldp
          e0a    node2 (00:a0:98:da:16:44) e0a        -
          e0b    node2 (00:a0:98:da:16:44) e0b        -
node2/cdp
          e0a    node1                      e0a        AFF-A300
          e0b    node1                      e0b        AFF-A300
node2/lldp
          e0a    node1 (00:a0:98:da:87:49) e0a        -
          e0b    node1 (00:a0:98:da:87:49) e0b        -
8 entries were displayed.
```

2. Re-enable auto-revert for the cluster LIFs:

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

3. Verify that all LIFs are home. This might take a few seconds.

```
network interface show -vserver Cluster -lif lif_name
```

Show example

The LIFs have been reverted if the “Is Home” column is `true`, as shown for `node1_clus2` and `node2_clus2` in the following example:

```
cluster::> network interface show -vserver Cluster -fields curr-  
port,is-home  
vserver  lif                curr-port is-home  
-----  
Cluster  node1_clus1        e0a      true  
Cluster  node1_clus2        e0b      true  
Cluster  node2_clus1        e0a      true  
Cluster  node2_clus2        e0b      true  
4 entries were displayed.
```

If any cluster LIFS have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver Cluster -lif lif_name
```

4. Check the cluster status of the nodes from the system console of either node:

```
cluster show
```

Show example

The following example shows `epsilon` on both nodes to be `false`:

```
Node  Health  Eligibility Epsilon  
-----  
node1 true    true       false  
node2 true    true       false  
2 entries were displayed.
```

5. Confirm connectivity between the cluster ports:

```
cluster ping-cluster local
```

6. If you suppressed automatic case creation, reenable it by invoking an AutoSupport message:

```
system node autosupport invoke -node * -type all -message MAINT=END
```

For more information, see [NetApp KB Article 1010449: How to suppress automatic case creation during scheduled maintenance windows](#).

7. Change the privilege level back to admin:

```
set -privilege admin
```

Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.