



Install or upgrade the RCF

Cluster and storage switches

NetApp
January 17, 2025

Table of Contents

- Install or upgrade the RCF 1
- Install or upgrade the Reference Configuration File (RCF) overview 1
- Install the Reference Configuration File (RCF) 2
- Upgrade you Reference Configuration File (RCF) 11

Install or upgrade the RCF

Install or upgrade the Reference Configuration File (RCF) overview

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 switch for the first time. You upgrade your RCF version when you have an existing version of the RCF file installed on your switch.

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.
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.

Suggested documentation

- [Cisco Ethernet Switches \(NSS\)](#)

Consult the switch compatibility table for the supported ONTAP and RCF versions on the NetApp Support Site. Note that there can be command dependencies between the command syntax in the RCF and the

syntax found in specific versions of NX-OS.

- [Cisco Nexus 3000 Series Switches](#)

Refer to the appropriate software and upgrade guides available on the Cisco website for complete documentation on the Cisco switch upgrade and downgrade procedures.

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 four 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.

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

Commands used

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

What's next?

[Install the RCF](#) or [Upgrade your RCF](#).

Install the Reference Configuration File (RCF)

You install the Reference Configuration File (RCF) after setting up the Nexus 9336C-FX2 switch for the first time.

Before you begin

Verify the following installations and connections:

- A console connection to the switch. The console connection is optional if you have remote access to the switch.
- Switch cs1 and switch cs2 are powered up and the initial switch setup is complete (the Management IP address and SSH is setup).
- The desired NX-OS version has been installed.
- ISL connections between switches are connected.
- ONTAP node cluster ports are not connected.

Step 1: Install the RCF on the switches

1. Login to switch cs1 using SSH or by using a serial console.
2. Copy the RCF to the bootflash of switch cs1 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 cs1:

```
cs1# 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)...
```

3. 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 cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

4. 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

```
cs1# 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
*
*****
*****
```

5. 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.

6. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
7. After you verify the RCF versions and switch settings are correct, copy the running-config file to the startup-config file.

```
copy running-config startup-config
```

Show example

```
cs1# copy running-config startup-config
[#####] 100% Copy complete
```

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

8. Reboot switch cs1.

```
cs1# reload
```

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

9. Repeat steps 1 through 7 on switch cs2.
10. Connect the cluster ports of all nodes in the ONTAP cluster to switches cs1 and cs2.

Step 2: Verify the switch connections

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 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                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         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                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port          Native  Status      Port
              Vlan              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
Po1          1
..
..
..
..
..
```



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

3. 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      Po1 (SU)       Eth       LACP        Eth1/35 (P)       Eth1/36 (P)
cs1#
```

Step 3: Set up your ONTAP cluster

NetApp recommends that you use System Manager to set up new clusters.

System Manager provides a simple and easy workflow for cluster set up and configuration including assigning a node management IP address, initializing the cluster, creating a local tier, configuring protocols and provisioning initial storage.

Go to [Configure ONTAP on a new cluster with System Manager](#) for setup instructions.

What's next?

[Verify SSH configuration.](#)

Upgrade your Reference Configuration File (RCF)

You upgrade your RCF version when you have an existing version of the RCF file installed on your operational switches.

Before you begin

Make sure you have the following:

- A current backup of the switch configuration.
- A fully functioning cluster (no errors in the logs or similar issues).
- The current RCF.
- If you are updating your RCF version, you need a boot configuration in the RCF that reflects the desired boot images.

If you need to change the boot configuration to reflect the current boot images, you must do so before reapplying the RCF so that the correct version is instantiated on future reboots.



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 or have preserved basic configuration information prior to erasing the switch settings.

Step 1: Prepare for the upgrade

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

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

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

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  cluster1-01_clus1  up/up      169.254.3.4/23
              e0a      true
cluster1-01  cluster1-01_clus2  up/up      169.254.3.5/23
              e0d      true
cluster1-02  cluster1-02_clus1  up/up      169.254.3.8/23
              e0a      true
cluster1-02  cluster1-02_clus2  up/up      169.254.3.9/23
              e0d      true
cluster1-03  cluster1-03_clus1  up/up      169.254.1.3/23
              e0a      true
cluster1-03  cluster1-03_clus2  up/up      169.254.1.1/23
              e0b      true
cluster1-04  cluster1-04_clus1  up/up      169.254.1.6/23
              e0a      true
cluster1-04  cluster1-04_clus2  up/up      169.254.1.7/23
              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.

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

Step 2: Configure ports

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

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

```
cs1(config-if-range)# shutdown
```



Make sure to shutdown **all** connected cluster ports to avoid any network connection issues. See the Knowledge Base article [Node out of quorum when migrating cluster LIF during switch OS upgrade](#) for further details.

2. 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 -role cluster
```

Show example

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

Current Is	Logical	Status	Network	Current
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
```

- a. Record any custom additions between the current running-config and the RCF file in use (such as an SNMP configuration for your organization).
- b. For NX-OS 10.2 and newer use the `show diff running-config` command to compare with the saved RCF file in the bootflash. Otherwise, use a third part diff/compare tool.

5. Save basic configuration details to the `write_erase.cfg` file on the bootflash.

```
switch# show run | i "username admin password" > bootflash:write_erase.cfg
```

```
switch# show run | section "vrf context management" >>
bootflash:write_erase.cfg
```

```
switch# show run | section "interface mgmt0" >> bootflash:write_erase.cfg
```

```
switch# show run | section "switchname" >> bootflash:write_erase.cfg
```

6. Issue the write erase command to erase the current saved configuration:

```
switch# write erase
```

Warning: This command will erase the startup-configuration.

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

7. Copy the previously saved basic configuration into the startup configuration.

```
switch# copy write_erase.cfg startup-config
```

8. Perform a reboot of the switch:

```
switch# reload
```

This command will reboot the system. (y/n)? [n] y

9. After the management IP address is reachable again, log in to the switch through SSH.

You may need to update host file entries related to the SSH keys.

10. Copy the RCF to the bootflash of switch cs1 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 cs1:

```
cs1# 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)...
```

11. 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 cs1:

```
cs1# copy Nexus_9336C_RCF_v1.6-Cluster-HA-Breakout.txt running-
config echo-commands
```

12. 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

```
cs1# 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
*
*****
*****
```

13. 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.

14. Reapply any previous customizations to the switch configuration. Refer to [Review cabling and configuration considerations](#) for details of any further changes required.
15. After you verify the RCF versions, custom additions, 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.

```
cs1# copy running-config startup-config
```

```
[] 100% Copy complete
```

16. Reboot switch cs1. You can ignore the “cluster switch health monitor” alerts and “cluster ports down” events reported on the nodes while the switch reboots.

```
cs1# reload
```

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

17. Verify the health of cluster ports on the cluster.
 - a. Verify that cluster 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.

```
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
-----
-----
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.

```

18. 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::*>

```

19. Repeat steps 1 to 18 on switch cs2.

20. Enable auto-revert on the cluster LIFs.

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

Step 3: Verify the cluster network configuration and cluster health

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 Port ID	Local Intrfce	Hldtme	Capability	Platform
node1 e0a	Eth1/1	133	H	FAS2980
node2 e0a	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. 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                               Eth1/26, Eth1/27,
Eth1/31                               Eth1/29, Eth1/30,
Eth1/34                               Eth1/32, Eth1/33,
33   VLAN0033                         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                               Eth1/23, Eth1/24,
34   VLAN0034                         active   Eth1/26, Eth1/27,
Eth1/25                               Eth1/29, Eth1/30,
Eth1/28                               Eth1/32, Eth1/33,
Eth1/31
Eth1/34

```

```
cs1# show interface trunk
```

```

-----
Port                Native  Status      Port
                   Vlan              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
Po1         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      Po1 (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::*>
```

If any cluster LIFs have not returned to their home ports, revert them manually from the local node:

```
network interface revert -vserver vservice_name -lif lif_name
```

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. 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
```

Packet	Source	Destination
Node	Date	LIF
Loss		
-----	-----	-----
node1		
3/5/2022 19:21:18 -06:00	cluster1-01_clus2	cluster1-02-
clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-01_clus2	cluster1-
02_clus2 none		
node2		
3/5/2022 19:21:18 -06:00	cluster1-02_clus2	cluster1-
01_clus1 none		
3/5/2022 19:21:20 -06:00	cluster1-02_clus2	cluster1-
01_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 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)

```

What's next?

[Verify SSH configuration.](#)

Copyright information

Copyright © 2025 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.