

Network Configuration Example

Four-Member QFX Series VCF Upgrade Procedure

Author: Rugved Bidkar



Published
2021-07-13

Juniper Networks, Inc.
1133 Innovation Way
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered marks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Network Configuration Example Four-Member QFX Series VCF Upgrade Procedure
Copyright © 2021 Juniper Networks, Inc. All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <https://support.juniper.net/support/eula/>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

1

[About This Guide | iv](#)

[Example: How to Upgrade a Four-Member QFX VCF](#)

[How to Upgrade a Four-Member QFX Series VCF | 2](#)

[About This Network Configuration Example | 2](#)

[Configuration Example | 2](#)

[Requirements | 3](#)

[Overview | 3](#)

[Configuration | 5](#)

[Conclusion | 24](#)

About This Guide

Use this network configuration example to manually upgrade a four-member QFX Series Virtual Chassis Fabric.

1

CHAPTER

Example: How to Upgrade a Four-Member QFX VCF

[How to Upgrade a Four-Member QFX Series VCF | 2](#)

How to Upgrade a Four-Member QFX Series VCF

IN THIS SECTION

- [About This Network Configuration Example | 2](#)
- [Configuration Example | 2](#)
- [Conclusion | 24](#)

About This Network Configuration Example

This network configuration example (NCE) shows how to upgrade a four-member QFX Series Virtual Chassis Fabric (VCF) when the nonstop software upgrade (NSSU) process is either not available or undesirable. This process minimizes service disruption and has minimal impact on data center workloads.

SEE ALSO

[*Virtual Chassis Fabric Overview*](#)

[*Upgrading Software on a Virtual Chassis Fabric Using Nonstop Software Upgrade*](#)

[*Understanding Nonstop Software Upgrade on a Virtual Chassis Fabric*](#)

[*Upgrading Two-Member QFX Series Virtual Chassis*](#)

Configuration Example

IN THIS SECTION

- [Requirements | 3](#)
- [Overview | 3](#)
- [Configuration | 5](#)

Requirements

We use the following in this example:

- A two-spine and two-leaf VCF composed of QFX5100 switches running Junos OS Release 14.1X53-D47.6
- Pre-provisioned mode VCF that is configured using VCF best practices such as Virtual Chassis graceful Routing Engine switchover (GRES) and nonstop bridging (NSB)
- Layer 2-only VCF
- MX Series router as the uplink device
- Serial console access (mandatory)
- Junos OS Release 18.4R1.8

You can use this approach for an upgrade between any releases as long as all devices in the VCF are running the same release version.

You can use this procedure for the following QFX Series VCFs:

- A four-member QFX5100 VCF consisting of only QFX5100s
- A four-member QFX5110 VCF consisting of:
 - Only QFX5110s, or
 - Two QFX5110s as spine devices and two QFX5100s in mixed mode as leaf devices, or
 - Two QFX5110s as spine devices and one QFX5100 and one QFX5110 in mixed mode as leaf devices

The uplink device can be any device with routing functions.

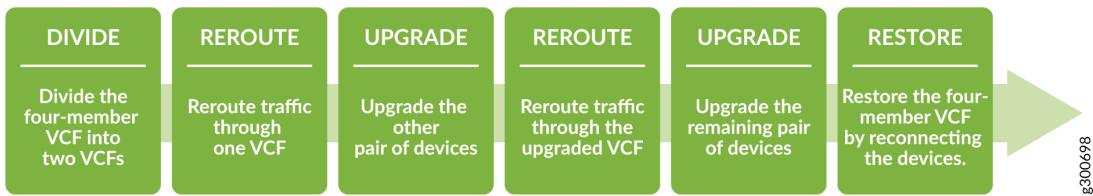
Overview

IN THIS SECTION

- [Topology | 4](#)

Sometimes it is either not possible or not desirable to upgrade a VCF to another software release using NSSU. This document shows an alternative method to upgrade a four-member QFX Series VCF with

minimal downtime. This method is not a replacement for NSSU, but a minimally invasive method that must be implemented when necessary with proper planning as noted in the following steps.



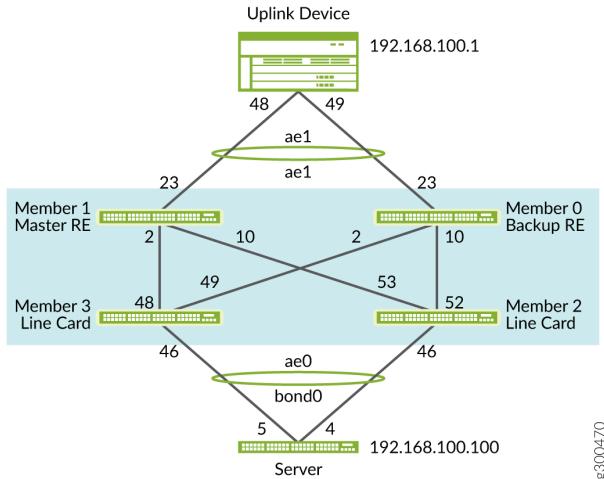
To upgrade the VCF, first divide it into two VCFs that each consist of one Routing Engine and one line card. After rerouting traffic through one VCF, upgrade the other pair of devices. Reroute traffic through the upgraded VCF before upgrading the remaining pair of devices. Restore the four-member VCF by reconnecting the devices one at a time to the new two-member VCF.

You might see alerts during this procedure, including SNMP traps and system log messages.

Topology

[Figure 1 on page 4](#) illustrates the topology of the VCF. Members 1 and 0 are connected to the uplink device, while the line cards are connected to the server.

Figure 1: Topology of the VCF



Configuration

IN THIS SECTION

- Prepare for the Upgrade | 5
- Reroute Traffic Through Member 1 and Member 3 | 10
- Upgrade Member 0 and Member 2 | 11
- Reroute Traffic Through Member 0 and Member 2 | 13
- Upgrade Member 1 and Member 3 | 14
- Restore the Four-Member VCF | 17

Prepare for the Upgrade

Step-by-Step Procedure

1. Log into the device using the root user or another login user with administrative privileges you have configured.
2. Check the status of the VCF before you begin the upgrade. Note the serial numbers, member IDs, and associated roles of the devices.

3. Check the Virtual Chassis ports (VCPs) and create a topology diagram for reference. [Figure 1 on page 4](#) shows the topology of the VCF in this example.

```
user@switch> show virtual-chassis vc-port

fpc0:
-----
Interface  Type          Trunk  Status   Speed      Neighbor
or          or             ID      ID       (mbps)    ID   Interface
PIC / Port
0/10       Configured    -1     Up       40000     2   vcp-255/0/52
0/2        Configured    -1     Up       40000     3   vcp-255/0/49

fpc1:
-----
Interface  Type          Trunk  Status   Speed      Neighbor
or          or             ID      ID       (mbps)    ID   Interface
PIC / Port
0/10       Configured    -1     Up       40000     2   vcp-255/0/53
0/2        Configured    -1     Up       40000     3   vcp-255/0/48

fpc2:
-----
Interface  Type          Trunk  Status   Speed      Neighbor
or          or             ID      ID       (mbps)    ID   Interface
PIC / Port
0/52       Configured    -1     Up       40000     0   vcp-255/0/10
0/53       Configured    -1     Up       40000     1   vcp-255/0/10

fpc3:
-----
Interface  Type          Trunk  Status   Speed      Neighbor
or          or             ID      ID       (mbps)    ID   Interface
PIC / Port
0/48       Configured    -1     Up       40000     1   vcp-255/0/2
0/49       Configured    -1     Up       40000     0   vcp-255/0/2
```

4. Check that all four members are present. Check the Junos OS image running on each device. Each device must be running the same Junos OS version. If there is a version mismatch, the device should show as inactive.

```

user@switch> show version
fpc0:
-----
Hostname: switch
Model: qfx5100-24q-2p
Junos: 14.1X53-D47.6
JUNOS Base OS Software Suite [14.1X53-D47.6]
JUNOS Base OS boot [14.1X53-D47.6]

fpc1:
-----
Hostname: switch
Model: qfx5100-24q-2p
Junos: 14.1X53-D47.6
JUNOS Base OS Software Suite [14.1X53-D47.6]
JUNOS Base OS boot [14.1X53-D47.6]

fpc2:
-----
Hostname: switch
Model: qfx5100-48s-6q
Junos: 14.1X53-D47.6
JUNOS Base OS Software Suite [14.1X53-D47.6]
JUNOS Base OS boot [14.1X53-D47.6]

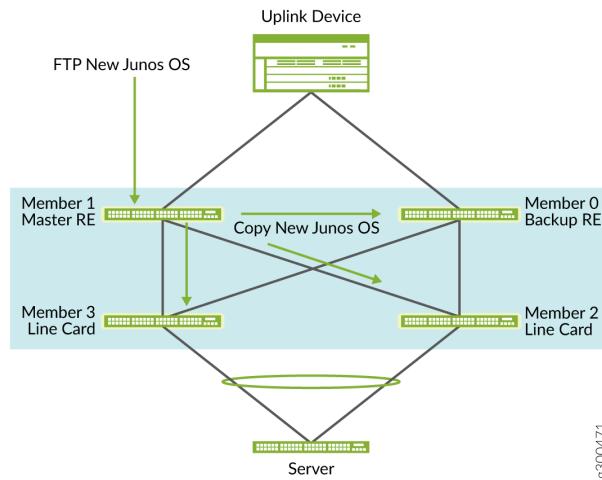
fpc3:
-----
Hostname: switch
Model: qfx5100-48s-6q
Junos: 14.1X53-D47.6
JUNOS Base OS Software Suite [14.1X53-D47.6]
JUNOS Base OS boot [14.1X53-D47.6]
JUNOS Crypto Software Suite [14.1X53-D47.6]
JUNOS Online Documentation [14.1X53-D47.6]

```

5. Use FTP to copy the new Junos OS image to the primary Routing Engine. Then copy the new image from the primary Routing Engine to the other VCF members. See *Remote Access Overview* for how to configure FTP.

Figure 2 on page 8 illustrates how the new Junos OS image is distributed among the members.

Figure 2: Copy Junos OS Image to VCF members



To copy the image from the `/var/tmp` directory on the primary Routing Engine to Member 3, also called **fpc 3's** `/var/tmp`, use the following statement:

```
file copy /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-signed.tgz fpc3:/var/tmp
```

NOTE: Copying the image can take a little while, so be patient.

Do the same for the other members. The FPC number is the same as the member number.

```
file copy /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-signed.tgz fpc0:/var/tmp  
file copy /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-signed.tgz fpc2:/var/tmp
```

6. Access each member from the VCF primary Routing Engine and confirm that the file was copied to each member. For example, to access Member 3:

```
{master:1}  
user@switch> request session member 3
```

--- JUNOS 14.1X53-D47.6 built 2018-09-08 01:46:47 UTC

Next, check the `/var/tmp` directory on this VCF member for the new Junos OS image.

```
user@switch:LC:3% cd /var/tmp/
user@switch:LC:3% ls -ltr
total 1222684
-r--r--r-- 1 root  field      505 Apr 18 19:05 preinstall_boot_loader.conf
-rw-r--r-- 1 root  field      42  Apr 18 19:07 vjunos-install.log
drwxr-xr-x  2 root  field     512 Apr 18 19:14 gres-tp
drwxrwxrwt  2 root  wheel     512 Apr 18 19:14 vi.recover
drwxrwxrwx  2 root  wheel     512 Apr 18 19:14 pics
drwxrwxrwx  2 root  wheel     512 Apr 18 19:14 install
-rw-r--r--  1 root  field      0 Apr 18 19:27 stable
-rw-r----- 1 root  field    1043 Apr 18 19:30 juniper.conf+.gz
-rw-r--r--  1 root  field  625814976 Apr 19 21:28 jinstall-host-qfx-5-flex-18.4R1.8-signed.tgz
```

When you are done, use `exit` to get back to the primary device.

```
user@switch:LC:3% exit
```

Repeat the image check on each device in the VCF.

- When you split the VCF in half, you will temporarily form two Virtual Chassis with two members each. We recommend disabling split detection whenever you form a Virtual Chassis with only two members. If you do not disable split detection, the primary device may take on a linecard role and stop the control and data planes when you disconnect it from the backup Routing Engine later in this example.

Disable split detection on the primary device.

```
user@switch# set virtual-chassis no-split-detection
```

- To check for any traffic loss during the procedure, start a continuous ping from the server to IRB 192.168.100.1 on the uplink MX Series router.

```
user@server> ping 192.168.100.1

PING 192.168.100.1 (192.168.100.1) 56(84) bytes of data.
64 bytes from 192.168.100.1: icmp_seq=1 ttl=64 time=3.33 ms
64 bytes from 192.168.100.1: icmp_seq=2 ttl=64 time=6.84 ms
64 bytes from 192.168.100.1: icmp_seq=3 ttl=64 time=7.87 ms
```

```
64 bytes from 192.168.100.1: icmp_seq=4 ttl=64 time=5.91 ms
. . .
```

Reroute Traffic Through Member 1 and Member 3

Step-by-Step Procedure

1. Using the figure above, identify the Link Aggregation Control Protocol (LACP) member interfaces and VCPs you will need to disable on Members 0 and 2 to isolate them from the rest of the VCF. The VCPs you will be disabling are port 2 on Member 0 and port 53 on Member 2.

Use the command below on the primary Routing Engine (Member 1) to determine the names of the relevant interfaces. You will be disabling the LACP member interfaces towards the uplink device and servers. In this case, et-0/0/23.0 is the Member 0 upstream interface and xe-2/0/46.0 is the Member 2 downstream interface.

```
{master:1}
user@switch> show interfaces terse | match ae

et-0/0/23.0          up    up    aenet    --> ae1.0
et-1/0/23.0          up    up    aenet    --> ae1.0
xe-2/0/46.0          up    up    aenet    --> ae0.0
xe-3/0/46.0          up    up    aenet    --> ae0.0
ae0                  up    up
ae0.0                up    up    eth-switch
ae1                  up    up
ae1.0                up    up    eth-switch
```

2. Access the primary device (Member 1) console and do the following:

Disable the interface on Member 0 to the uplink device.

```
user@switch# set interfaces et-0/0/23 disable
```

Disable the interface from Member 2 to the server.

```
user@switch# set interfaces xe-2/0/46 disable
```

Commit the configuration for it to take effect.

3. On Member 1:

Delete the VCP from Member 0 towards Member 3.

Refer to Step 3 in ["Prepare for the Upgrade" on page 5](#) and your topology diagram to determine which VCP you need to disable. Under fpc0 in the table, identify the VCP in the Interface Type or PIC/Port column going to Neighbor ID 3. In this case, disable the VCP identified as PIC/Port 0/2, which is vcp-255/0/2.

```
user@switch> request virtual-chassis vc-port delete pic-slot 0 port 2 member 0
```

Delete the VCP from Member 2 towards Member 1.

```
user@switch> request virtual-chassis vc-port delete pic-slot 0 port 53 member 2
```

4. Check that the members were removed from the VCF and marked as NotPrsnt.

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled

                                Mstr      Mixed Route Neighbor List
Member ID  Status  Serial No  Model  prio  Role  Mode  Mode ID  Interface
0 (FPC 0)  NotPrsnt  XXXXXXXX000  ... 
1 (FPC 1)  Prsnt   XXXXXXXX001  ...  129   Master*  N  F   3  vcp-255/0/2
2 (FPC 2)  NotPrsnt  XXXXXXXX002  ... 
3 (FPC 3)  Prsnt   XXXXXXXX003  ...  0   Linecard N  F   1  vcp-255/0/48
```

Upgrade Member 0 and Member 2

Step-by-Step Procedure

1. Access the consoles for Members 0 and 2. Enter the following command to upgrade the members to the Junos OS image that was copied onto the devices.

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-
signed.tgz no-copy no-validate reboot
```

- Once each isolated member is upgraded, verify that the isolated members, Member 0 and Member 2, are present.

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled

          Mstr      Mixed Route Neighbor List
Member ID Status  Serial No  Model  prio  Role  Mode  Mode ID  Interface
  0 (FPC 0)  Prsnt  XXXXXXXX000  ...    129   Master   N  F    2  vcp-255/0/10
  1 (FPC 1)  NotPrsnt XXXXXXXX001  ...
  2 (FPC 2)  Prsnt  XXXXXXXX002  ...      0  Linecard* N  F    0  vcp-255/0/52
  3 (FPC 3)  NotPrsnt XXXXXXXX003  ...
```

A new VCF automatically formed because Member 0 was already configured as a backup Routing Engine, so it took on the primary Routing Engine role when it was disconnected from the original primary device. Member 2 was already configured in the linecard role.

The output above displays the VCP interfaces that link the devices. If the output does not display VCP interfaces in the last column, complete Step 3.

- If the output in the previous step does not show that Member 0 and Member 2 are connected and that they are the members of a new VCF, configure the VCP link between them.

On Member 0, enable VCP 10.

```
user@switch> request virtual-chassis vc-port set pic-slot 0 port 10
```

On Member 2, enable VCP 52.

```
user@switch> request virtual-chassis vc-port set pic-slot 0 port 52
```

- Confirm the upgrade was successful.

```
{linecard:2}
user@switch> show version

fpc0:
-----
Hostname: switch
```

```

Model: qfx5100-24q-2p
Junos 18.4R1.8
JUNOS Base OS Software Suite [18.4R1.8]
JUNOS Base OS boot [18.4R1.8]
. . .

fpc2:
-----
Hostname: switch
Model: qfx5100-48s-6q
Junos: 18.4R1.8
JUNOS Base OS Software Suite [18.4R1.8]
. . .

```

Reroute Traffic Through Member 0 and Member 2

Step-by-Step Procedure

1. Simultaneously disable the uplink and server-facing ports on the old VCF pair (Member 1 and Member 3) and enable the server and uplink interfaces in the new VCF that we formed from the upgraded Member 0 and Member 2. This redirects the traffic through the new VCF.

It is very important to commit the configuration *at the same time* on both devices so that the LACP states on the host and the uplink MX Series router are maintained. You can do this with scripting, for example with Ansible tooling.

If you do not commit the configurations at the same time, traffic will be dropped and service impacted for as long as it takes you to disable the ports on the old VCF and enable the interfaces on the new VCF.

On Member 1, remove the residual configuration from when it was the primary device of the four-member VCF.

```

user@switch# delete interfaces et-0/0/23 disable
user@switch# delete interfaces xe-2/0/46 disable

```

Disable the uplink and server-facing ports on Member 1.

```

user@switch# set interfaces et-1/0/23 disable
user@switch# set interfaces xe-3/0/46 disable

```

Enable the uplink and server-facing ports on Member 0.

```
user@switch# delete interfaces et-0/0/23 disable
user@switch# delete interfaces xe-2/0/46 disable
```

2. Run `commit` at the same time on Member 1 and Member 0.
3. Check that the continuous ping from the server to IRB 192.168.100.1 on the uplink MX Series router is still running successfully. This confirms the traffic path was switched successfully.

Upgrade Member 1 and Member 3

Step-by-Step Procedure

1. Check that the old VCF consists of one primary device and one device in a linecard role.

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled



| Member ID | Status   | Serial No   | Model | prio | Role | Mixed Route Neighbor List |      | Interface      |
|-----------|----------|-------------|-------|------|------|---------------------------|------|----------------|
|           |          |             |       |      |      | Mstr                      | Mode |                |
| 0 (FPC 0) | NotPrsnt | XXXXXXXX000 |       | ...  |      |                           |      |                |
| 1 (FPC 1) | Prsnt    | XXXXXXXX001 |       | ...  | 129  | Master*                   | N F  | 3 vcp-255/0/2  |
| 2 (FPC 2) | NotPrsnt | XXXXXXXX002 |       | ...  |      |                           |      |                |
| 3 (FPC 3) | Prsnt    | XXXXXXXX003 |       | ...  | 0    | Linecard                  | N F  | 1 vcp-255/0/48 |


```

2. Break the old VCF by deleting the VCPs between Member 1 and Member 3. Since Member 1 is the primary device, you can run these commands on Member 1.

To delete the VCP from Member 3 towards Member 1:

```
user@switch> request virtual-chassis vc-port delete pic-slot 0 port 48 member 3
```

To delete the VCP from Member 1 towards Member 3:

```
user@switch> request virtual-chassis vc-port delete pic-slot local 0 port 2
```

Verify this was successful on each device.

On Member 1:

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled

          Mstr      Mixed Route Neighbor List
Member ID Status  Serial No  Model  prio  Role    Mode  Mode ID  Interface
0 (FPC 0) NotPrsnt XXXXXXXX000 ...
1 (FPC 1) Prsnt   XXXXXXXX001 ...    129   Master*  N  F
2 (FPC 2) NotPrsnt XXXXXXXX002 ...
3 (FPC 3) NotPrsnt XXXXXXXX003 ...
```

Access the Member 3 console:

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled

          Mstr      Mixed Route Neighbor List
Member ID Status  Serial No  Model  prio  Role    Mode  Mode ID  Interface
0 (FPC 0) NotPrsnt XXXXXXXX000 ...
1 (FPC 1) NotPrsnt XXXXXXXX001 ...
2 (FPC 2) NotPrsnt XXXXXXXX002 ...
3 (FPC 3) Prsnt   XXXXXXXX003 ...    0   Linecard* N  F
```

3. Upgrade Member 3 to Junos OS Release 18.4R1.

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-
signed.tgz no-copy no-validate reboot
```

Confirm the upgrade was successful.

```
user@switch> show version

fpc3:
-----
Hostname: switch
```

```
Model: qfx5100-24q-2p
Junos: 18.4R1.8
. . .
```

4. Upgrade Member 1 to Junos OS Release 18.4R1.

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-5-flex-18.4R1.8-
signed.tgz no-copy no-validate reboot
```

Confirm the upgrade was successful.

```
user@switch> show version

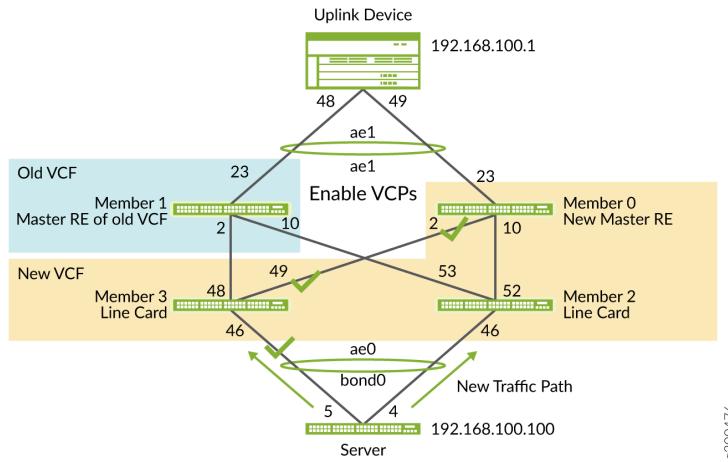
fpc1:
-----
Hostname: switch
Model: qfx5100-24q-2p
Junos: 18.4R1.8
. . .
```

Restore the Four-Member VCF

Step-by-Step Procedure

1. Add Member 3 to the new VCF by enabling VCP 49 on Member 3 and VCP 2 on Member 0. [Figure 3 on page 17](#) shows the status of the new VCF after these ports have been enabled.

Figure 3: Add Member 3 to the New VCF



On Member 3, enable the VCP towards Member 0:

```
user@switch> request virtual-chassis vc-port set pic-slot local 0 port 49
```

On Member 0:

- Enable the VCP towards Member 3:

```
user@switch> request virtual-chassis vc-port set pic-slot local 0 port 2
```

- Verify that vcp-255/0/2 is enabled on Member 0 and vcp-255/0/49 is enabled on Member 3:

```
user@switch> show virtual-chassis
```

Preprovisioned Virtual Chassis Fabric

Fabric ID: 123a.123b.123c

Fabric Mode: Enabled

Mixed Route Neighbor List										
Member ID	Status	Serial No	Model	prio	Role	Mode	Mode ID	Interface		
0 (FPC 0)	Prsnt	XXXXXXXXX000	...	129	Master*	N F	2	vcp-255/0/10		
							3	vcp-255/0/2		
1 (FPC 1)	NotPrsnt	XXXXXXXXX001	...							
2 (FPC 2)	Prsnt	XXXXXXXXX002	...	0	Linecard	N F	0	vcp-255/0/52		
3 (FPC 3)	Prsnt	XXXXXXXXX003	...	0	Linecard	N F	0	vcp-255/0/49		

2. Since Member 1 was the primary Routing Engine of the original VCF, it will have some residual configurations for Members 0, 2, and 3. These configurations might interfere with the VCF when you add it to the new VCF, especially if Member 1 preempts Member 0 as the primary device for the new VCF.

On Member 0, the primary device of the new VCF, use the following command to re-enable the server-facing interface for Member 3 and keep Member 1 from accidentally shutting it down.

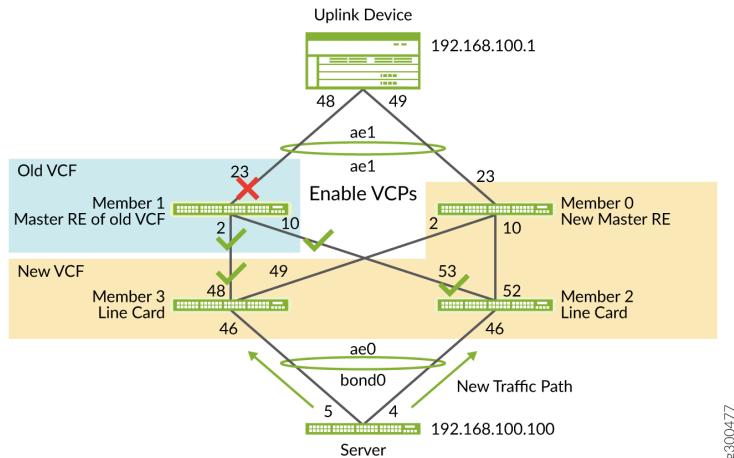
```
{master: 0}
user@switch# delete interfaces xe-3/0/46 disable
```

3. On Member 1, keep the uplink-facing interface et-1/0/23 disabled. Traffic passes to the uplink MX Series router from the neighboring new VCF primary device.

NOTE: If Member 1 preempts Member 0 as the primary device of the new VCF during the next step, the set interfaces et-1/0/23 disable statement is carried forward to the new VCF. That could cause traffic disruption, in which case this statement would need to be removed immediately.

4. To add Member 1 to the new VCF, restore the VCP links from Member 1 to Members 2 and 3, as shown in [Figure 4 on page 19](#).

Figure 4: Add Member 1 to the New VCF



On Member 1:

- Set the VCP that connects to Member 3.

```
user@switch> request virtual-chassis vc-port set pic-slot local 0 port 2
```

- Set the VCP that connects to Member 2.

```
user@switch> request virtual-chassis vc-port set pic-slot local 0 port 10
```

On Member 0, the new VCF primary device:

- Set the VCP on Member 2 that connects to Member 1.

```
user@switch> request virtual-chassis vc-port set pic-slot 0 port 53 member 2
```

- Set the VCP on Member 3 that connects to Member 1.

```
user@switch> request virtual-chassis vc-port set pic-slot 0 port 48 member 3
```

5. In most cases, the configuration of the new VCF primary Routing Engine is applied to the newly joined backup Routing Engine. Sometimes the newly joined backup Routing Engine (which was the original VCF primary Routing Engine) may preempt and take over the primary role from the newer VCF primary device. Check whether this has occurred.

```
user@switch> show virtual-chassis

Preprovisioned Virtual Chassis Fabric
Fabric ID: 123a.123b.123c
Fabric Mode: Enabled

          Mstr      Mixed Route Neighbor List
Member ID Status  Serial No  Model  prio  Role    Mode  Mode ID Interface
  0 (FPC 0)  Prsnt  XXXXXXXX000  ...    129   Backup   N   F    2  vcp-255/0/10
                                         3  vcp-255/0/2
  1 (FPC 1)  Prsnt  XXXXXXXX001  ...    129   Master*  N   F    2  vcp-255/0/10
  2 (FPC 2)  Prsnt  XXXXXXXX002  ...     0   Linecard  N   F    0  vcp-255/0/52
                                         1  vcp-255/0/53
  3 (FPC 3)  Prsnt  XXXXXXXX003  ...     0   Linecard  N   F    0  vcp-255/0/49
```

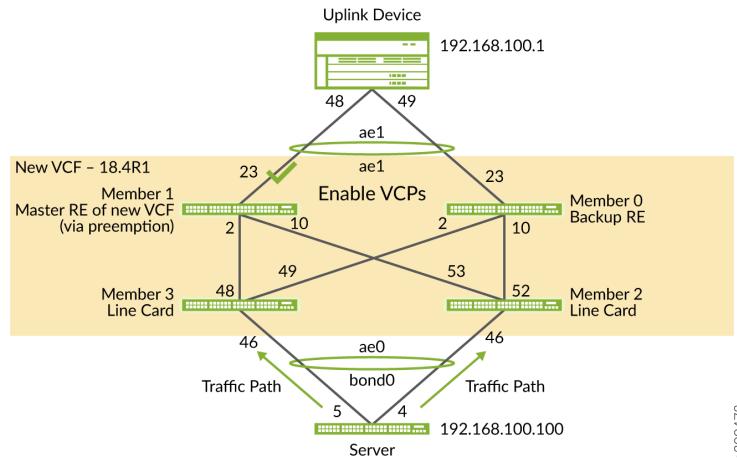
Member 1 has taken over the primary role. This may disrupt the traffic flow. If you observe this, quickly enable the uplink shown in the next step.

6. On Member 1, enable the uplink-facing interface et-1/0/23 on the new VCF.

```
user@switch# delete interfaces et-1/0/23 disable
```

You have now formed a four-member VCF, as shown in [Figure 5 on page 21](#).

Figure 5: Restore the Four-Member VCF



7. Expect less than a minute of traffic disruption as the LACP states reset when Member 1 joins the new VCF. Monitor the ongoing ping from the server.

```

64 bytes from 192.168.100.1: icmp_seq=9910 ttl=64 time=1.84 ms
64 bytes from 192.168.100.1: icmp_seq=9911 ttl=64 time=6.83 ms
64 bytes from 192.168.100.1: icmp_seq=9912 ttl=64 time=0.938 ms
64 bytes from 192.168.100.1: icmp_seq=9913 ttl=64 time=9.03 ms
64 bytes from 192.168.100.1: icmp_seq=9914 ttl=64 time=7.84 ms
From 192.168.100.100 icmp_seq=9954 Destination Host Unreachable
...

```

8. Since Member 1 has re-taken primary role, check that the uplink and server-facing interfaces were not automatically disabled because of a residual configuration. Run the following commands on Member 1 and check that the LACP child interfaces are up and back in the collecting distributing state.

```

{master:1}
user@switch> show lacp interfaces

Aggregated interface: ae0
  LACP state:      Role   Exp   Def   Dist   Col   Syn   Aggr   Timeout   Activity
  xe-2/0/46      Actor   No    No    Yes   Yes   Yes   Yes    Fast    Active

```

xe-2/0/46	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-3/0/46	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
xe-3/0/46	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
LACP protocol: Receive State Transmit State Mux State									
xe-2/0/46		Current		Fast	periodic	Collecting	distributing		
xe-3/0/46		Current		Fast	periodic	Collecting	distributing		
Aggregated interface: ae1									
LACP state: Role Exp Def Dist Col Syn Aggr Timeout Activity									
et-1/0/23	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
et-1/0/23	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
et-0/0/23	Actor	No	No	Yes	Yes	Yes	Yes	Fast	Active
et-0/0/23	Partner	No	No	Yes	Yes	Yes	Yes	Fast	Active
LACP protocol: Receive State Transmit State Mux State									
et-1/0/23		Current		Fast	periodic	Collecting	distributing		
et-0/0/23		Current		Fast	periodic	Collecting	distributing		

```
{master:1}
user@switch> show interfaces terse | match ae
```

```
et-0/0/23.0      up    up    aenet    --> ae1.0
et-1/0/23.0      up    up    aenet    --> ae1.0
xe-2/0/46.0      up    up    aenet    --> ae0.0
xe-3/0/46.0      up    up    aenet    --> ae0.0
ae0              up    up
ae0.0            up    up    eth-switch
ae1              up    up
ae1.0            up    up    eth-switch
```

9. On Member 1, the new primary device of the VCF, confirm that all VCF members are running the intended Junos OS release.

```
{master:1}
user@switch> show version

fpc0:
-----
Hostname: switch
Model: qfx5100-24q-2p
Junos: 18.4R1.8
```

```
JUNOS Base OS Software Suite [18.4R1.8]
```

```
JUNOS Base OS boot [18.4R1.8]
```

```
...
```

```
fpc1:
```

```
-----
```

```
Hostname: switch
```

```
Model: qfx5100-24q-2p
```

```
Junos: 18.4R1.8
```

```
JUNOS Base OS Software Suite [18.4R1.8]
```

```
JUNOS Base OS boot [18.4R1.8]
```

```
...
```

```
fpc2:
```

```
-----
```

```
Hostname: switch
```

```
Model: qfx5100-48s-6q
```

```
Junos: 18.4R1.8
```

```
JUNOS Base OS Software Suite [18.4R1.8]
```

```
JUNOS Base OS boot [18.4R1.8]
```

```
...
```

```
fpc3:
```

```
-----
```

```
Hostname: switch
```

```
Model: qfx5100-48s-6q
```

```
Junos: 18.4R1.8
```

```
JUNOS Base OS Software Suite [18.4R1.8]
```

```
JUNOS Base OS boot [18.4R1.8]
```

```
...
```

10. On the ongoing ping, verify that traffic from the server is flowing normally through the VCF. Expect a downtime of 40-50 seconds.

```
64 bytes from 192.168.100.1: icmp_seq=10057 ttl=64 time=4.91 ms
```

```
64 bytes from 192.168.100.1: icmp_seq=10058 ttl=64 time=4.44 ms
```

```
64 bytes from 192.168.100.1: icmp_seq=10059 ttl=64 time=10.9 ms
```

```
...
```

```
64 bytes from 192.168.100.1: icmp_seq=10074 ttl=64 time=6.98 ms
```

```
64 bytes from 192.168.100.1: icmp_seq=10075 ttl=64 time=4.94 ms
```

```
^C
```

```
--- 192.168.100.1 ping statistics ---
```

```
10075 packets transmitted, 9970 received, +39 errors, 1% packet loss, time 10089682ms
rtt min/avg/max/mdev = 0.261/7.778/151.226/10.579 ms, pipe 3
```

The traffic is flowing normally through the VCF. Your four-member VCF is upgraded and fully functional.

Conclusion

This procedure outlines one of the recommended ways to upgrade an entire VCF with minimal impact to data center workloads when NSSU is not available or not desirable.