Virtual Port Channel Software Upgrade Technical Note

Part Number: OL-29017-01
Publication Date: March 2013

This document describes the procedure, best practices, and known behaviors that are associated with upgrading or downgrading a virtual port channel (vPC) domain that includes vPC peer switches in an environment of Cisco Nexus switches with dual supervisor modules and single supervisor modules.

We recommend that you use an In-Service Software Upgrade (ISSU) or In-Service Software Downgrade (ISSD) to upgrade or downgrade a vPC domain. A nondisruptive software upgrade through an ISSU or downgrade through ISSD is possible when each of the peer switches has two supervisor modules. When a Cisco Nexus switch has a single supervisor module, traffic disruption might occur.

The primary focus of this document is to describe how to upgrade or downgrade a vPC domain when the recommended ISSU or ISSD procedure cannot be followed because there are many software releases between the current and the target software releases, or because the Cisco Nexus switches do not have dual supervisor modules.

Intended Audience

The intended audience for this document includes network administrators, network operations engineers, and network engineers who wish to perform software upgrades in a vPC domain with the least possible disruption to the production network.

Document Scope

This document describes how to upgrade or downgrade software in a vPC domain. Information about the vPC functionality or operation are beyond the scope of this document. In addition, this document does not address hardware upgrades.

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Virtual Port Channel Overview

A vPC allows links that are physically connected to two different Cisco Nexus devices to appear as a single port channel to a third device. The third device can be a Cisco Nexus 2000 Series Fabric Extender or a switch, server, or any other networking device.

A vPC can provide Layer 2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes and load-balancing traffic where alternative paths exist.

The vPC domain includes both vPC peer devices, the vPC peer-keepalive link, the vPC peer link, and all the port channels in the vPC domain that are connected to the downstream device. You can have only one vPC domain ID on each device.

After the two Cisco Nexus switches are connected into a vPC domain, they appear as one logical entity (switch) to the third device (typically, an access switch or a host). As a result, the access switch or host can establish a port channel with this one logical switch instead of using two different physical Cisco Nexus switches.

Figure 1 Physical Topology and Logical Topology

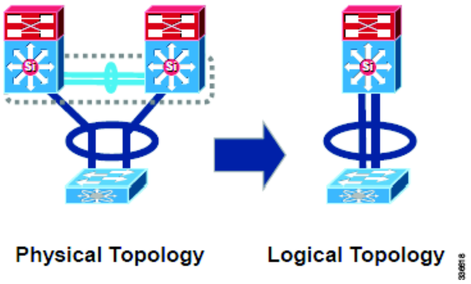

The benefits of using a vPC include the following:

- Enables you to use a single device that can use a port channel across two upstream devices.
- Eliminates Spanning Tree Protocol (STP)-blocked ports.
- Provides a loop-free topology.
- Uses all available uplink bandwidth.
- Enable fast convergence if either the link or a device fails.
- Ensures high availability.
- Provides link-level resiliency.

Upgrading or Downgrading a vPC Domain

Some of the most common reasons for upgrading or downgrading software in a vPC domain are as follows:

- Support for new hardware.
- Support for new software features.
- Improved software reliability through bug fixes.
- Deployment of software for a new installation.

- Maintenance of a uniform software version across the network.

Methods to Upgrade or Downgrade a vPC Domain

You can perform a software upgrade or downgrade in a vPC domain by using one of these methods:

- Use the Cisco recommended ISSU or ISSD procedure.
- Use a manual vPC upgrade/downgrade procedure.

ISSU Requirements

To provide a fully nondisruptive upgrade, you should know or follow these requirements:

- You must have a stable control plane on both peers, which means that you cannot make a network, topology, or configuration change while an ISSU is in progress.
- You cannot change configuration settings or network connections during an upgrade. Any changes in the network settings might cause a disruptive upgrade.
- The system needs to have dual supervisors.
- The configuration mode is blocked during an ISSU to prevent any change.
- The configured features on the system should be compatible and supported with the target release.
- The standby supervisor bootflash needs to have sufficient space to accept the target image.
- The specified system and kickstart images must be compatible.
- There must not be any power disruption while the upgrade is in progress.
- Do not add or remove a module or supervisor while an upgrade is in progress.
- The current and target releases should be fully ISSU compatible. Any software releases that are not available for download from Cisco.com, such as engineering or beta releases, are not supported for an ISSU.

Other ISSU failure conditions are described in the *Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide.*

If any of the preceding requirements are not met, an ISSU cannot be used for a nondisruptive upgrade. Instead, you must perform a manual upgrade of the vPC peers.

To determine if the current and target software releases support an ISSU or ISSD, see the Upgrade/Downgrade Caveats section listed in the release notes for the target software version. This section has a table that lists the releases that support a nondisruptive ISSU to or ISSD from the current release. Releases that are not listed for a particular release train do not support a direct ISSU to or ISSD from the current release.

Release Notes for Cisco Nexus 7000 Series switches are available here:

http://www.cisco.com/en/US/products/ps9402/prod_release_notes_list.html

For more information about the ISSU process, see the *Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide* for the desired release, available here:

http://www.cisco.com/en/US/products/ps9402/prod_installation_guides_list.html

In-Service Software Upgrade

In a Cisco Nexus 7000 Series switch chassis with dual supervisors, you can perform an in-service software upgrade (ISSU) or in-service software downgrade (ISSD) to upgrade or downgrade the system software while the system continues to forward traffic. An ISSU or ISSD uses nonstop forwarding (NSF) with a stateful switchover (SSO) to perform the software upgrade with no system downtime.

Note In this document, the term ISSU is used to represent both ISSU and ISSD unless stated otherwise.

An ISSU is initiated through the command-line interface (CLI). When initiated, an ISSU updates (as needed) the following components on the system:

- Supervisor BIOS, kickstart image, and system image
- Module BIOS and image
- Connectivity Management Processor (CMP) BIOS and image

In a redundant system with two supervisors, one of the supervisors is active while the other operates in the standby mode. During an ISSU, the new software is loaded onto the standby supervisor while the active supervisor continues to operate using the old software. As part of the upgrade, a switchover occurs between the active and standby supervisors, and the standby supervisor becomes active and begins running the new software. After the switchover, the new software is loaded onto the (formerly active) standby supervisor.

An ISSU-based upgrade is a system-wide upgrade that applies the same image and versions across the entire system, including all configured virtual device contexts (VDCs). VDCs are a control-plane and user-interface virtualization and cannot run independent image versions per virtualized resource.

ISSU in a vPC Environment

A vPC domain can be fully upgraded (or downgraded) nondisruptively without any packet loss using the Cisco recommended ISSU (or ISSD) feature.

When you perform an ISSU for vPC system upgrade, you should serialize the upgrade, which means that you should upgrade the peers one after the other and not simultaneously. In other words, after the first peer fully completes the ISSU process, you should perform an ISSU on the second peer.

It is also important that you do not change any configurations on either of the peers while an ISSU is in progress; however, if you change the vPC configuration during the upgrade, it causes an inconsistency between the vPC peer devices (the one being upgraded and the other device) and the ISSU fails.

In addition, an ISSU is not supported when the vPC peer link is down. An ISSU with a vPC requires a fully functional and active vPC system.

To avoid this undesirable situation, a vPC peer implements an automatic lock of the configuration on the device that is not undergoing the upgrade and releases the lock when the upgrade is complete.

vPC ISSU Process Steps

The steps for an ISSU upgrade procedure are as follows. In this example, the upgrade is from Cisco NX-OS Release 5.2.x software to Release 6.0.x software using an ISSU.

To perform an ISSU, follow these steps:

Step 1 Establish console connections to both the active and standby supervisor modules.

Step 2 Copy the running configuration for all VDCs to the startup configuration and to an external device for backup purposes by entering the **copy running-config** *destination-path* **vdc all** command.

Step 3 Determine the preupgrade hardware and software compatibility check. This step should be performed before you enter the **install all** command.

Enter the **show install-all impact** command to display hardware and software information.

If you discover any compatibility issues, you must address these issues before you proceed to step 4.

After entering the **install all** command on the first switch (7K1), the system performs a compatibility check. (See the figure in Step 5.) After the check finishes, the system shows the result of the compatibility check and shows if there is any possible impact of the ISSU upgrade (disruptive/nondisruptive) and prompts you to continue.

Step 4 Enter yes or no after verifying the results of the preupgrade compatibility check.

Note Starting with Cisco NX-OS Release 5.2(1), you can upgrade multiple line cards simultaneously by using the *parallel* argument with the **install all** command. A parallel upgrade decreases the ISSU time; an ISSU upgrade that is done serially (one line card at a time) takes longer.

To start a parallel upgrade, enter the i**nstall all kickstart image system image** *parallel* command.

If the parallel upgrade is not desired, omit the *parallel* argument and the system will use the default serial upgrade method (one line card at a time).

Step 5 If you answer yes answer to the preceding prompt, a configuration lock is implemented on the other peer switch (7K2) and the code upgrade to the Release 6.0(x) software on first switch (7K1) progresses.

A log message similar to the following appears:

```
'%VPC-2-VPC_ISSU_START: Peer vPC switch ISSU start, locking configuration"
```

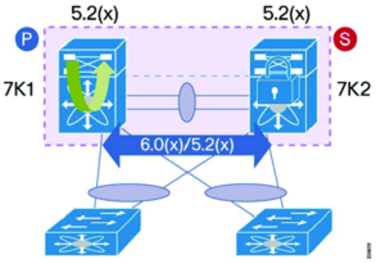

Step 6 After the first vPC peer switch (7K1) is upgraded successfully to Release 6.0(x) software, the configuration lock is released and a log message similar to the following appears:

```
"%VPC-2-VPC_ISSU_END: Peer vPC switch ISSU end, unlocking configuration"
```

Step 7 Repeat Steps 1 through 6 on the second switch (7K2). The upgrade on the second peer switch (7K2) to Release 6.0(x) software is performed while the configuration lock is implemented on switch 7K1.

During this phase, the vPC system can run with different software versions on each of the peers without any problems.

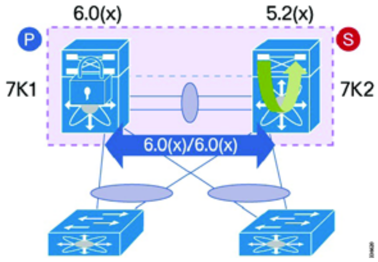

Step 8 The software upgrade to Release 6.0(x) software on switch 7K2 finishes and the configuration lock is released from switch 7K1.

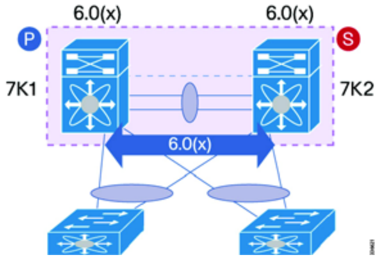

Step 9 The ISSU process is complete. After the upgrade process completes on both peers, they have the same vPC roles (primary/secondary) as they had before the upgrade process started. In other words, the ISSU upgrade does not involve a role change.

Traditional (Manual) Reload for a vPC Domain Upgrade

This section describes how to manually upgrade the software.

Reasons for a Manual Upgrade

The most common reasons for performing a manual upgrade include the following:

- The system has a single supervisor.
- The current and target releases are not ISSU compatible. See the previous section for more details about compatible releases.
- A parallel upgrade is required to reduce the time that is taken for the upgrade. A parallel upgrade is an upgrade of up to three modules at the same time. This type of upgrade with an ISSU was not supported in releases earlier than Cisco NX-OS Release 5.2(x). Therefore, if you are using a software release that is earlier than Cisco NX-OS Release 5.2(x), consider using a manual upgrade to upgrade all modules at the same time. More information about a parallel upgrade can be found in the *Cisco Nexus 7000 Series NX-OS Software Upgrade and Downgrade Guide, Release 6.x.*
- Other time constraints.
- The user wants to upgrade the vPC system while the vPC peer link is down.
- The Cisco Nexus 7000 Series switches are running software releases that are not available on Cisco.com, such as an engineering release or a beta release.

Upgrading a vPC Domain with the Traditional Reload Method

When you use a traditional reload method to upgrade or downgrade your vPC domain, convergence times can be expected.

Before You Begin

Before you start the procedure, review the following information so that you know what to expect:

- Expect and plan for some downtime during the traditional reload upgrade process. Layer 3 protocols converge during the reload and bringup process. Exact convergence times depend on the traffic flow, scale, and configuration. We recommend that you perform the traditional reload upgrade process during a time when downtime can be tolerated. If a nondisruptive software upgrade is desired, use the ISSU process.
- Devices that are connected to orphan ports lose connectivity for the duration of the reload process for the respective switch to which they are single-homed.
- We recommend that you use various vPC enhancements whenever possible to optimize the vPC domain in your environment. Review all vPC enhancements individually to determine if they apply to your environment.
  - Delay restore (default)
  - Graceful consistency check (default)
  - ARP synchronization
  - Auto-recovery, and auto-recovery reload delay
  - Peer gateway
  - PIM prebuild (not supported on F2 Series modules)

  You can find additional enhancements and information in the following guides:

  - *Design and Configuration Guide: Best Practices for Virtual Port Channels (vPC) on Cisco Nexus 7000 Series Switches*

    http://www.cisco.com/en/US/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf
  - *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*

    http://www.cisco.com/en/US/partner/products/ps9402/products_installation_and_configuration_guides_list.html
- vPC does not support role preemption.
- If you use a different procedure for your specific environment, it is possible that the vPC peers end up in the roles of "primary, operational secondary" and "secondary, operational primary." These states are fully functional and the vPC peers can be left in that state. There is no need to alter those roles.
- You do not need to alter the role priority on either vPC peer during the upgrade process. Altering the role priority does not affect the system until the peer link is flapped. If altering the role priority is desired, be aware that altering the role priority and then flapping the vPC peer link causes vPCs to reinitialize, can possibly trigger STP root convergence, and result in traffic loss. The following warning message appears when you attempt this type of change.

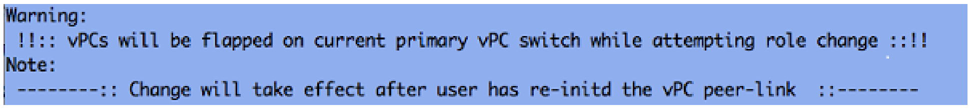

```
Warning:
 !!:: vPCs will be flapped on current primary vPC switch while attempting role change ::!!
Note:
 --------:: Change will take effect after user has re-initd the vPC peer-link  ::--------
```

- When you are bringing up a vPC, the VLANs on the vPC peer link go to a suspended state while the peers negotiate the allowed VLANs. Only the VLANs on the vPC peer link are suspended. All other ports should not alter the STP state, and you should not alter the traffic flow through the vPCs.
- If possible, move all single-homed devices that are connected through orphan ports to the device that is not being upgraded. Devices that are connected to orphan ports on the vPC switch being upgraded lose connectivity for the duration of the upgrade.
- Save the running configuration for all VDCs to bootflash and to an external device.

  ```
  switch# copy running-config destination vdc-all
  ```

- Power down or remove any unsupported line cards (in cases of downgrading).

  ```
  switch# poweroff module x
  ```

  ```
  switch# poweroff xbar x
  ```

- Have a console connection to the supervisor module(s).

Procedure

To perform a manual upgrade, follow these steps:

Step 1 Log in to the vPC primary switch.

Note There is no functional difference between starting with the primary switch or the secondary switch. However, by beginning with the primary switch, you can ensure that both devices end up in their originally configured primary and secondary roles. There is no functional difference between the role that the switches are in at the end of the upgrade.

Step 2 Configure the boot variable for the Cisco NX-OS kickstart and NX-OS system images.

```
switch# boot kickstart bootflash:n7000-s1-kickstart.6.1.1.bin
switch# boot system bootflash:n7000-s1-dk9.6.1.1.bin
```

After setting the boot variables, the active supervisor checks to verify that the image is also present on the standby supervisor. If it is not present, the following message appears:

```
VDC-1 %$ %BOOTVAR-5-IMAGE_NOTEXISTS: Warning: image <filename> doesn't exist on sup2
```

The active supervisor copies the image to the standby supervisor. Manual intervention is not required; the active supervisor starts the process.

Step 3 Verify that the image was copied to the standby supervisor by entering the **show boot auto-copy list** command.

When the auto-copy process completes, the following message appears:

```
%BOOTVAR-5-AUTOCOPY_SUCCEED: auto-copy of file /bootflash/<filename> to standby supervisor
succeed
```

Step 4 Confirm the images are on the standby supervisor.

```
switch# dir bootflash://sup-standby/
  208476721    Aug 16 18:09:19 2012  n7000-s1-dk9.6.1.1.bin
   30035968    Aug 16 18:04:19 2012  n7000-s1-kickstart.6.1.1.bin
```

Step 5 Save the running configuration with the new boot variables.

```
switch# copy running-config startup-config vdc-all
```

Step 6 Verify that the Current Boot Variables and the Boot Variables on the next reload match the expected image.

```
switch# show boot
Current Boot Variables:
sup-1
kickstart variable = bootflash:/n7000-s1-kickstart.6.1.1.bin
system variable = bootflash:/n7000-s1-dk9.6.1.1.bin
sup-2
kickstart variable = bootflash:/n7000-s1-kickstart.6.1.1.bin
system variable = bootflash:/n7000-s1-dk9.6.1.1.bin
No module boot variable set
Boot Variables on next reload:
sup-1
kickstart variable = bootflash:/n7000-s1-kickstart.6.1.1.bin
system variable = bootflash:/n7000-s1-dk9.6.1.1.bin
sup-2
kickstart variable = bootflash:/n7000-s1-kickstart.6.1.1.bin
system variable = bootflash:/n7000-s1-dk9.6.1.1.bin
No module boot variable set
```

Step 7 After verifying that the image is on the bootflash of both supervisor modules, reload the configured primary vPC peer. The configured vPC secondary peer should take over as the “secondary, operational primary,” and any Layer 3 protocols will converge.

Step 8 Verify that the configured vPC secondary takes over this role by entering the **show vpc brief** command. The vPC peer link will be down, but the vPCs will remain up on the “secondary, operational primary.” This situation allows traffic to continue to flow through the remaining vPCs, using approximately half of the total bandwidth previously available before the configured primary vPC reload.

Note Any orphan ports that have connectivity through the switch that is reloaded lose connectivity for the duration of the reload process.

Step 9 After the vPC peer link is established during the bootup process for the upgraded peer, the “secondary, operational primary” role is maintained for the upgraded peer. There is no preemptive action in the vPC roles, and therefore the configured primary vPC peer that is brought up on the new version of code takes the role of “primary, operational secondary” after the bootup completes. This role assignment occurs regardless of what vPC role priority is configured on the vPC peers.

Step 10 Repeat Steps 2 through 9 for the “secondary, operational primary” vPC switch.

Step 11 After bringing up of the secondary vPC peer, both switches should be in their original state of primary and secondary vPC peer and running the new version of code.

vPC Upgrade Procedure FAQs

1. **Should I expect any downtime during the traditional reload upgrade process.**?

You should expect and plan for downtime during the traditional reload upgrade process. Layer 3 protocols converge during the reload and bringup process; the time varies depending on the traffic flow, scale, timers, and other configuration.

2. **Are there any specific vPC features that I should enable before I start the reload upgrade process?**

We recommend that you use various vPC enhancements whenever possible to optimize the vPC domain in a specific environment. Review all vPC enhancements individually to determine if they are applicable to your environment.

- Delay restore (default)
- Graceful consistency check (default)
- ARP synchronization
- Auto-recovery, and auto-recovery reload delay
- Peer gateway
- PIM prebuild (not supported on F2/F2e Series modules)

You can find additional information in the following guides:

- *Design and Configuration Guide: Best Practices for Virtual Port Channels (vPC) on Cisco Nexus 7000 Series Switches*

  http://www.cisco.com/en/US/docs/switches/datacenter/sw/design/vpc_design/vpc_best_practices_design_guide.pdf
- *Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide*

  http://www.cisco.com/en/US/partner/products/ps9402/products_installation_and_configuration_guides_list.html

3. **What roles will my vPC peers be in after the upgrade process is complete?**

If you follow the suggested procedure, both devices should end up in their original vPC primary and secondary roles. However, if you use a different procedure, it is possible that your devices can end up in “secondary, operational primary” and “primary, operational secondary” roles. There is no functional impact if the vPC devices end up in these roles and they can remain in these roles and be fully functional.

4. **After my upgrade is completed, why are my vPC peer devices in “secondary, operational primary” and “primary, operational secondary” roles?**

The roles at the procedure depend on the roles that the vPC devices start in, and the upgrade procedure that you used. There is no preemption support in vPC role election. After the bootup of the original primary vPC peer, the vPC peer that is now “secondary, operational primary” is considered “sticky” and to minimize role changes, it does not give up its operational primary role. Again, there is no functional impact while the vPC peer devices are in these roles, and you do not need to take extra downtime to alter these roles.

5. **Can I get my vPC peer devices into the original “Primary” and “Secondary” roles?**

Because there is no preemption support as of this writing, we suggest that you perform the recommended upgrade procedure to get the devices into their original primary and secondary roles. You can also reload the “secondary, operational primary” vPC peer one additional time to get the devices into the original roles. Instead of an entire reload, you can reconfigure the role priorities and then flap the vPC peer link. Flapping the vPC peer link does not return the devices to their original roles. Review the “Before You Begin” section on page 7 that discusses altering the role priority and flapping the peer link.

6. **Should I change the role priority, and if so, does this change cause any issues?**

vPC role priorities do not need to be altered at any point in the process. Changing the vPC role priority on either device does not take effect until the vPC peer link is flapped. This action does cause more downtime.

7. **Should I shut down any links prior to the upgrade process?**

You do not need to shut down any links before you reload either vPC peer unless the tested procedure is approved for your specific network design.

8. **Are there any other caveats I should be aware of while manually upgrading the vPC system using traditional reload method?**

There is an existing caveat when Cisco Fabric Extender (FEX) modules are connected to a Cisco Nexus 7000 Series switch. When the switch is manually upgraded, FEX host interfaces (HIFs) lose the configuration after the reload of the Cisco Nexus 7000 Series switch.

The configuration loss occurs because the system reapplies (replays) the FEX configuration as soon as the Cisco Nexus 7000 Series switch comes up, without waiting for the FEX to come online. As a result, the FEX HIF configuration is rejected because the system does not see the FEX online and the FEX interfaces do not appear to be available.

This issue occurs only when the Cisco Nexus 7000 Series switch is upgraded or downgraded manually and does not occur when the Cisco Nexus 7000 Series switch is reloaded with the same NX-OS software release. This issue does not occur when the ISSU or ISSD procedure is used.

The Cisco Nexus 7000 Series FEX preprovisioning feature will fix this issue in a future Cisco NX-OS software release. Until the issue is fixed, if you are manually upgrading the vPC system, you must save the FEX HIF (FEX host interfaces connected to hosts) configurations to both the startup configuration file and to an external device before starting the reload, and reapply the configuration once the FEX module is fully online.

9. **How do I introduce an isolated device back into the vPC system?**

Check the Link Aggregation Control Protocol (LACP) roles on both peers. If the roles are same, disable the auto-recovery using the **no auto-recovery** command in vPC domain on both peers, and reload the isolated device. The isolated device will have the **none established** LACP role and can be introduced into the vPC system without LACP role re-election. You can find additional information about LACP roles in *Nexus 7000 Chassis Replacement Procedure*.

Conclusion

You should use the Cisco recommended ISSU or ISSD procedure when upgrading or downgrading a Cisco Nexus 7000 Series vPC domain. You must meet certain conditions before you can perform the ISSU/ISSD procedure described in this document.

If you use the manual or traditional reload method to upgrade the vPC domain, follow the step-by-step procedure provided in this document and be aware of the listed caveats.

The manual upgrade procedure documented in this Tech Note has been tested and validated by the Cisco Cisco Nexus 7000 Series quality assurance group.

Obtain Documentation and Submit a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. The RSS feeds are a free service.