

Configure In-Service Software Upgrade on Cisco Nexus 9000 and 3100 Series Switches

Overview

Cisco® In-Service Software Upgrade (ISSU) service allows you to upgrade your device software while the switch continues to forward traffic. ISSU reduces or eliminates the downtime typically caused by software upgrades. Cisco NX-OS Software Release 7.0(3)I3(1) and later supports **non-disruptive** ISSU upgrade on Cisco Nexus® 9000 and 3100 Series Switches running in standalone mode. The default upgrade process is disruptive, so ISSU needs to be enabled using the command-line interface (CLI), as described in the configuration section of this document. Use of the non-disruptive option helps ensure a non-disruptive upgrade. The guest shell is disabled during the ISSU process and then reactivated after the upgrade.

When you perform an ISSU process, some Layer 2 and Layer 3 protocols will extend their values to accommodate the upgrade. For example, Unidirectional Link Detection (UDLD) and Bidirectional Forwarding Detection (BFD) will increase their hello timers so that adjacency is maintained during the ISSU process. Also, Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), Intermediate-System - Intermediate-System (ISIS) and Enhanced Interior Gateway Routing Protocol (EIGRP) perform graceful restart. Because an allocated time is needed for ISSU to successfully complete, aggressive timers are not supported for any Layer 2 and Layer 3 protocols. For example, Link Aggregation Control Protocol (LACP) fast timer or OSPF fast timer is not supported. For Layer 2 and Layer 3 protocols with sensitive timers, the timeout value should be increased. For applications for which you cannot increase the timeout value, the upgrade will be disruptive.

NX-OS version 7.0(3)I4(1) introduces support for ISSU on FEX, VXLAN segment routing and Network Address Translation (NAT). At the time of this writing, the following features are not supported for ISSU: OpenFlow and Layer 2 and Layer 3 aggressive timers. If the user wants to use ISSU and these features are already enabled on the system, the user can either perform a disruptive upgrade or disable the unsupported features and proceed with a non-disruptive upgrade. Features that do not support ISSU should warn the user when ISSU is triggered.

Hardware That Supports ISSU

The following Cisco Nexus 3100 and 9000 Series Switches support ISSU:

- Cisco Nexus 3164Q and 31128PQ Switches
- Cisco Nexus 9332PQ, 9372PX, 9372PX-E, 9372TX, 9372TX-E, 9396PX, 9396TX, 93120TX, and 93128TX Switches
- Cisco Nexus 9504, 9508, and 9516 Switches with Cisco X9432PQ, X9464PX, X9464TX, X9536PQ, X9564PX, X9564TX, or X9636PQ line cards, dual supervisor modules, and a minimum of two system controllers and two fabric modules

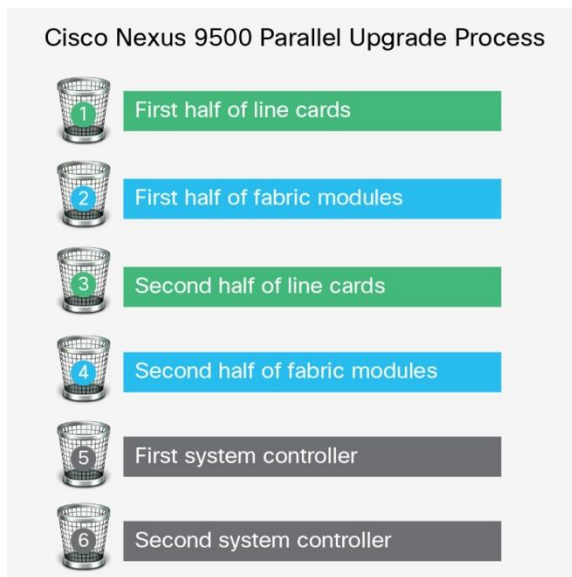
Cisco Nexus 9500 Platform

End-of-row (EoR) Cisco Nexus 9500 platform switches are modular switches that require two supervisors for ISSU. The minimum configuration required is two system controllers and two fabric modules.

Modular Cisco Nexus 9000 Series Switches support parallel upgrade as the default method. The parallel method upgrades the modules in batches instead of one after the other for a faster upgrade. In the upgrade sequence, first the supervisors are upgraded (requires switchover), then the line cards, fabric modules, system controllers and the FEX. After switchover is performed in a parallel upgrade, the secondary supervisor takes over, and the installer determines the current line cards and fabric modules. It then divides the components into buckets. It places the first half of the line cards in the first bucket, the first half of the fabric modules in the second bucket, the second half of line cards in the third bucket, the second half of the fabric modules in the fourth bucket, the first system controller in the fifth bucket, and the second system controller in the sixth bucket (Figure 1.1). Each bucket is upgraded successfully before the upgrade process starts for the next bucket. The console shows which modules are assigned to which bucket and status of the upgrade. The user has the option to choose a serial upgrade using the CLI.

Note: The minimum requirement for a modular Cisco Nexus 9000 Series Switch undergoing ISSU is two supervisors, two system controllers and two fabric modules. The Cisco Nexus 9400 linecards can have a partially connected fabric module. In this case, if only two fabric modules are used with the 9400 linecards, the fabric modules should not be in slots 21 and 25. They should be in slots 22, 23, 24, or 26. This allows for the system to maintain high availability during ISSU.

Figure 1. Nexus 9500 Parallel Upgrade Process



Cisco Nexus 9300 Platform and 3100 Series

The ToR Cisco Nexus 9300 platform switches and Cisco Nexus 3100 Series Switches are standalone switches with single supervisors. ISSU on the Cisco Nexus 9000 and 3100 Series switch causes the supervisor CPU to reset and load the new software version. The control plane is inactive in this duration, but the data plane keeps forwarding packets leading to an upgrade with no service disruption. After the CPU loads the updated version of NX-OS, the system restores the control plane to previous known configuration and runtime state and gets in-sync with the data plane, thereby completing the ISSU process. Since the data plane kept forwarding packets while the control plane was upgraded, any servers connected to the Cisco Nexus 9000 and 3100 Series switch access layer should see no traffic disruption.

When an upgrade is performed on TOR the control plane is reset. This reset will cause spanning tree to time out to its neighboring devices, which will result in a spanning-tree topology change. In addition, if the switch undergoing ISSU is the spanning-tree root, it may not be able to send Bridge Protocol Data Units (BPDUs) during the ISSU process. If there are downstream switches connected to devices undergoing ISSU, as a best practice, you should use a virtual port channel (vPC) design for ISSU. vPC offers the advantage of running even if the two vPC peer devices operate with different NX-OS releases. For instance, vPC peer device 1 can run NX-OS 7.0(3)I3(1), and the other vPC peer device can NX-OS 7.0(3)I4(x). During the transition phase, the vPC continues to work even if the peer devices use different NX-OS code. Configuration lock during ISSU prevents synchronous upgrades on both vPC peer devices simultaneously (configuration is automatically locked on the other vPC peer device when ISSU is initiated). This feature enables support for non-disruptive upgrade for the vPC domain.

When the spanning-tree primary switch undergoes ISSU, it notifies the spanning-tree secondary switch so that it can tune its spanning-tree timers. If the ToR switches are spanning-tree root switches, the peer switch should be enabled. The peer switch allows both devices to share a common bridge ID when sending BPDUs. Because the peer switch is enabled, the spanning-tree secondary switch will continue to send BPDUs to its connected devices to avoid a spanning-tree topology change while the spanning-tree primary switch undergoes ISSU.


ISSU on TOR best practices:

- vPC best practices include the following:
 - If TOR is spanning-tree root bridge ensure peer-switch is enabled
 - If TOR is not root bridge, ensure all the ports are edge/edge trunk ports
 - Known caveat: CSCuy12559. Has been fixed in 7.0(3)I3(2)
- The ToR switch performs a supervisor restart so that graceful restart is used for BGP, OSPF, ISIS and EIGRP routing protocols.
- BFD and UDLD timers are automatically adjusted.
- During the ISSU process on a TOR, all FHRPs will cause the other peer to become active if the node undergoing ISSU is active. If the failover is not desired during ISSU and no new hosts need to be serviced during the down control plane, the FHRP hello interval can be set to at least 120 seconds. Note: If BFD is enabled for FHRP and a failover occurs, BFD will detect the failure.

Design Considerations

Cisco Nexus 9000 Series Switches can be used in a number of designs in a data center. Cisco Nexus 9000 Series Switches can be used as end-of-row or top-of-rack access layer switches, as aggregation or core switches in traditional hierarchical two- or three-tier network designs, or deployed in a modern leaf-spine architecture. Below are examples of common Nexus 9000 and 3100 Series designs that are fully compatible with ISSU.


Figure 2. Traditional Design with Cisco vPC



Layer 2 ISSU Best Practices:

- There should not be any STP topology change during ISSU
- Bridge assurance (BA) should not be active on any port (except Peer Link)
- TOR vPC best practices include the following:
 - If TOR is spanning-tree root bridge ensure peer-switch is enabled
 - If TOR is not root bridge, ensure all the ports are edge/edge trunk ports
 - Known caveat: CSCuy12559. Has been fixed in 7.0(3)I3(2)
- If resilient hashing is configured, the upgrade would be disruptive
- Aggressive timers for L2 protocols are not supported during ISSU


Figure 3. Two-Tier Routed Access Layer Design



Layer 3 ISSU Best Practices:


- BFD will automatically increase its timer so the adjacency is maintained during the ISSU process.
- Graceful Restart for BGP, OSPF, ISIS and EIGRP is enabled by default. ISSU warns to be disruptive if one or more BGP and OSPF Peer does not have Graceful restart enabled.
- ISSU is assumed to be in a stable topology (i.e. we don't expect link flaps).
- For EOR platform, ISSU will warn if the negotiated hold time is smaller than the maximum expected switchover time. In TOR we adjust the timers automatically to a higher value.
- ISSU aborts in case of OSPF/ISIS topology changes occur during ISSU.
- Aggressive timers for L3 protocols are not supported during ISSU.

Figure 4. Fabric Extender Design



The Nexus 9000 EOR and TOR support Cisco FEX to connect to one parent switch. Servers should be dual-homed to two different fabric extenders. During ISSU, any servers connected to FEX should see no traffic disruption.

Figure 5. Two-Tier Design Leaf-Spine and VXLAN Overlay



VXLAN is an industry-standard protocol and uses underlay IP networks. It extends Layer 2 segments over a Layer 3 infrastructure to build Layer 2 overlay logical networks. Follow the same ISSU best practices for the Layer 2 or Layer 3 underlay configuration that is used in VXLAN.

Configuration Overview

You cannot configure a device during an upgrade. Use this command to verify that you have no active configuration sessions:

show configuration session summary

Use this command on TOR to validate the spanning-tree configuration for ISSU:

show spanning-tree issu-impact

Check the impact of upgrading the software before actually performing the upgrade:

show install all impact nxos bootflash:nxos.7.0.3.I2.1.bin

Then begin the upgrade process:

install all nxos bootflash:nxos.7.0.3.I2.1.bin

The following options are available:

- **No-reload:** This option exits the software upgrade process before the device is reloaded.
- **Non-disruptive:** This option performs ISSU to prevent the disruption of data traffic. (By default, the software upgrade process is disruptive.)
- **Non-interruptive:** This option upgrades the software without any prompts. This option skips all error and sanity checks.

-
- **Serial:** This option upgrades the I/O modules in Cisco Nexus 9500 platform switches one at a time. (By default, the I/O modules are upgraded in parallel, which reduces the overall upgrade time. Specifically, the I/O modules are upgraded in parallel in this order: the first half of the line cards and fabric modules, then the second half of the line cards and fabric modules, then the first system controller, and then the second system controller.)

Recommendations

The following procedures are strongly recommended:

- Use ISSU to change the NX-OS code release for the vPC domain. Perform the operation sequentially, one vPC peer device at a time.
- Refer to the NX-OS release notes to select the correct target NX-OS code release based on the running code.
- Carefully check the NX-OS release notes and follow the recommended guidelines to perform ISSU successfully.
- **Note:** The current NX-OS release for Nexus 9000 and 3100 platforms does not support non-disruptive ISSD (in service software downgrade).

For More Information

For General Software Upgrade and Downgrade process including pre-requisites, please refer to the following link:

http://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/upgrade/guide/b_Cisco_Nexus_9000_Series_NX-OS_Software_Upgrade_and_Downgrade_Guide_Release_7x/b_Cisco_Nexus_9000_Series_NX-OS_Software_Upgrade_and_Downgrade_Guide_Release_7x_chapter_010.html#topic_9135155FAFF94AC380960BCB748A443B



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)