



High Availability Configuration Guide, Cisco IOS XE Fuji 16.8.x (Catalyst 9500 Switches)

First Published: 2018-04-06

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883



CONTENTS

CHAPTER 1

Configuring Nonstop Forwarding with Stateful Switchover 1

- Prerequisites for Nonstop Forwarding with Stateful Switchover 1
- Restrictions for Cisco Nonstop Forwarding with Stateful Switchover 1
- Information About NSF with SSO 2
 - Overview of Nonstop Forwarding with Stateful Switchover 2
 - SSO Operation 3
 - NSF Operation 3
 - Cisco Express Forwarding 4
 - Routing Protocols 4
 - BGP Operation 4
 - EIGRP Operation 5
 - OSPF Operation 6
- How to Configure Cisco NSF with SSO 6
 - Configuring SSO 6
- Configuration Examples for Nonstop Forwarding with Stateful Switchover 7
 - Example: Configuring SSO 7
 - Verifying Cisco Express Forwarding with NSF 8
- Additional References for Nonstop Forwarding with Stateful Switchover 9
- Feature History Information for Nonstop Forwarding with Stateful Switchover 9

CHAPTER 2

Configuring Cisco StackWise Virtual 11

- Finding Feature Information 11
- Prerequisites for Cisco StackWise Virtual 11
- Restrictions for Cisco StackWise Virtual 12
- Information About Cisco StackWise Virtual 13
 - StackWise Virtual Overview 13

Cisco StackWise Virtual Topology	13
Cisco StackWise Virtual Redundancy	15
SSO Redundancy	15
Nonstop Forwarding	16
Multichassis EtherChannels	16
MEC Minimum Latency Load Balancing	16
MEC Failure Scenarios	17
Cisco StackWise Virtual Packet Handling	17
Traffic on a StackWise Virtual link	18
Layer 2 Protocols	18
Layer 3 Protocols	19
Dual-Active Detection	20
Dual-Active-Detection Link with Fast Hello	21
Dual-Active Detection with enhanced PAgP	21
Recovery Actions	22
Implementing Cisco StackWise Virtual	22
How to Configure Cisco StackWise Virtual	23
Configuring Cisco StackWise Virtual Settings	23
Configuring Cisco StackWise Virtual Link	24
Configuring StackWise Virtual Fast Hello Dual-Active-Detection Link	26
Enabling ePAgP Dual-Active-Detection	27
Disabling Cisco StackWise Virtual	29
Verifying Cisco StackWise Virtual Configuration	31
Additional References for StackWise Virtual	31
Feature History for Cisco StackWise Virtual	32

CHAPTER 3

Configuring Graceful Insertion and Removal	33
Restrictions for Graceful Insertion and Removal	33
Information about Graceful Insertion and Removal	33
Overview	33
Layer 2 interface shutdown	34
Custom Template	34
How to Configure Graceful Insertion and Removal	34
Creating maintenance template	34

Configuring System Mode Maintenance	35
Starting and Stopping Maintenance Mode	36
Configuration Examples for Graceful Removal and Insertion	36
Example: Configuring maintenance template	36
Example: Configuring system mode maintenance	37
Example: Starting and Stopping maintenance mode	37
Example: Displaying system mode settings	37
Monitoring Graceful Insertion and Removal	38
Additional References for Graceful Insertion and Removal	38
Feature History and Information for GIR	39



CHAPTER 1

Configuring Nonstop Forwarding with Stateful Switchover

- [Prerequisites for Nonstop Forwarding with Stateful Switchover, on page 1](#)
- [Restrictions for Cisco Nonstop Forwarding with Stateful Switchover, on page 1](#)
- [Information About NSF with SSO, on page 2](#)
- [How to Configure Cisco NSF with SSO, on page 6](#)
- [Configuration Examples for Nonstop Forwarding with Stateful Switchover, on page 7](#)
- [Additional References for Nonstop Forwarding with Stateful Switchover, on page 9](#)
- [Feature History Information for Nonstop Forwarding with Stateful Switchover, on page 9](#)

Prerequisites for Nonstop Forwarding with Stateful Switchover

- NSF must be configured on a networking device that has been configured for SSO.
- Border Gateway Protocol (BGP) support in NSF requires that neighbor networking devices be NSF-aware; that is, devices must have the graceful restart capability and advertise that capability in their OPEN message during session establishment. If an NSF-capable device discovers that a particular BGP neighbor does not have graceful restart capability, it does not establish an NSF-capable session with that neighbor. All other neighbors that have graceful restart capability continue to have NSF-capable sessions with this NSF-capable networking device.
- Open Shortest Path First (OSPF) support in NSF requires that all neighbor networking devices be NSF-aware. If an NSF-capable device discovers that it has non-NSF-aware neighbors on a particular network segment, it disables NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware devices continue to provide NSF capabilities.

Restrictions for Cisco Nonstop Forwarding with Stateful Switchover

The following are restrictions for configuring NSF with SSO:

- For NSF operation, you must have SSO configured on the device.

- All Layer 3 neighboring devices must be an NSF helper or NSF-capable to support graceful restart capability.
- For IETF, all neighboring devices must be running an NSF-aware software image.
- The Hot Standby Routing Protocol (HSRP) is not supported with NSF SSO.
- An NSF-aware device cannot support two NSF-capable peers performing an NSF restart operation at the same time. However, both neighbors can reestablish peering sessions after the NSF restart operation is complete.

Information About NSF with SSO

Overview of Nonstop Forwarding with Stateful Switchover

Cisco Nonstop Forwarding (NSF) works with the Stateful Switchover (SSO) feature. The device supports fault resistance by allowing a standby switch to take over if the active device becomes unavailable. NSF works with SSO to minimize the amount of time a network is unavailable.

Usually, when a networking device restarts, all routing peers of that device detect that the device went down and then came back up. This transition results in what is called a routing flap, which could spread across multiple routing domains. Routing flaps caused by routing restarts create routing instabilities, which are detrimental to the overall network performance. Cisco NSF helps to suppress routing flaps in SSO-enabled devices, thus reducing network instability.

Cisco NSF with SSO allows for the forwarding of data packets to continue along known routes while the routing protocol information is being restored following a switchover. With NSF/SSO, peer networking devices do not experience routing flaps. Data traffic is forwarded through intelligent line cards or dual forwarding processors (FPs) while the standby router processor (RP) assumes control from the failed active RP during a switchover. NSF with SSO operation provides the ability of line cards and FPs to remain active through a switchover and to be kept current with the Forwarding Information Base (FIB) on the active RP.

NSF provides the following benefits:

- Improved network availability—NSF continues forwarding network traffic and application state information so that user session information is maintained after a switchover.
- Overall network stability—Network stability can be improved with the reduction in the number of route flaps that are created when devices in the network fail, and lose their routing tables.
- Neighboring devices do not detect a link flap—Because interfaces remain active during a switchover, neighboring devices do not detect a link flap (the link does not go down and come back up).
- Prevents routing flaps—Because SSO continues forwarding network traffic during a switchover, routing flaps are avoided.
- Maintains user sessions established prior to the switchover.
- If the standby device does not respond, a new standby device is elected as the standby.
- If the active device does not respond, the standby device becomes the active device.
- If a stack member does not respond, that member is removed from the stack.

- If the standby device does not respond, a new standby device is elected.
- If the active device does not respond, the standby device becomes the active device.

SSO Operation

When a standby device runs in SSO mode, the standby device starts up in a fully-initialized state and synchronizes with the persistent configuration and the running configuration on the active device. It subsequently maintains the state of the protocols, and all changes in hardware and software states for features that support SSO are kept in synchronization. Consequently, it offers minimum interruption to Layer 2 sessions in a redundant active device configuration.

If the active device fails, the standby device becomes the active device. This new active device uses existing Layer 2 switching information to continue forwarding traffic. Layer 3 forwarding is delayed until routing tables are repopulated in the newly active device.

NSF Operation

NSF always runs with SSO, and provides redundancy for Layer 3 traffic. NSF is supported by BGP, Enhanced Interior Gateway Routing Protocol (EIGRP), and OSPF routing protocols and also by Cisco Express Forwarding for forwarding. These routing protocols have been enhanced with NSF-capability and awareness, which means that devices running these protocols can detect a switchover and take necessary actions to continue forwarding network traffic and to recover route information from peer devices.

Each protocol depends on Cisco Express Forwarding to continue forwarding packets during switchover while routing protocols rebuild the Routing Information Base (RIB) tables. After the convergence of routing protocols, Cisco Express Forwarding updates the FIB table and removes stale route entries. Cisco Express Forwarding then updates the hardware with the new FIB information.

If the active device is configured (with the **graceful-restart** command) for BGP, OSPF, or EIGRP routing protocols, routing updates are automatically sent during the active device election.

NSF has two primary components:

- **NSF-aware:** A networking device is NSF-aware if it is running NSF-compatible software. If neighboring devices detect that an NSF device can still forward packets when an active device election happens, this capability is referred to as NSF-awareness. Enhancements to the Layer 3 routing protocols (BGP, OSPF, and EIGRP) are designed to prevent route-flapping so that the Cisco Express Forwarding routing table does not time out or the NSF device does not drop routes. An NSF-aware device helps to send routing protocol information to the neighboring NSF device. NSF-awareness is enabled by default for EIGRP-stub, EIGRP, and OSPF protocols. NSF-awareness is disabled by default for BGP.
- **NSF-capability:** A device is NSF-capable if it is configured to support NSF; it rebuilds routing information from NSF-aware or NSF-capable neighbors. NSF works with SSO to minimize the amount of time that a Layer 3 network is unavailable following an active device election by continuing to forward IP packets. Reconvergence of Layer 3 routing protocols (BGP, OSPFv2, and EIGRP) is transparent to the user and happens automatically in the background. Routing protocols recover routing information from neighbor devices and rebuild the Cisco Express Forwarding table.

Cisco Express Forwarding

A key element of Cisco IOS Nonstop Forwarding (NSF) is packet forwarding. In a Cisco networking device, packet forwarding is provided by Cisco Express Forwarding. Cisco Express Forwarding maintains the Forwarding Information Base (FIB), and uses the FIB information that is current at the time of a switchover to continue forwarding packets during a switchover, to reduce traffic interruption during the switchover.

During normal NSF operation, Cisco Express Forwarding on the active device synchronizes its current FIB and adjacency databases with the FIB and adjacency databases on the standby device. Upon switchover, the standby device initially has FIB and adjacency databases that are mirror images of those that were current on the active device. Cisco Express Forwarding keeps the forwarding engine on the standby device current with changes that are sent to it by Cisco Express Forwarding on the active device. The forwarding engine can continue forwarding after a switchover as soon as the interfaces and a data path are available.

As the routing protocols start to repopulate the RIB on a prefix-by-prefix basis, the updates cause prefix-by-prefix updates to Cisco Express Forwarding, which it uses to update the FIB and adjacency databases. Existing and new entries receive the new version (“epoch”) number, indicating that they have been refreshed. The forwarding information is updated on the forwarding engine during convergence. The device signals when the RIB has converged. The software removes all FIB and adjacency entries that have an epoch older than the current switchover epoch. The FIB now represents the newest routing protocol forwarding information.

Routing Protocols

Routing protocols run only on the active RP, and receive routing updates from neighbor devices. Routing protocols do not run on the standby RP. Following a switchover, routing protocols request that the NSF-aware neighbor devices send state information to help rebuild routing tables. Alternately, the Intermediate System-to-Intermediate System (IS-IS) protocol can be configured to synchronize state information from the active to the standby RP to help rebuild the routing table on the NSF-capable device in environments where neighbor devices are not NSF-aware.

**Note**

For NSF operation, routing protocols depend on Cisco Express Forwarding to continue forwarding packets while routing protocols rebuild the routing information.

BGP Operation

When a NSF-capable device begins a BGP session with a BGP peer, it sends an OPEN message to the peer. Included in the message is a declaration that the NSF-capable device has “graceful restart capability.” Graceful restart is the mechanism by which BGP routing peers avoid a routing flap following a switchover. If the BGP peer has this capability, it is aware that the device sending the message is NSF-capable. Both the NSF-capable device and its BGP peer(s) need to exchange the Graceful Restart Capability in their OPEN messages, at the time of session establishment. If both peers do not exchange the Graceful Restart Capability, the session is not graceful restart capable.

If the BGP session is lost during the RP switchover, the NSF-aware BGP peer marks all routes associated with the NSF-capable device as stale; however, it continues to use these routes to make forwarding decisions for a set period of time. This functionality means that no packets are lost while the newly active RP is waiting for convergence of the routing information with the BGP peers.

After an RP switchover occurs, the NSF-capable device reestablishes the session with the BGP peer. In establishing the new session, it sends a new graceful restart message that identifies the NSF-capable device as having restarted.

At this point, the routing information is exchanged between two BGP peers. Once this exchange is complete, the NSF-capable device uses the routing information to update the RIB and the FIB with the new forwarding information. The NSF-aware device uses the network information to remove stale routes from its BGP table. Following that, the BGP protocol is fully converged.

If a BGP peer does not support the graceful restart capability, it will ignore the graceful-restart capability in an OPEN message; but will establish a BGP session with the NSF-capable device. This function allows interoperability with non-NSF-aware BGP peers (and without NSF functionality), but the BGP session with non-NSF-aware BGP peers will not be graceful restart capable.



Note BGP support in NSF requires that neighbor networking devices be NSF-aware; that is, devices must have the Graceful Restart Capability and advertise that capability in their OPEN message during session establishment. If an NSF-capable device discovers that a particular BGP neighbor does not have Graceful Restart Capability, it will not establish an NSF-capable session with that neighbor. All other neighbors that have Graceful Restart Capability will continue to have NSF-capable sessions with this NSF-capable networking device.

EIGRP Operation

Enhanced Interior Gateway Routing Protocol (EIGRP) NSF capabilities are exchanged by EIGRP peers in hello packets. The NSF-capable device notifies its neighbors that an NSF restart operation has started by setting the restart (RS) bit in a hello packet. When an NSF-aware device receives notification from an NSF-capable neighbor that an NSF-restart operation is in progress, the NSF-capable and NSF-aware devices immediately exchange their topology tables. The NSF-aware device sends an end-of-table update packet when the transmission of its topology table is complete. The NSF-aware device then performs the following actions to assist the NSF-capable device:

- The EIGRP hello hold timer is expired to reduce the time interval set for hello packet generation and transmission. This allows the NSF-aware device to reply to the NSF-capable device more quickly reducing the amount of time required for the NSF-capable device to rediscover neighbors and rebuild the topology table.
- The route-hold timer is started. This timer is used to set the period of time that the NSF-aware device will hold known routes for the NSF-capable neighbor. This timer is configured with the **timers nsf route-hold** command. The default time period is 240 seconds.
- In the peer list, the NSF-aware device notes that the NSF-capable neighbor is restarting, maintains adjacency, and holds known routes for the NSF-capable neighbor until the neighbor signals that it is ready for the NSF-aware device to send its topology table, or the route-hold timer expires. If the route-hold timer expires on the NSF-aware device, the NSF-aware device discards held routes and treats the NSF-capable device as a new device joining the network and reestablishes adjacency accordingly.
- The NSF-aware device continues to send queries to the NSF-capable device which is still in the process of converging after a switchover, effectively extending the time before a stuck-in-active condition can occur.

When the switchover operation is complete, the NSF-capable device notifies its neighbors that it has reconverged and has received all of their topology tables by sending an end-of-table update packet to assisting devices. The NSF-capable device then returns to normal operation. The NSF-aware device will look for alternate paths

(go active) for any routes that are not refreshed by the NSF-capable (restarting device). The NSF-aware device will then return to normal operation. If all paths are refreshed by the NSF-capable device, the NSF-aware device will immediately return to normal operation.



Note NSF-aware devices are completely compatible with non-NSF aware or -capable neighbors in an EIGRP network. A non-NSF aware neighbor will ignore NSF capabilities and reset adjacencies and otherwise maintain the peering sessions normally.

OSPF Operation

When an OSPF NSF-capable device performs a supervisor engine switchover, it must perform the following tasks in order to resynchronize its link state database with its OSPF neighbors:

- Relearn the available OSPF neighbors on the network without causing a reset of the neighbor relationship.
- Reacquire the contents of the link state database for the network.

As quickly as possible after a supervisor engine switchover, the NSF-capable device sends an OSPF NSF signal to neighboring NSF-aware devices. Neighbor networking devices recognize this signal as an indicator that the neighbor relationship with this device should not be reset. As the NSF-capable device receives signals from other devices on the network, it can begin to rebuild its neighbor list.

After neighbor relationships are reestablished, the NSF-capable device begins to resynchronize its database with all of its NSF-aware neighbors. At this point, the routing information is exchanged between the OSPF neighbors. Once this exchange is complete, the NSF-capable device uses the routing information to remove stale routes, update the RIB, and update the FIB with the new forwarding information. The OSPF protocols are then fully converged.



Note OSPF support in NSF requires that all neighbor networking devices be NSF-aware. If an NSF-capable device discovers that it has non-NSF -aware neighbors on a particular network segment, it disables NSF capabilities for that segment. Other network segments composed entirely of NSF-capable or NSF-aware devices continue to provide NSF capabilities.

How to Configure Cisco NSF with SSO

Configuring SSO

You must configure SSO in order to use NSF with any supported protocol.

Procedure

	Command or Action	Purpose
Step 1	enable Example:	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.

	Command or Action	Purpose
	Device> enable	
Step 2	show redundancy states Example: Device# show redundancy states	Displays the operating redundancy mode.
Step 3	redundancy Example: Device(config)# redundancy	Enters redundancy configuration mode.
Step 4	mode sso Example: Device(config-red)# mode sso	Configures stateful switchover. • When this command is entered, the standby switch is reloaded and begins to work in SSO mode.
Step 5	end Example: Device(config-red)# end	Exits redundancy configuration mode and returns to privileged EXEC mode.
Step 6	show redundancy states Example: Device# show redundancy states	Displays the operating redundancy mode.
Step 7	debug redundancy status Example: Device# debug redundancy status	Enables the debugging of redundancy status events.

Configuration Examples for Nonstop Forwarding with Stateful Switchover

Example: Configuring SSO

This example shows how to configure the system for SSO and displays the redundancy state:

```
Device(config)# redundancy
Device(config-red)# mode sso
Device(config-red)# end
Device#
```

The following is sample output from the **show redundancy states** command:

```
show redundancy states
my state = 13 -ACTIVE
peer state = 8 -STANDBY HOT
Mode = Duplex
```

```

Unit = Primary
Unit ID = 5
Redundancy Mode (Operational) = sso
Redundancy Mode (Configured) = sso
Split Mode = Disabled
Manual Swact = Enabled
Communications = Up
client count = 29
client_notification TMR = 30000 milliseconds
keep_alive TMR = 9000 milliseconds
keep_alive count = 1
keep_alive threshold = 18
RF debug mask = 0x0

```

Verifying Cisco Express Forwarding with NSF

Procedure

show cef state

Displays the state of Cisco Express Forwarding on a networking device.

Example:

```
Device# show cef state
```

```

CEF Status:
RP instance
common CEF enabled
IPv4 CEF Status:
CEF enabled/running
dCEF enabled/running
CEF switching enabled/running
universal per-destination load sharing algorithm, id DEA83012
IPv6 CEF Status:
CEF disabled/not running
dCEF disabled/not running
universal per-destination load sharing algorithm, id DEA83012
RRP state:
I am standby RRP: no
RF Peer Presence: yes
RF PeerComm reached: yes
RF Progression blocked: never
Redundancy mode: rpr(1)
CEF NSF sync: disabled/not running
CEF ISSU Status:
FIBHWIDB broker
No slots are ISSU capable.
FIBIDB broker
No slots are ISSU capable.
FIBHWIDB Subblock broker
No slots are ISSU capable.
FIBIDB Subblock broker
No slots are ISSU capable.
Adjacency update
No slots are ISSU capable.
IPv4 table broker
No slots are ISSU capable.
CEF push

```

No slots are ISSU capable.

Additional References for Nonstop Forwarding with Stateful Switchover

Related Documents

Related Topic	Document Title
For complete syntax and usage information for the commands used in this chapter.	Catalyst 9400 Command Reference

MIBs

MIB	MIBs Link
All the supported MIBs for this release.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/support

Feature History Information for Nonstop Forwarding with Stateful Switchover

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use the Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Nonstop Forwarding with Stateful Switchover

Feature Name	Release	Feature Information
Nonstop Forwarding with Stateful Switchover	Cisco IOS XE Everest 16.6.1	Cisco NSF works with the SSO feature. NSF works with SSO to minimize the amount of time a network is unavailable to users following a switchover. The main objective of NSF SSO is to continue forwarding IP packets following a Route Processor (RP) switchover.



CHAPTER 2

Configuring Cisco StackWise Virtual

- [Finding Feature Information, on page 11](#)
- [Prerequisites for Cisco StackWise Virtual, on page 11](#)
- [Restrictions for Cisco StackWise Virtual, on page 12](#)
- [Information About Cisco StackWise Virtual, on page 13](#)
- [How to Configure Cisco StackWise Virtual, on page 23](#)
- [Verifying Cisco StackWise Virtual Configuration, on page 31](#)
- [Additional References for StackWise Virtual, on page 31](#)
- [Feature History for Cisco StackWise Virtual, on page 32](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see Bug Search Tool and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.

Prerequisites for Cisco StackWise Virtual

- Both switches in the Cisco StackWise Virtual pair must be directly connected to each other.
- Both switches in the Cisco StackWise Virtual pair must be of the same switch model.
- Both switches in the Cisco StackWise Virtual pair must be running the same license level.
- Both switches in the Cisco StackWise Virtual pair must be running the same software version.
- Both switches in the Cisco StackWise Virtual pair must be running the same SDM template.
- All the ports used for configuring a StackWise Virtual Link (SVL) must share the same speed. For example, you cannot configure a 10G or a 40G port to form an SVL, simultaneously.

Restrictions for Cisco StackWise Virtual

- The Federal Information Processing Standards (FIPS) is not supported on Cisco StackWise Virtual links.
- Cisco StackWise Virtual is supported only on the following models:
 - C9500-24Q
 - C9500-12Q
 - C9500-40X
 - C9500-16X
- Cisco StackWise Virtual links can be configured only on the following ports:

Table 2: Switch Models and Supported Ports

Switch Model	Supported Ports
C9500-24Q	23, 24
C9500-12Q	11, 12
C9500-40X	37, 38, 39, 40
C9500-16X	13, 14, 15, 16

- Cisco StackWise Virtual configuration commands will be recognised only on a switch running Network Advantage license. The configuration commands will not be recognised on a Network Essentials license.
- When configuring StackWiseVirtual links (SVLs) on 9500-40X and C9500-16X models of the Cisco Catalyst 9500 Series Switches, note that you cannot create SVLs on any of the network modules.
- 4x10G breakout cables are not supported with SVLs.
- When deploying Cisco StackWise Virtual, ensure that VLAN ID 4094 is not used anywhere on the network. All inter-chassis system control communication between stack members is carried over the reserved VLAN ID 4094 from the global range.
- The dual active and StackWise Virtual link configuration must be performed manually and the device should be rebooted for the configuration changes to take effect.
- Only Cisco Transceiver Modules are supported.
- The interface VLAN MAC address that is assigned by default, can be overridden using the **mac-address** command. If this command is configured on a single SVI or router port that requires Layer 3 injected packets, all other SVIs or routed ports on the device also must be configured with the same first four most significant bits (4MSB) of the MAC address. For example, if you set the MAC address of any SVI to xxxx.yyyy.zzzz, set the MAC address of all other SVIs to start with xxxx.yyyy. If Layer 3 injected packets are not used, this restriction does not apply.



Note This applies to all Layer 3 ports, SVIs, and routed ports. This does not apply to GigabitEthernet0/0 port.

Information About Cisco StackWise Virtual

StackWise Virtual Overview

Cisco StackWise Virtual is a network system virtualization technology that pairs two switches into one virtual switch. Switches in a Cisco StackWise Virtual solution increase operational efficiency by using single control and management plane, scale system bandwidth with distributed forwarding plane, and assist in building resilient networks using the recommended network design. Cisco StackWise Virtual allows two physical switches to operate as a single logical virtual switch using a 40G or 10G Ethernet connection.

Cisco StackWise Virtual Topology

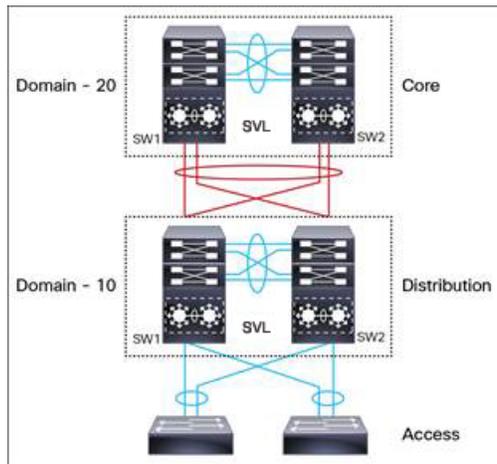
A typical network design consists of core, distribution, and access layers. The default mode of a switch is standalone. When two redundant switches are deployed in the distribution layer, the following network challenges arise:

- If VLAN IDs are reused between access layers then, it will introduce a spanning tree loop that will impact the overall performance of the network.
- Spanning tree protocols and configuration are required to protect Layer 2 network against spanning tree protocol loop, and root and bridge protocol data unit management.
- Additional protocols such as first hop redundancy protocol are required to virtualize the IP gateway function. This should align with STP root priorities for each VLAN.
- The Protocol independent multicast designated router (PIM DR) configuration should be fine-tuned to selectively build a multicast forwarding topology on a VLAN.
- The standalone distribution layer system provides protocol-driven remote failure and detection, which results in slower convergence time. Fine-tune FHRP and PIM timers for rapid fault detection and recovery process.

We recommend Cisco StackWise Virtual model for aggregation layers and collapsed aggregation and core layers. The stack can be formed over a redundant 40G or 10G fiber links to ensure that the distribution or the aggregation switches can be deployed over a large distance.

Note that STP keeps one of the ports connected to the distribution switches blocked on the access switches. As a result of this, an active link failure causes STP convergence and the network suffers from traffic loss, flooding, and a possible transient loop in the network. On the other hand, if the switches are logically merged into one switch, all the access switches might form an EtherChannel bundle with distribution switches, and a link failure within an EtherChannel would not have any impact as long as at least one member within the EtherChannel is active.

Figure 1: Typical Network Design using Cisco StackWise Virtual



Etherchannel in StackWise Virtual is capable of implementing Multi-chassis EtherChannel (MEC) across the stack members. When access layer and aggregation layer are collapsed into a single StackWise Virtual system, MEC across the different access layer domain members and across distribution and access layer switches will not be supported. MEC is designed to forward the traffic over the local link irrespective of the hash result.

Since the control plane, management plane, and data plane are integrated, the system behaves as a single switch.

The virtualization of multiple physical switches into a single logical switch is from a control and management plane perspective only. Because of the control plane being common, it may look like a single logical entity to peer switches. The data plane of the switches is distributed. Each switch is capable of forwarding over its local interfaces without involving other members. However, when a packet coming into a switch has to be forwarded over a different member's port, the forwarding context of the packet is carried over to the destination switch after ingress processing is performed in the ingress switch. Egress processing is done only in the egress switch. This provides a uniform data plane behavior to the entire switch irrespective whether of the destination port is in a local switch or in a remote switch. However, the common control plane ensures that all the switches have equivalent data plane entry for each forwarding entity.

An election mechanism elects one of the switches to be Cisco StackWise Virtual active and the other switch to be Cisco StackWise Virtual standby in terms of Control Plane functions. The active switch is responsible for all the management, bridging and routing protocols, and software data path. The standby switch is in hot standby state ready to take over the role of active, if the active switch fails over.

The following are the components of the Cisco StackWise Virtual solution:

- Stack members
- StackWise Virtual link: 10G or 40G Ethernet connections. SVL is established using the 10G or 40G interfaces on the supported switch models. However, a combination of both the interfaces is not supported.

StackWise Virtual link is the link that connects the switches over Ethernet. Typically, Cisco StackWise Virtual consists of multiple 10-G or 40-G physical links. It carries all the control and data traffic between the switching units. You can configure a StackWise Virtual link on a supported port. When a switch is powered up and the hardware is initialized, it looks for a configured StackWise Virtual link before the initialization of the control plane.

The Link Management Protocol (LMP) is activated on each link of the StackWise Virtual links as soon as the links are established. LMP ensure the integrity of SVL links and monitors and maintains the health of the

links. The redundancy role of each switch is resolved by the StackWise Discovery Protocol (SDP). It ensures that the hardware and software versions are compatible to form the SVL and determines which switch becomes active or standby from a control plane perspective.

Cisco StackWise Virtual Header (SVH) is 64-byte frame header that is prepended over all control, data, and management plane traffic that traverse over each SVL between the two stack members of the Cisco StackWise Virtual domain. The SVH-encapsulated traffic operates at OSI Layer 2 and can be recognized and processed only by Cisco StackWise Virtual-enabled switches. SVL interfaces are non-bridgeable and non-routeable, and allows non-routeable traffic over L2 or L3 network.

Cisco StackWise Virtual Redundancy

Cisco StackWise Virtual operates stateful switchover (SSO) between the active and standby switches. The following are the ways in which Cisco StackWise Virtual's redundancy model differs from that of the standalone mode:

- The Cisco StackWise Virtual active and standby switches are hosted in separate switches and use a StackWise Virtual link to exchange information.
- The active switch controls both the switches of Cisco StackWise Virtual. The active switch runs the Layer 2 and Layer 3 control protocols and manages the switching modules of both the switches.
- The Cisco StackWise Virtual active and standby switches perform data traffic forwarding.



Note If the Cisco StackWise Virtual active switch fails, the standby switch initiates a switchover and assumes the Cisco StackWise Virtual active switch role.

SSO Redundancy

A StackWise Virtual system operates with SSO redundancy if it meets the following requirements:

- Both the switches must be running the same software version, unless they are in the process of software upgrade.
- StackWise Virtual link-related configuration in the two switches must match.
- License type must be same on both the switch models.
- Both the switch models must be in the same StackWise Virtual domain.

With SSO redundancy, the StackWise Virtual standby switch is always ready to assume control if a fault occurs on the StackWise Virtual active switch. Configuration, forwarding, and state information are synchronized from the StackWise Virtual active switch to the redundant switch at startup, and whenever changes to the StackWise Virtual active switch configuration occur. If a switchover occurs, traffic disruption is minimized.

If StackWise Virtual does not meet the requirements for SSO redundancy, it will be incapable of establishing a relationship with the peer switch. StackWise Virtual runs stateful switchover (SSO) between the StackWise Virtual active and standby switches. The StackWise Virtual determines the role of each switch during initialization.

The CPU in the StackWise Virtual standby switch runs in hot standby state. StackWise Virtual uses a StackWise Virtual link to synchronize configuration data from the StackWise Virtual active switch to the StackWise Virtual standby switch. Also, protocols and features that support high availability synchronize their events and state information to the StackWise Virtual standby switch.

Nonstop Forwarding

While implementing Nonstop Forwarding (NSF) technology in systems using SSO redundancy mode, network disruptions are minimized for campus users and applications. High availability is provided even when the control-plane processing stack-member switch is reset. During a failure of the underlying Layer 3, NSF-capable protocols perform graceful network topology resynchronization. The preset forwarding information on the redundant stack-member switch remains intact; this switch continues to forward the data in the network. This service availability significantly lowers the mean time to repair (MTTR) and increases the mean time between failure (MTBF) to achieve a high level of network availability.

Multichassis EtherChannels

Multichassis EtherChannel (MEC) is an EtherChannel bundled with physical ports having common characteristics such as speed and duplex, that are distributed across each Cisco StackWise Virtual system. A Cisco StackWise Virtual MEC can connect to any network element that supports EtherChannel (such as a host, server, router, or switch). Cisco StackWise Virtual supports up to 128 MECs deployed in Layer 2 or Layer 3 modes. EtherChannel 128 is reserved for SVL connections. Hence, the maximum available MEC count is 127.

In a Cisco StackWise Virtual system, an MEC is an EtherChannel with additional capability. A multichassis EtherChannel link reduces the amount of traffic that requires transmission across the StackWise Virtual link by populating the index port only with the ports local to the physical switch. This allows the switch to give precedence to the local ports of the multichassis EtherChannel link over those on the remote switch.

Each MEC can optionally be configured to support either Cisco PAgP, IEEE LACP, or Static ON mode. We recommend that you implement EtherChannel using Cisco PAgP or LACP with a compatible neighbor. If a remotely connected neighbor such as Cisco Wireless LAN Controller (WLC) does not support this link-bundling protocol, then a Static ON mode can be deployed. These protocols run only on the Cisco StackWise Virtual active switch.

An MEC can support up to eight physical links that can be distributed in any proportion between the Cisco StackWise Virtual active switch and the Cisco StackWise Virtual standby switch. We recommend that you distribute the MEC ports across both switches evenly.

MEC Minimum Latency Load Balancing

The StackWise Virtual environment is designed such that data forwarding always remains within the switch. The Virtual Stack always tries to forward traffic on the locally available links. This is true for both Layer 2 and Layer 3 links. The primary motivation for local forwarding is to avoid unnecessarily sending data traffic over the StackWise Virtual link and thus reduce the latency (extra hop over the SVL) and congestion. The bidirectional traffic is load-shared between the two StackWise Virtual members. However, for each StackWise Virtual member, ingress and egress traffic forwarding is based on locally-attached links that are part of MEC. This local forwarding is a key concept in understanding convergence and fault conditions in a StackWise Virtual enabled campus network.

The active and standby switches support local forwarding that will individually perform the desired lookups and forward the traffic on local links to uplink neighbors. If the destination is a remote switch in the StackWise

Virtual domain, ingress processing is performed on the ingress switch and then traffic is forwarded over the StackWise Virtual link to the egress switch where only egress processing is performed.

MEC Failure Scenarios

We recommend that you configure a MEC with at least one link to each switch. This configuration ensures that there is always an alternate path for data traffic in case of a switch failure.

The following sections describe issues that may arise and the resulting impact:

Single MEC Link Failure

If a link within a MEC fails (and other links in the MEC are still operational), the MEC redistributes the load among the operational links, as in a regular port.

All MEC Links to the Cisco StackWise Virtual Active Switch Fail

If all the links to the Cisco StackWise Virtual active switch fail, a MEC becomes a regular EtherChannel with operational links to the Cisco StackWise Virtual standby switch.

Data traffic that terminates on the Cisco StackWise Virtual active switch reaches the MEC by crossing a StackWise Virtual link to the Cisco StackWise Virtual standby switch. Control protocols continue to run in the Cisco StackWise Virtual active switch. Protocol messages reach the MEC by crossing a StackWise Virtual link.

All MEC Links Fail

If all the links in an MEC fail, the logical interface for the EtherChannel is set to Unavailable. Layer 2 control protocols perform the same corrective action as for a link-down event on a regular EtherChannel.

On adjacent switches, routing protocols and the Spanning Tree Protocol (STP) perform the same corrective action as for a regular EtherChannel.

Cisco StackWise Virtual Standby Switch Failure

If the Cisco StackWise Virtual standby switch fails, a MEC becomes a regular EtherChannel with operational links on the Cisco StackWise Virtual active switch. Connected peer switches detect the link failures, and adjust their load-balancing algorithms to use only the links to the StackWise Virtual active switch.

Cisco StackWise Virtual Active Switch Failure

Cisco StackWise Virtual active switch failure results in a stateful switchover (SSO). After the switchover, a MEC is operational on the new Cisco StackWise Virtual active switch. Connected peer switches detect the link failures (to the failed switch), and adjust their load-balancing algorithms to use only the links to the new Cisco StackWise Virtual active switch.

Cisco StackWise Virtual Packet Handling

In Cisco StackWise Virtual, the Cisco StackWise Virtual active switch runs the Layer 2 and Layer 3 protocols and features and manages the ports on both the switches.

Cisco StackWise Virtual uses StackWise Virtual link to communicate system and protocol information between the peer switches and to carry data traffic between the two switches.

The following sections describe packet handling in Cisco StackWise Virtual.

Traffic on a StackWise Virtual link

A StackWise Virtual link carries data traffic and in-band control traffic between two switches. All the frames that are forwarded over the StackWise Virtual link are encapsulated with a special StackWise Virtual Header (SVH). The SVH adds an overhead of 64 bytes for control and data traffic, which provides information for Cisco StackWise Virtual to forward the packet on the peer switch.

A StackWise Virtual link transports control messages between two switches. Messages include protocol messages that are processed by the Cisco StackWise Virtual active switch, but received or transmitted by interfaces on the Cisco StackWise Virtual standby switch. Control traffic also includes module programming between the Cisco StackWise Virtual active switch and the switching modules on the Cisco StackWise Virtual standby switch.

Cisco StackWise Virtual transmits data traffic over a StackWise Virtual link under the following circumstances:

- Layer 2 traffic flooded over a VLAN (even for dual-homed links).
- Packets processed by software on the Cisco StackWise Virtual active switch where the ingress interface is on the Cisco StackWise Virtual standby switch.
- The packet destination is on the peer switch, as described in the following examples:
 - Traffic within a VLAN where the known destination interface is on the peer switch.
 - Traffic that is replicated for a multicast group and the multicast receivers are on the peer switch.
 - The known unicast destination MAC address is on the peer switch.
 - The packet is a MAC notification frame destined for a port on the peer switch.

A StackWise Virtual link also transports system data, such as NetFlow export data and SNMP data, from the Cisco StackWise Virtual standby switch to the Cisco StackWise Virtual active switch.

Traffic on the StackWise Virtual link is load balanced with the same global hashing algorithms available for EtherChannels (the default algorithm is source-destination IP).

Layer 2 Protocols

The Cisco StackWise Virtual active switch runs the Layer 2 protocols (such as STP and VTP) for the switching modules on both the switches. Protocol messages that are received on the standby switch ports must traverse SVL links to reach the active switch where they are processed. Similarly, protocol messages that are transmitted from the standby switch ports originate on the active switch, and traverse the SVL links to reach the standby ports.

All the Layer 2 protocols in Cisco StackWise Virtual work similarly in standalone mode. The following sections describe the difference in behavior for some protocols in Cisco StackWise Virtual.

Spanning Tree Protocol

The Cisco StackWise Virtual active switch runs the STP. The Cisco StackWise Virtual standby switch redirects the STP BPDUs across a StackWise Virtual link to the StackWise Virtual active switch.

The STP bridge ID is commonly derived from the switch MAC address. To ensure that the bridge ID does not change after a switchover, Cisco StackWise Virtual continues to use the original switch MAC address for the STP Bridge ID.

EtherChannel Control Protocols

Link Aggregation Control Protocol (LACP) and Port Aggregation Protocol (PAgP) packets contain a device identifier. Cisco StackWise Virtual defines a common device identifier for both the switches. Use either PAgP or LACP on Multi EtherChannels instead of mode ON, even if all the three modes are supported.



Note A new PAgP enhancement has been defined for assisting with dual-active scenario detection.

Switched Port Analyzer

Switched Port Analyzer (SPAN) on StackWise Virtual link ports is not supported; SVL ports can be neither a SPAN source, nor a SPAN destination. Cisco StackWise Virtual supports all the SPAN features for non-SVL interfaces. The number of SPAN sessions that are available on Cisco StackWise Virtual matches that on a single switch running in standalone mode.

Private VLANs

Private VLANs on StackWise Virtual work the same way as in standalone mode. The only exception is that the native VLAN on isolated trunk ports must be configured explicitly.

Apart from STP, EtherChannel Control Protocols, SPAN, and private VLANs, the Dynamic Trunking Protocol (DTP), Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), and Unidirectional Link Detection Protocol (UDLD) are the additional Layer 2 control-plane protocols that run over the SVL connections.

Layer 3 Protocols

The Cisco StackWise Virtual active switch runs the Layer 3 protocols and features for the StackWise Virtual. All the Layer 3 protocol packets are sent to and processed by the Cisco StackWise Virtual active switch. Both the member switches perform hardware forwarding for ingress traffic on their interfaces. When software forwarding is required, packets are sent to the Cisco StackWise Virtual active switch for processing.

The same router MAC address assigned by the Cisco StackWise Virtual active switch is used for all the Layer 3 interfaces on both the Cisco StackWise Virtual member switches. After a switchover, the original router MAC address is still used. The router MAC address is chosen based on chassis-mac and is preserved after switchover by default. The following sections describe the Layer 3 protocols for Cisco StackWise Virtual.

IPv4 Unicast

The CPU on the Cisco StackWise Virtual active switch runs the IPv4 routing protocols and performs any required software forwarding. All the routing protocol packets received on the Cisco StackWise Virtual standby switch are redirected to the Cisco StackWise Virtual active switch across the StackWise Virtual link. The Cisco StackWise Virtual active switch generates all the routing protocol packets to be sent out over ports on either of the Cisco StackWise Virtual member switches.

Hardware forwarding is distributed across both members on Cisco StackWise Virtual. The CPU on the Cisco StackWise Virtual active switch sends Forwarding Information Base (FIB) updates to the Cisco StackWise Virtual standby switch, which in turn installs all the routes and adjacencies into hardware.

Packets intended for a local adjacency (reachable by local ports) are forwarded locally on the ingress switch. Packets intended for a remote adjacency (reachable by remote ports) must traverse the StackWise Virtual link.

The CPU on the Cisco StackWise Virtual active switch performs all software forwarding and feature processing (such as fragmentation and Time to Live exceed functions). If a switchover occurs, software forwarding is disrupted until the new Cisco StackWise Virtual active switch obtains the latest Cisco Express Forwarding and other forwarding information.

In virtual switch mode, the requirements to support non-stop forwarding (NSF) match those in the standalone redundant mode of operation.

From a routing peer perspective, Multi-Chassis EtherChannels (MEC) remain operational during a switchover, that is, only the links to the failed switch are down, but the routing adjacencies remain valid.

Cisco StackWise Virtual achieves Layer 3 load balancing over all the paths in the Forwarding Information Base entries, be it local or remote.

IPv6

Cisco StackWise Virtual supports IPv6 unicast and multicast because it is present in the standalone system.

IPv4 Multicast

The IPv4 multicast protocols run on the Cisco StackWise Virtual active switch. Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) protocol packets received on the Cisco StackWise Virtual standby switch are transmitted across a StackWise Virtual link to the StackWise Virtual active switch. The latter generates IGMP and PIM protocol packets to be sent over ports on either of the Cisco StackWise Virtual members.

The Cisco StackWise Virtual active switch synchronizes the Multicast Forwarding Information Base (MFIB) state to the Cisco StackWise Virtual standby switch. On both the member switches, all the multicast routes are loaded in the hardware, with replica expansion table (RET) entries programmed for only local, outgoing interfaces. Both the member switches are capable of performing hardware forwarding.



Note To avoid multicast route changes as a result of a switchover, we recommend that all the links carrying multicast traffic be configured as MEC rather than Equal Cost Multipath (ECMP).

For packets traversing a StackWise Virtual link, all Layer 3 multicast replications occur on the egress switch. If there are multiple receivers on the egress switch, only one packet is replicated and forwarded over the StackWise Virtual link, and then replicated to all the local egress ports.

Software Features

Software features run only on the Cisco StackWise Virtual active switch. Incoming packets to the Cisco StackWise Virtual standby switch that require software processing are sent across a StackWise Virtual link to the Cisco StackWise Virtual active switch.

Dual-Active Detection

If the standby switch detects a complete loss of the StackWise Virtual link, it assumes the active switch has failed and will take over as the active switch. However, if the original Cisco StackWise Virtual active switch is still operational, both the switches will now be Cisco StackWise Virtual active switches. This situation is called a dual-active scenario. This scenario can have adverse effects on network stability because both the switches use the same IP addresses, SSH keys, and STP bridge IDs. Cisco StackWise Virtual detects a

dual-active scenario and takes recovery action. Dual-active-detection link is the dedicated link used to mitigate this.

If a StackWise Virtual link fails, the Cisco StackWise Virtual standby switch cannot determine the state of the Cisco StackWise Virtual active switch. To ensure that switchover occurs without delay, the Cisco StackWise Virtual standby switch assumes that the Cisco StackWise Virtual active switch has failed and initiates switchover to take over the Cisco StackWise Virtual active role. The original Cisco StackWise Virtual active switch enters recovery mode and brings down all the interfaces except the StackWise Virtual link and the management interfaces.

Dual-Active-Detection Link with Fast Hello

To use the dual-active fast hello packet detection method, you must provision a direct ethernet connection between the two Cisco StackWise Virtual switches. You can dedicate up to four links for this purpose.

The two switches periodically exchange special dual-active hello messages containing information about the switch state. If all SVLs fail and a dual-active scenario occurs, each switch recognizes that there is a dual-active scenario from the peer's messages. This initiates recovery actions as described in the [Recovery Actions, on page 22](#) section. If a switch does not receive an expected dual-active fast hello message from the peer before the timer expires, the switch assumes that the link is no longer capable of dual-active detection.



Note Do not use the same port for SVL and DAD link.

Dual-Active Detection with enhanced PAgP

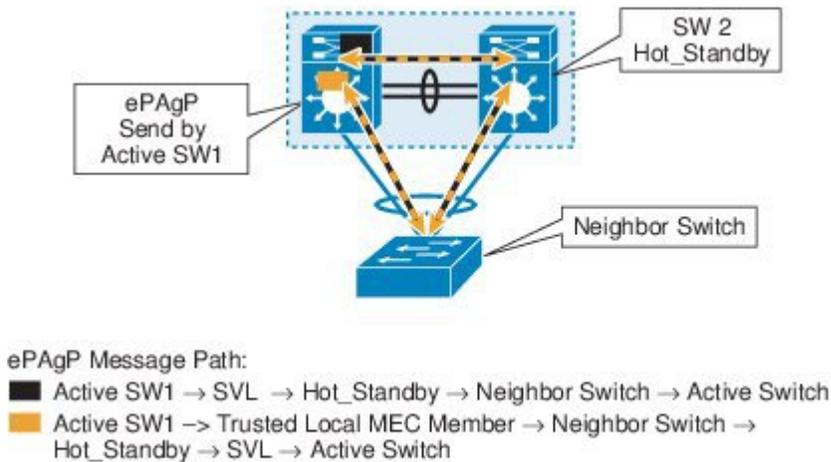
Port aggregation protocol (PAgP) is a Cisco proprietary protocol used for managing EtherChannels. If a StackWise Virtual MEC terminates on a Cisco switch, you can run PAgP protocol on the MEC. If PAgP is running on the MECs between the StackWise Virtual switch and an upstream or downstream switch, the StackWise Virtual can use PAgP to detect a dual-active scenario. The MEC must have at least one port on each switch of the StackWise Virtual setup.

Enhanced PAgP is an extension of the PAgP protocol. In virtual switch mode, ePAgP messages include a new type length value (TLV) which contains the ID of the StackWise Virtual active switch. Only switches in virtual switch mode send the new TLV.

When the StackWise Virtual standby switch detects SVL failure, it initiates SSO and becomes StackWise Virtual active. Subsequent ePAgP messages sent to the connected switch from the newly StackWise Virtual active switch contain the new StackWise Virtual active ID. The connected switch sends ePAgP messages with the new StackWise Virtual active ID to both StackWise Virtual switches.

If the formerly StackWise Virtual active switch is still operational, it detects the dual-active scenario because the StackWise Virtual active ID in the ePAgP messages changes.

Figure 2: Dual-active-detection with ePAgP



Note To avoid PAgP flaps and to ensure that dual-active detection functions as expected, the stack MAC persistent wait timer must be configured as indefinite using the command **stack-mac persistent timer 0**.

Recovery Actions

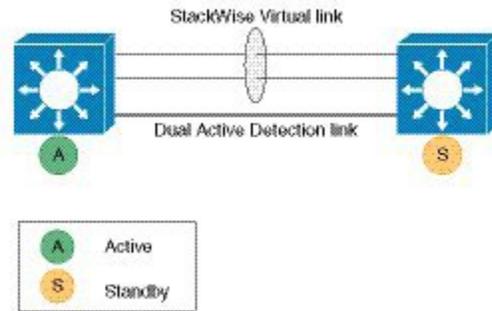
A Cisco StackWise Virtual active switch that detects a dual-active condition shuts down all of its non-SVL interfaces to remove itself from the network. The switch then waits in recovery mode until the SVLs have been recovered. You should physically repair the SVL failure and the recovery switch should be manually reloaded for it to be the standby switch.

Implementing Cisco StackWise Virtual

The two-node solution of Cisco StackWise Virtual is normally deployed at the aggregation layer. Two switches are connected over an SVL.

Cisco StackWise Virtual combines the two switches into a single logical switch with a large number of ports, offering a single point of management. One of the member switches is the active and works as the control and management plane, while the other one is the standby. The virtualization of multiple physical switches into a single logical switch is only from a control and management perspective. Because of the control plane being common, it may look like a single logical entity to peer switches. The data plane of the switches are converged, that is, the forwarding context of a switch might be passed to the other member switch for further processing when traffic is forwarded across the switches. However, the common control plane ensures that all the switches have equivalent data plane entry for each forwarding entity.

Figure 3: Two-Node Solution



An election mechanism that determines which switch is Cisco StackWise Virtual active and which one is a control plane standby, is available. The active switch is responsible for management, bridging and routing protocols, and software data path. These are centralized on the active switch supervisor of the Cisco StackWise Virtual active switch.

How to Configure Cisco StackWise Virtual

Configuring Cisco StackWise Virtual Settings

To enable StackWise Virtual, perform the following procedure on both the switches:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none">• Enter your password if prompted.
Step 2	switch switch-number renumber new switch -number Example: Device# switch 1 renumber 2	(Optional) Reassigns the switch number. The default switch number will be 1. The valid values for the new switch number are 1 and 2.
Step 3	switch switch-number priority priority-number Example: Device# switch 1 priority 5	(Optional) Assigns the priority number. The default priority number is 1. The highest priority number is 15.

	Command or Action	Purpose
Step 4	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 5	stackwise-virtual Example: Device (config) # stackwise-virtual	Enables Cisco StackWise Virtual and enters stackwise-virtual submode.
Step 6	domain id Example: Device (config-stackwise-virtual) # domain 2	(Optional) Specifies the Cisco StackWise Virtual domain ID. The domain ID range is from 1 to 255. The default value is one.
Step 7	end Example: Device (config-stackwise-virtual) # end	Returns to privileged EXEC mode.
Step 8	show stackwise-virtual Example: Device# show stackwise-virtual	
Step 9	write memory Example: Device# write memory	Saves the running-configuration which resides in the system RAM and updates the ROMmon variables. If you do not save the changes, the changes will no longer be part of the startup configuration when the switch reloads. Note that the configurations for stackwise-virtual and domain are saved to the running-configuration and the startup-configuration after the reload.
Step 10	reload Example: Device# reload	Restarts the switch and forms the stack.

Configuring Cisco StackWise Virtual Link



Note Depending on the switch model, SVL is supported on all 10G interfaces and 40G interfaces of the Cisco Catalyst 9500 Series switches and on all the 100G, 40G, 25G and 10G interfaces of the Cisco Catalyst 9500 Series High Performance switches. However, a combination of different interface speeds is not supported.

To configure a switch port as an SVL port, perform the following procedure on both the switches:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	Perform one the of the following depending on the switch that you are configuring. <ul style="list-style-type: none"> • If you are configuring a Cisco Catalyst 9500 Series Switch, use interface {TenGigabitEthernet FortyGigabitEthernet} <interface> • If you are configuring a Cisco Catalyst 9500 Series High Performance Switch, use interface {HundredGigE FortyGigabitEthernet TwentyFiveGigE} <interface> Example: Device(config)# interface TenGigabitEthernet1/0/2	Enters ethernet interface configuration mode.
Step 4	stackwise-virtual link <i>link value</i> Example: Device(config-if)# stackwise-virtual link 1	Associates the interface with configured SVL.
Step 5	end Example: Device(config-if)# end	Returns to privileged EXEC mode.
Step 6	write memory Example: Device# write memory	Saves the running-configuration which resides in the system RAM and updates the ROMmon variables. If you do not save the changes, the changes will no longer be part of the startup configuration when the switch reloads. Note that the configuration for stackwise-virtual link <i>link value</i> is saved only in the

	Command or Action	Purpose
		running-configuration and not the startup-configuration.
Step 7	reload Example: Device# reload	Restarts the switch. Note When converting a Cisco Catalyst 9500 Series High Performance switch from standalone mode to SVL mode for the first time, one of the switches boots up or resets, for resolving the switch number conflict and sets the SWITCH_NUMBER environment variable to 2. The following message appears on the console prompt indicating this: <pre> Waiting for remote chassis to join ##### Chassis number is 2 All chassis in the stack have been discovered. Accelerating discovery Chassis is reloading, reason: Configured Switch num conflicts with peer, Changing local switch number to 2 and reloading to take effect </pre>

Configuring StackWise Virtual Fast Hello Dual-Active-Detection Link

To configure StackWise Virtual Fast Hello DAD link, perform the following procedure. This procedure is optional.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	Perform one the of the following depending on the switch that you are configuring.	Enters ethernet interface configuration mode.

	Command or Action	Purpose
	<ul style="list-style-type: none"> If you are configuring a Cisco Catalyst 9500 Series Switch, use interface {TenGigabitEthernet FortyGigabitEthernet} <interface> If you are configuring a Cisco Catalyst 9500 Series High Performance Switch, use interface {HundredGigE FortyGigabitEthernet TwentyFiveGigE} <interface> <p>Example:</p> <pre>Device(config)#interface TenGigabitEthernet1/0/40</pre>	
Step 4	<p>stackwise-virtual dual-active-detection</p> <p>Example:</p> <pre>Device(config-if)#stackwise-virtual dual-active-detection</pre>	<p>Associates the interface with StackWise Virtual dual-active-detection.</p> <p>Note This command will not be visible on the device after the configuration, but will continue to function.</p>
Step 5	<p>end</p> <p>Example:</p> <pre>Device(config-if)#end</pre>	Returns to privileged EXEC mode.
Step 6	<p>write memory</p> <p>Example:</p> <pre>Device#write memory</pre>	Saves the running-configuration which resides in the system RAM and updates the ROMmon variables. If you do not save the changes, the changes will no longer be part of the startup configuration when the switch reloads. Note that the configuration for stackwise-virtual dual-active-detection is saved only in the running-configuration and not the startup-configuration.
Step 7	<p>reload</p> <p>Example:</p> <pre>Device#reload</pre>	Restarts the switch and configuration takes effect.

Enabling ePAgP Dual-Active-Detection

To enable ePAgP dual-active-detection on a switch port, perform the following procedure on . This procedure is optional.

Procedure

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode.

	Command or Action	Purpose
	Example: Device> enable	<ul style="list-style-type: none"> Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	interface { TenGigabitEthernet FortyGigabitEthernet TwentyFiveGigE } < <i>interface</i> > Example: Device (config) # interface FortyGigabitEthernet 1/0/5	Enters the interface configuration mode.
Step 4	channel-group <i>group_ID</i> mode desirable Example: Device (config-if) # channel-group 1 mode desirable	Enables PAgP MEC with channel-group id in the range of 1 to 126 for 10 GigabitEthernet interfaces.
Step 5	exit Example: Device (config-if) # exit	Exits interface configuration.
Step 6	interface port-channel <i>channel-group-id</i> Example: Device (config) # interface port-channel 1	Selects a port channel interface to configure.
Step 7	shutdown Example: Device (config-if) # shutdown	Shuts down an interface.
Step 8	exit Example: Device (config-if) # exit	Exits interface configuration.
Step 9	stackwise-virtual Example: Device (config) # stackwise-virtual	Enters the StackWise Virtual configuration mode.
Step 10	dual-active detection pagp Example:	Enables pagp dual-active detection. This is enabled by default.

	Command or Action	Purpose
	Device (config-stackwise-virtual) # dual-active detection pagp	
Step 11	dual-active detection pagp trust channel-group <i>channel-group id</i> Example: Device (config-stackwise-virtual) # dual-active detection pagp trust channel-group 1	Enables dual-active detection trust mode on channel-group with the configured ID.
Step 12	exit Example: Device (config-stackwise-virtual) # exit	Exits the StackWise-Virtual configuration mode.
Step 13	interface port-channel <i>portchannel</i> Example: Device (config) # interface port-channel 1	Configured port-channel on the switch.
Step 14	no shutdown Example: Device (config-if) # no shutdown	Enables the configured port-channel on the switch.
Step 15	end Example: Device (config-if) # end	Exits interface configuration.
Step 16	write memory Example: Device# write memory	Saves the running-configuration which resides in the system RAM and updates the ROMmon variables. If you do not save the changes, the changes will no longer be part of the startup configuration when the switch reloads. Note that the configuration for dual-active detection pagp trust channel-group <i>channel-group id</i> is saved to the running-configuration and the startup-configuration after the reload.
Step 17	reload Example: Device# reload	Restarts the switch and configuration takes effect.

Disabling Cisco StackWise Virtual

To disable Cisco StackWise Virtual on a switch, perform the following procedure:

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	Perform one the of the following depending on the switch that you are configuring. <ul style="list-style-type: none"> • If you are configuring a Cisco Catalyst 9500 Series Switch, use interface {TenGigabitEthernet FortyGigabitEthernet}<interface> • If you are configuring a Cisco Catalyst 9500 Series High Performance Switch, use interface {HundredGigE FortyGigabitEthernet TwentyFiveGigE} <interface> Example: Device (config) # interface TenGigabitEthernet 1/0/41	Enters ethernet interface configuration mode.
Step 4	no stackwise-virtual dual-active-detection Example: Device (config-if) # no stackwise-virtual dual-active-detection	Dissociates the interface from StackWise Virtual DAD.
Step 5	Repeat step Step 3, on page 30 . Example: Device (config) # interface TenGigabitEthernet 1/0/5	Enters the interface configuration mode.
Step 6	no stackwise-virtual link link Example: Device (config-if) # no stackwise-virtual link 1	Dissociates the interface from SVL.
Step 7	exit Example: Device (config-if) # exit	Exits interface configuration.

	Command or Action	Purpose
Step 8	no stackwise-virtual Example: Device(config)# no stackwise-virtual	Disables StackWise Virtual configuration.
Step 9	exit Example: Device(config)# exit	Exits the global configuration mode.
Step 10	write memory Example: Device# write memory	Saves the running configuration.
Step 11	reload Example: Device# reload	Restarts the switch and the configuration takes effect.

Verifying Cisco StackWise Virtual Configuration

To verify your StackWise Virtual configuration, use the following **show** commands:

show stackwise-virtual switch <i>number</i> <1-2>	Displays information of a particular switch in the stack.
show stackwise-virtual link	Displays StackWise Virtual link information.
show stackwise-virtual bandwidth	Displays the bandwidth available for the Cisco StackWise Virtual.
show stackwise-virtual neighbors	Displays the Cisco StackWise Virtual neighbors.
show stackwise-virtual dual-active-detection	Displays StackWise Virtual dual-active-detection information.
show stackwise-virtual dual-active-detection pagp	Displays ePAGP dual-active-detection information.

Additional References for StackWise Virtual

Related Documents

Related Topic	Document Title
For complete syntax and usage information for the commands used in this chapter.	High Availability Command Reference for Catalyst 9500 Switches)

MIBs

MIB	MIBs Link
All the supported MIBs for this release.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/support

Feature History for Cisco StackWise Virtual

This table provides release and related information for features explained in this module.

These features are available on all releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Everest 16.6.1	Cisco StackWise Virtual	<p>Cisco StackWise Virtual is a network system virtualization technology that pairs two switches into one virtual switch to simplify operational efficiency with a single control and management plane.</p> <p>Support for this feature was introduced on the C9500-12Q, C9500-16X, C9500-24Q, C9500-40X models of the Cisco Catalyst 9500 Series Switches</p>

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>.



CHAPTER 3

Configuring Graceful Insertion and Removal

Graceful Insertion and Removal (GIR) provides an alternative method to minimize network service impact caused by device maintenance. GIR leverages redundant paths in the network to smoothly remove a device under maintenance, out of service, and insert it back to service when the maintenance is complete. This module describes the how to configure GIR.

- [Restrictions for Graceful Insertion and Removal, on page 33](#)
- [Information about Graceful Insertion and Removal, on page 33](#)
- [How to Configure Graceful Insertion and Removal, on page 34](#)
- [Configuration Examples for Graceful Removal and Insertion, on page 36](#)
- [Monitoring Graceful Insertion and Removal, on page 38](#)
- [Additional References for Graceful Insertion and Removal , on page 38](#)
- [Feature History and Information for GIR, on page 39](#)

Restrictions for Graceful Insertion and Removal

This feature is not supported on the C9500-32C, C9500-32QC, C9500-48Y4C, and C9500-24Y4C models of the Cisco Catalyst 9500 Series Switches.

In Cisco IOS XE Everest 16.6.1, GIR is supported for layer two interface shutdown and ISIS routing protocol. This is configured either by creating customized templates or without a template.

Information about Graceful Insertion and Removal

Overview

Graceful Insertion and Removal (GIR) isolates a switch from the network in order to perform debugging or an upgrade. When switch maintenance is complete, the switch will return to normal mode on either reaching the configured maintenance timeout, or by enabling the **stop maintenance** command.

A switch can be put into maintenance mode using default template or a custom template. The default template contains all the ISIS instances, along with **shut down l2**. In the custom template, you can configure the required ISIS instances and **shutdown l2** option. On entering maintenance mode, all participating protocols are isolated, and L2 ports are shut down. When normal mode is restored, all the protocols and L2 ports are brought back up.

Creating a maintenance mode template before you put the switch in maintenance mode is optional. The objective of maintenance mode for a device is to minimize traffic disruption at the time of removal from the network, as well as during the time of insertion. There are mainly three stages:

- Graceful removal of the node from network.
- Performing maintenance on the device.
- Graceful insertion into the network.

Snapshots are taken automatically while entering and exiting the maintenance mode. You can use the **snapshot create** *snapshot-name snapshot-description* command to capture and store snapshots for pre-selected features. Snapshots are useful to compare the state of a switch before it went into maintenance mode and after it came back to normal mode. The snapshot process consists of three parts:

- Creating a snapshot of the states of a few preselected features on the switch and storing them on the persistent storage media.
- Listing the snapshots taken at various time intervals and managing them.
- Comparing snapshots and showing the summary and details of each feature.

The maximum number of snapshots that may be stored on the switch is 10. You can use the command **snapshot delete** *snapshot-name* to delete a specific snapshot from the device.

Layer 2 interface shutdown

Layer 2 interfaces will be shut down when the system is transitioning into maintenance mode. Layer 2 interfaces are shut down using the **shutdown I2** command in the custom template.

Custom Template

The network administrator can create a template that will be applied when the system goes into maintenance mode. This allows the administrator to isolate specific protocols. All instances that need to be isolated must be explicitly specified.

The admin can create multiple templates with different configurations. However, only a single template will be applied to the maintenance mode CLI. Once applied, the template cannot be updated. If the template needs to be updated, then you must remove it, make the changes, and then re-apply.

How to Configure Graceful Insertion and Removal

Creating maintenance template

Procedure

	Command or Action	Purpose
Step 1	enable Example:	Enables privileged EXEC mode. Enter your password if prompted.

	Command or Action	Purpose
	Device> enable	
Step 2	configure terminal Example: Device# config t	Enters the global configuration mode.
Step 3	maintenance-template <i>template_name</i> Example: Device(config)# maintenance-template girl	Creates a template with the specified name. For example, see Examples: Creating customer profile.
Step 4	router <i>routing_protocol instance_id</i> shutdown I2 Example: Device(config-maintenance-templ)# router isis 1 Device(config-maintenance-templ)# shutdown I2	Creates instances that should be isolated under this template. <ul style="list-style-type: none"> • router: Configures routing protocols and associated instance id. • shutdown I2: Shuts down layer 2 interfaces.

Configuring System Mode Maintenance

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	configure terminal Example: Device# config t	Enters the global configuration mode.
Step 3	system mode maintenance Example: Device(config)# system mode maintenance	Enters system mode maintenance configuration mode. Different sub commands to create maintenance mode parameters are configured in this mode.
Step 4	timeout <i>timeout-value</i> template <i>template-name</i> failsafe <i>failsafe-timeout-value</i> on-reload <i>reset-reason</i> MAINTENANCE	Configures maintenance mode parameters. <ul style="list-style-type: none"> • timeout: Configures maintenance mode timeout period in minutes, after which the system automatically returns to normal mode. • template: Configures maintenance mode using the specified template.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • failsafe: Configures client-ack timeout value. If the system is going into maintenance mode, it will continue to reach maintenance. If the system is exiting from maintenance mode, then it will reach normal mode. • on-reload reset-reason MAINTENANCE: Reloads system on maintenance mode.

Starting and Stopping Maintenance Mode

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	start maintenance Example: Device# start maintenance	Puts the system into maintenance mode.
Step 3	stop maintenance Example: Device# stop maintenance	Puts the system back into normal mode.

Configuration Examples for Graceful Removal and Insertion

The following examples show the sequence followed to enable GIR during a maintenance window.

Example: Configuring maintenance template

This example shows how to configure a maintenance template t1 with an ISIS routing protocol instance.

```
Device# config terminal
Device (config)# maintenance-template t1
Device (config-maintenance-templ)# router isis 1
```

This example shows how to configure a maintenance template t1 with shutdown l2.

```
Device# config terminal
Device (config)# maintenance-template t1
Device (config-maintenance-templ)# shutdown 12
```

Example: Configuring system mode maintenance

This example shows how to create maintenance template and configure the maintenance mode parameters.

```
Device# config terminal
Device(config)# system mode maintenance
Device(config-maintenance)#timeout 20
Device(config-maintenance)#failsafe 30
Device(config-maintenance)#on-reload reset-reason MAINTENANCE
Device(config-maintenance)#template t1
Device(config-maintenance)#exit
```

Example: Starting and Stopping maintenance mode

This example shows how to put the system into maintenance mode.

```
Device# start maintenance
```

After the activity is completed, the system can be put out of maintenance mode.

This example shows how to put the system out of maintenance mode.

```
Device# stop maintenance
```

Example: Displaying system mode settings

This example shows how to display system mode settings using different options.

```
Device#show system mode
  System Mode: Normal

Device#show system mode maintenance
  System Mode: Normal
  Current Maintenance Parameters:
  Maintenance Duration: 15(mins)
  Failsafe Timeout: 30(mins)
  Maintenance Template: t1
  Reload in Maintenance: False

Device#show system mode maintenance clients
  System Mode: Normal
  Maintenance Clients:
  CLASS-EGP

  CLASS-IGP
  router isis 1: Transition None

  CLASS-MCAST

  CLASS-L2
```

```
Device#show system mode maintenance template default
System Mode: Normal
default maintenance-template details:
router isis 1
router isis 2
```

```
Device#show system mode maintenance template t1
System Mode: Normal
Maintenance Template t1 details:

router isis 1
```

Monitoring Graceful Insertion and Removal

Table 3: Privelege EXEC show commands

Command	Purpose
show system mode [maintenance [clients template <i>template-name</i>]]	Displays information about system mode.
show system snapshots [dump < <i>snapshot-file-name</i> >]	Displays all the snapshots present on the device. Using the keyword dump displays all snapshots in XML format.
show system snapshots compare <i>snapshot-name1</i> <i>snapshot-name2</i>	Displays differences between snapshots taken before entering maintenance mode and after exiting from the maintenance mode.

Table 4: Global Troubleshooting Commands

Command	Purpose
debug system mode maintenance	Displays information for troubleshooting GIR feature.

Additional References for Graceful Insertion and Removal

Related Documents

Related Topic	Document Title
For complete syntax and usage information for the commands used in this chapter.	High Availability Command Reference, Cisco IOS XE Everest 16.6.1.

MIBs

MIB	MIBs Link
All the supported MIBs for this release.	To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	http://www.cisco.com/support

Feature History and Information for GIR

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Release	Modification
Cisco IOS XE Everest 16.6.1	This feature was introduced.

