



Configuring Cisco NSF with SSO

- [Finding Feature Information, on page 1](#)
- [Prerequisites for NSF with SSO, on page 1](#)
- [Restrictions for NSF with SSO, on page 2](#)
- [Information About NSF with SSO, on page 2](#)
- [How to Configure Cisco NSF with SSO , on page 7](#)

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 NSF with SSO

The following are prerequisites and considerations for configuring NSF with SSO.

- Use of the routing protocols requires the IP Services license level. EIGRP-stub and OSPF for routed access are supported on IP Base license level.
- BGP support in NSF requires that neighbor networking devices be NSF-aware; that is, the devices must have the graceful restart capability and advertise that capability in their OPEN message during session establishment. If an NSF-capable router 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.
- OSPF support in NSF requires that all neighbor networking devices be NSF-aware. If an NSF-capable router 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 routers continue to provide NSF capabilities.

Restrictions for NSF with SSO

The following are restrictions for configuring NSF with SSO:

- NSF does not support IP Multicast Routing, as it is not SSO-aware.
- NSF is not supported if the IOS-XE software is running in the LAN Base mode.
- For NSF operation, you must have SSO configured on the device.
- All Layer 3 neighboring devices must be NSF Helper or NSF-capable to support graceful restart capability.
- For IETF, all neighboring devices must be running an NSF-aware software image.

Information About NSF with SSO

Overview of NSF with SSO

The switch supports fault resistance by allowing a standby switch to take over if the active switch becomes unavailable. Cisco nonstop forwarding (NSF) works with stateful switchover (SSO) to minimize the amount of time a network is unavailable.

NSF provides these 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 may be improved with the reduction in the number of route flaps, which were created when routers in the network failed and lost their routing tables.
- Neighboring routers do not detect a link flap—Because the interfaces remain up during a switchover, neighboring routers 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.

Keepalive messages are sent and received between the active and standby switches.

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

In addition, hello messages are sent and received by all stack members.

- If a stack member does not respond, that member is removed from the stack.
- If the standby switch does not respond, a new standby switch is elected.
- If the active switch does not respond, the standby switch becomes the active switch.

SSO Operation

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

If the active switch fails, the standby switch becomes the active switch. This new active switch uses existing Layer 2 switching information to continue forwarding traffic. Layer 3 forwarding will be delayed until the routing tables have been repopulated in the newly active switch.



Note SSO Layer 2 Only is supported if the IOS-XE software is running the LAN Base license level.

The state of these features is preserved between both the active and standby switches:

- 802.3
- 802.3u
- 802.3x (Flow Control)
- 802.3ab (GE)
- 802.3z (Gigabit Ethernet including CWDM)
- 802.3ad (LACP)
- 802.1p (Layer 2 QoS)
- 802.1q
- 802.1X (Authentication)
- 802.1D (Spanning Tree Protocol)
- 802.3af (Inline power)
- PAgP
- VTP
- Dynamic ARP Inspection
- DHCP
- DHCP snooping
- IP source guard
- IGMP snooping (versions 1 and 2)
- DTP (802.1q and ISL)
- MST
- PVST+

- Rapid-PVST
- PortFast/UplinkFast/BackboneFast
- BPDU guard and filtering
- Voice VLAN
- Port security
- Unicast MAC filtering
- ACL (VACLs, PACLS, RACLS)
- QOS (DBL)
- Multicast storm control/broadcast storm control

SSO is compatible with the following list of features. However, the protocol database for these features is not synchronized between the standby and active switches:

- 802.1Q tunneling with Layer 2 Protocol Tunneling (L2PT)
- Baby giants
- Jumbo frame support
- CDP
- Flood blocking
- UDLD
- SPAN/RSPAN
- NetFlow

All Layer 3 protocols on a switch are learned on the standby switch if SSO is enabled.

NSF Operation

Cisco IOS Nonstop Forwarding (NSF) always runs with stateful switchover (SSO) and provides redundancy for Layer 3 traffic. NSF is supported by the BGP, OSPF, and EIGRP routing protocols and is supported by Cisco Express Forwarding (CEF) for forwarding. The routing protocols have been enhanced with NSF-capability and awareness, which means that routers running these protocols can detect a switchover and take the necessary actions to continue forwarding network traffic and to recover route information from the peer devices.

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

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

The switch supports NSF-awareness and NSF-capability for the BGP, OSPF, and EIGRP protocols in IP Services license level and NSF-awareness for the EIGRP-stub in IP Base license level.

NSF has two primary components:

- NSF-awareness

A networking device is NSF-aware if it is running NSF-compatible software. If neighboring router devices detect that an NSF router can still forward packets when an active switch election happens, this capability is referred to as NSF-awareness. Cisco IOS enhancements to the Layer 3 routing protocols (BGP, OSPF, and EIGRP) are designed to prevent route-flapping so that the CEF routing table does not time out or the NSF router does not drop routes. An NSF-aware router helps to send routing protocol information to the neighboring NSF router. 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 has been 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 switch 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. The routing protocols recover routing information from neighbor devices and rebuild the Cisco Express Forwarding (CEF) 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 (CEF). CEF maintains the FIB and uses the FIB information that was current at the time of the switchover to continue forwarding packets during a switchover. This feature reduces traffic interruption during the switchover.

During normal NSF operation, CEF on the active supervisor switch synchronizes its current FIB and adjacency databases with the FIB and adjacency databases on the standby switch. Upon switchover, the standby switch initially has FIB and adjacency databases that are mirror images of those that were current on the active switch. CEF keeps the forwarding engine on the standby switch current with changes that are sent to it by CEF on the active switch. 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 CEF, 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 switch 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.

BGP Operation

When an NSF-capable router begins a BGP session with a BGP peer, it sends an OPEN message to the peer. Included in the message is a statement 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 received this capability, it is aware that the device sending the message is NSF-capable. Both the NSF-capable router and its BGP peers need to exchange the graceful restart capability in their OPEN messages at the time of session establishment. If both the peers do not exchange the graceful restart capability, the session will not be capable of a graceful restart.

If the BGP session is lost during the active switch switchover, the NSF-aware BGP peer marks all the routes associated with the NSF-capable router as stale; however, it continues to use these routes to make forwarding

decisions for a set period of time. This functionality prevents packets from being lost while the newly active switch is waiting for convergence of the routing information with the BGP peers.

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

At this point, the routing information is exchanged between the two BGP peers. After 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; the BGP protocol then is fully converged.

If a BGP peer does not support the graceful restart capability, it ignores the graceful restart capability in an OPEN message but establishes 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 is not capable of a graceful restart.



Note BGP support in NSF requires that neighbor networking devices be NSF-aware; that is, the devices must have the graceful restart capability and advertise that capability in their OPEN message during session establishment. If an NSF-capable router 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.

OSPF Operation

When an OSPF NSF-capable router performs an active switch 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 an active switch switchover, the NSF-capable router 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 router should not be reset. As the NSF-capable router receives signals from other routers on the network, it can begin to rebuild its neighbor list.

After neighbor relationships are reestablished, the NSF-capable router 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 router 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 routers continue to provide NSF capabilities.

EIGRP Operation

When an EIGRP NSF-capable router initially re-boots after an NSF restart, it has no neighbor and its topology table is empty. The router is notified by the standby (now active) switch when it needs to bring up the interfaces, reacquire neighbors, and rebuild the topology and routing tables. The restarting router and its peers must accomplish these tasks without interrupting the data traffic directed toward the restarting router. EIGRP peer routers maintain the routes learned from the restarting router and continue forwarding traffic through the NSF restart process.

To prevent an adjacency reset by the neighbors, the restarting router uses a new Restart (RS) bit in the EIGRP packet header to indicate a restart. The RS bit is set in the hello packets and in the initial INIT update packets during the NSF restart period. The RS bit in the hello packets allows the neighbors to be quickly notified of the NSF restart. Without seeing the RS bit, the neighbor can only detect an adjacency reset by receiving an INIT update or by the expiration of the hello hold timer. Without the RS bit, a neighbor does not know if the adjacency reset should be handled using NSF or the normal startup method.

When the neighbor receives the restart indication, either by receiving the hello packet or the INIT packet, it recognizes the restarting peer in its peer list and maintains the adjacency with the restarting router. The neighbor then sends its topology table to the restarting router with the RS bit set in the first update packet indicating that it is NSF-aware and is helping out the restarting router. The neighbor does not set the RS bit in their hello packets, unless it is also a NSF restarting neighbor.



Note A router may be NSF-aware but may not be helping the NSF restarting neighbor because booting from a cold start.

If at least one of the peer routers is NSF-aware, the restarting router would then receive updates and rebuild its database. The restarting router must then find out if it had converged so that it can notify the routing information base (RIB). Each NSF-aware router is required to send an end of table (EOT) marker in the last update packet to indicate the end of the table content. The restarting router knows it has converged when it receives the EOT marker. The restarting router can then begin sending updates.

An NSF-aware peer would know when the restarting router had converged when it receives an EOT indication from the restarting router. The peer then scans its topology table to search for the routes with the restarted neighbor as the source. The peer compares the route timestamp with the restart event timestamp to determine if the route is still available. The peer then goes active to find alternate paths for the routes that are no longer available through the restarted router.

When the restarting router has received all EOT indications from its neighbors or when the NSF converge timer expires, EIGRP notifies the RIB of convergence. EIGRP waits for the RIB convergence signal and then floods its topology table to all awaiting NSF-aware peers.

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	redundancy Example: Device(config)# redundancy	Enters redundancy configuration mode.
Step 2	mode sso Example: Device(config-red)# mode sso	Configures SSO. When this command is entered, the standby switch is reloaded and begins to work in SSO mode.
Step 3	end Example: Device(config-red)# end	Returns to EXEC mode.
Step 4	show running-config Example: Device# show running-config	Verifies that SSO is enabled.
Step 5	show redundancy states Example: Device# show redundancy states	Displays the operating redundancy mode.

Configuring SSO Example

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

```

Device(config)# redundancy
  Device(config)# mode sso
  Device(config)# end
Device# 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 CEF NSF

To verify CEF NSF, use the **show cef state** privileged EXEC command.

```
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.
```

Configuring BGP for NSF

You must configure BGP graceful restart on all peer devices participating in BGP NSF.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: Device(config)# configure terminal	Enters global configuration mode.
Step 2	router bgp as-number Example: Device(config)# router bgp 300	Enables a BGP routing process, which places the switch in switch configuration mode.

	Command or Action	Purpose
Step 3	bgp graceful-restart Example: Device(config)# bgp graceful-restart	Enables the BGP graceful restart capability, starting BGP NSF. If you enter this command after the BGP session has been established, you must restart the session for the capability to be exchanged with the BGP neighbor. Use this command on the restarting switch and all of its peers.

Verifying BGP NSF

To verify BGP NSF, you must check that BGP graceful restart is configured on the SSO-enabled networking device and on the neighbor devices. To verify, follow these steps:

Procedure

- Step 1** Verify that “bgp graceful-restart” appears in the BGP configuration of the SSO-enabled switch by entering the **show running-config** command:
- Example:**
- ```
Device# show running-config
.
.
.
router bgp 120
.
.
.
bgp graceful-restart
neighbor 192.0.2.0 remote-as 300
.
.
.
```
- Step 2** Repeat Step 1 on each of the BGP neighbors.
- Step 3** On the SSO device and the neighbor device, verify that the graceful restart function is shown as both advertised and received, and confirm the address families that have the graceful restart capability. If no address families are listed, BGP NSF does not occur either:

**Example:**

```
Device# show ip bgp neighbors
BGP neighbor is 192.0.2.3, remote AS 1, internal link
BGP version 4, remote router ID 192.0.2.4
BGP state = Established, up for 00:02:38
Last read 00:00:38, last write 00:00:35, hold time is 180, keepalive interval is 60
seconds
Neighbor capabilities:
Route refresh: advertised and received(new)
Address family IPv4 Unicast: advertised and received
Message statistics:
InQ depth is 0
OutQ depth is 0
Sent Rcvd
Opens: 1 1
```

```

Notifications: 0 0
Updates: 0 0
Keepalives: 4 4
Route Refresh: 0 0
Total: 5 5
Default minimum time between advertisement runs is 0 seconds
.....
(Remaining output deleted)

```

## Configuring OSPF NSF

All peer devices participating in OSPF NSF must be made OSPF NSF-aware, which happens automatically when you install an NSF software image on the device.

### Procedure

|               | Command or Action                                                                               | Purpose                                                                                |
|---------------|-------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure terminal</b><br><b>Example:</b><br>Device(config)# <b>configure terminal</b>       | Enters global configuration mode.                                                      |
| <b>Step 2</b> | <b>router ospf processID</b><br><b>Example:</b><br>Device(config)# <b>router ospf processID</b> | Enables an OSPF routing process, which places the switch in router configuration mode. |
| <b>Step 3</b> | <b>nsf</b><br><b>Example:</b><br>Device(config)# <b>nsf</b>                                     | Enables NSF operations for OSPF.                                                       |

## Verifying OSPF NSF

### Procedure

**Step 1** Verify that ‘nsf’ appears in the OSPF configuration of the SSO-enabled device by entering the show running-config command:

**Example:**

```

Device(config)#show running-config
route ospf 120
log-adjacency-changes
nsf
network 192.0.2.0 192.0.2.255 area 0
network 192.0.2.1 192.0.2.255 area 1
network 192.0.2.2 192.0.2.255 area 2
.
.

```

**Step 2** Enter the **show ip ospf** command to verify that NSF is enabled on the device:

**Example:**

```
Device show ip ospf
Routing Process "ospf 1" with ID 192.0.2.1
Start time: 00:02:07.532, Time elapsed: 00:39:05.052
Supports only single TOS(TOS0) routes
Supports opaque LSA
Supports Link-local Signaling (LLS)
transit capable is 0
External flood list length 0
IETF Non-Stop Forwarding enabled
restart-interval limit: 120 sec
IETF NSF helper support enabled
Cisco NSF helper support enabled
Reference bandwidth unit is 100 mbps
Area BACKBONE(0)
Number of interfaces in this area is 3 (1 loopback)
Area has no authentication
SPF algorithm last executed 00:08:53.760 ago
SPF algorithm executed 2 times
Area ranges are
Number of LSA 3. Checksum Sum 0x025BE0
Number of opaque link LSA 0. Checksum Sum 0x000000
Number of DCbitless LSA 0
Number of indication LSA 0
Number of DoNotAge LSA 0
Flood list length 0
```

## Configuring EIGRP NSF

### Procedure

|               | Command or Action                                                                                     | Purpose                                                                                     |
|---------------|-------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------|
| <b>Step 1</b> | <b>configure terminal</b><br><br><b>Example:</b><br>Device <b>configure terminal</b>                  | Enters global configuration mode.                                                           |
| <b>Step 2</b> | <b>router eigrp as-number</b><br><br><b>Example:</b><br>Device(config)# <b>router eigrp as-number</b> | Enables an EIGRP routing process, which places the switch in router configuration mode.     |
| <b>Step 3</b> | <b>nsf</b><br><br><b>Example:</b><br>Device(config-router)# <b>nsf</b>                                | Enables EIGRP NSF.<br><br>Use this command on the “restarting” switch and all of its peers. |

## Verifying EIGRP NSF

### Procedure

---

- Step 1** Verify that “nsf” appears in the EIGRP configuration of the SSO-enabled device by entering the **show running-config** command:

**Example:**

```
Device show running-config
..
.
router eigrp 100
auto-summary
nsf
..
.
```

- Step 2** Enter the **show ip protocols** command to verify that NSF is enabled on the device:

**Example:**

```
Device show ip protocols
*** IP Routing is NSF aware ***
Routing Protocol is "ospf 1"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
Router ID 192.0.2.3
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Maximum path: 1
Routing for Networks:
Routing on Interfaces Configured Explicitly (Area 0):
Loopback0
GigabitEthernet5/3
TenGigabitEthernet3/1
Routing Information Sources:
Gateway Distance Last Update
192.0.2.1 110 00:01:02
Distance: (default is 110)
Routing Protocol is "bgp 601"
Outgoing update filter list for all interfaces is not set
Incoming update filter list for all interfaces is not set
IGP synchronization is disabled
Automatic route summarization is disabled
Neighbor(s):
Address FiltIn FiltOut DistIn DistOut Weight RouteMap
192.0.2.0
Maximum path: 1
Routing Information Sources:
Gateway Distance Last Update
192.0.2.0 20 00:01:03
Distance: external 20 internal 200 local 200
```

---

