



CHAPTER 9

Configuring Cisco NSF with SSO Supervisor Engine Redundancy

This chapter describes how to configure supervisor engine redundancy using Cisco nonstop forwarding (NSF) with stateful switchover (SSO).

This chapter consists of these sections:

- [About NSF with SSO Supervisor Engine Redundancy, page 9-1](#)
- [Configuring NSF with SSO Supervisor Engine Redundancy, page 9-7](#)
- [Cisco High Availability Features in Cisco IOS XE 3.1.0SG, page 9-13](#)



Note

For complete syntax and usage information for the switch commands used in this chapter, look at the *Cisco Catalyst 4500 Series Switch Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/hw/switches/ps4324/index.html>

If the command is not found in the Catalyst 4500 Command Reference, it is located in the larger Cisco IOS library. Refer to the *Catalyst 4500 Series Switch Cisco IOS Command Reference* and related publications at this location:

<http://www.cisco.com/en/US/products/ps6350/index.html>

About NSF with SSO Supervisor Engine Redundancy

These sections describe supervisor engine redundancy using NSF with SSO:

- [About Cisco IOS NSF-Aware and NSF-Capable Support, page 9-2](#)
- [NSF with SSO Supervisor Engine Redundancy Overview, page 9-3](#)
- [SSO Operation, page 9-3](#)
- [NSF Operation, page 9-4](#)
- [Cisco Express Forwarding, page 9-4](#)
- [Routing Protocols, page 9-5](#)
- [NSF Guidelines and Restrictions, page 9-7](#)

About Cisco IOS NSF-Aware and NSF-Capable Support

Cisco IOS Nonstop Forwarding (NSF) has two primary components:

- NSF-awareness—If neighboring router devices detect that an NSF router can still forward packets when a supervisor engine switchover happens, this capability is referred to as NSF-awareness. Cisco IOS enhancements to the Layer 3 routing protocols (OSPF, BGP, 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-capability—NSF works with SSO to minimize the amount of time that a Layer 3 network is unavailable following a supervisor engine switchover by continuing to forward IP packets. Reconvergence of Layer 3 routing protocols (BGP, EIGRP, and OSPF v2) 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.




Note NSF does not support IPv6 and is IPv4 Unicast only.



Note NSF capable devices include Catalyst 4500 series switches, Catalyst 6500 series switches, Cisco 7500 series routers, Cisco 10000 series routers, and Cisco 12000 series routers.

A typical topology for NSF and NSF-aware routers is given below.

Figure 9-1 Topology for NSF and NSF-Capable Switches



The Catalyst 4500 series switch supports NSF capability and NSF-awareness for the EIGRP, OSPF, and BGP protocols in Enterprise Services mode and NSF-awareness for the EIGRP-stub in IP Base mode. NSF-awareness is turned on by default for EIGRP-stub, EIGRP, and OSPF protocols. For BGP, you need to turn it on manually.

If the supervisor engine is configured for BGP (with the **graceful-restart** command), EIGRP, or OSPF routing protocols, routing updates are automatically sent during the supervisor engine switchover of a neighboring NSF capable switch.

Table 9-1 lists the supervisor engines and Catalyst 4500 series switches that support NSF-capable.

Table 9-1 NSF-Capable Supervisor Engines

NSF-Capable Supervisor Engine	Switch Support
Supervisor Engine 7-E (WS-X45-SUP7-E)	WS-C4507R-E
Supervisor Engine 7-E (WS-X45-SUP7-E)	WS-C4510R-E
Supervisor Engine 7-E (WS-X45-SUP7-E)	WS-C4507R+E
Supervisor Engine 7-E (WS-X45-SUP7-E)	WS-C4510R+E

NSF with SSO Supervisor Engine Redundancy Overview

Catalyst 4500 series switches support fault resistance by allowing a standby supervisor engine to take over if the active supervisor engine fails. NSF works with SSO to minimize the amount of time a network is unavailable to its users following a switchover.

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

Catalyst 4500 series switches also support route processor redundancy (RPR). For information about these redundancy modes, see [Chapter 5, “Configuring Supervisor Engine Redundancy Using RPR and SSO.”](#)

SSO Operation

SSO establishes one of the supervisor engines as active while the other supervisor engine is designated as standby, and then SSO synchronizes information between them. A switchover from the active to the standby supervisor engine occurs when the active supervisor engine fails, or is removed from the switch, or is manually shut down for maintenance.

In networking devices running SSO, both supervisor engines must be running the same Cisco IOS software version (except during the ISSU upgrade/downgrade process) and ROMMON version so that the standby supervisor engine is always ready to assume control following a fault on the active supervisor engine. SSO switchover also preserves FIB and adjacency entries and can forward Layer 3 traffic after a switchover. Configuration information and data structures are synchronized from the active

to the standby supervisor engine and whenever changes to the active supervisor engine configuration occur. Following an initial synchronization between the two supervisor engines, SSO maintains state information between them, including forwarding information.

During switchover, system control and routing protocol execution is transferred from the active supervisor engine to the standby supervisor engine.

**Note**

Be aware that you can use the **[no] service slave-log** configuration command to forward all error messages from the standby supervisor engine to the active engine. By default, this capability is enabled.

NSF Operation

NSF always runs with 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.

A networking device is NSF-aware if it is running NSF-compatible software. A device is NSF-capable if it has been configured to support NSF; it rebuilds routing information from NSF-aware or NSF-capable neighbors.

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.

Cisco Express Forwarding

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

**Note**

Use of the routing protocols require the Enterprise Services (entservices) level of Cisco IOS-XE Software for the Catalyst 4500 series switch. EIGRP-stub and OSPF for routed access are supported on IP Base level of Cisco IOS XE.

The routing protocols run only on the active supervisor engine, and they receive routing updates from their neighbor routers. Routing protocols do not run on the standby supervisor engine. Following a switchover, the routing protocols request that the NSF-aware neighbor devices send state information to help rebuild the routing tables. NSF supports the BGP, OSPF, and EIGRP protocols.

**Note**

For NSF operation, the routing protocols depend on CEF to continue forwarding packets while the routing protocols rebuild the routing 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 supervisor engine 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 supervisor engine is waiting for convergence of the routing information with the BGP peers.

After a supervisor engine 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 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 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) supervisor engine 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.

NSF Guidelines and Restrictions

NSF with SSO has these restrictions:

- For NSF operation, you must have SSO configured on the device.
- NSF with SSO supports IP Version 4 traffic and protocols only; NSF with SSO does not support IPv6 traffic.
- The Virtual Redundancy Routing Protocols (VRRP) is not SSO-aware, meaning state information is not maintained between the active and standby supervisor engine during normal operation. VRRP and SSO can coexist but both features work independently. Traffic that relies on VRRP may switch to the VRRP standby in the event of a supervisor engine switchover.
- All neighboring devices participating in BGP NSF must be NSF-capable and configured for BGP graceful restart.
- OSPF NSF for virtual links is not supported.
- All OSPF networking devices on the same network segment must be NSF-aware (running an NSF software image).
- For IETF, all neighboring devices must be running an NSF-aware software image.

Configuring NSF with SSO Supervisor Engine Redundancy

The following sections describe the configuration tasks for the NSF feature:

- [Configuring SSO, page 9-8](#)
- [Configuring CEF NSF, page 9-8](#)
- [Verifying CEF NSF, page 9-9](#)
- [Configuring BGP NSF, page 9-9](#)
- [Verifying BGP NSF, page 9-10](#)
- [Configuring OSPF NSF, page 9-11](#)
- [Verifying OSPF NSF, page 9-11](#)
- [Configuring EIGRP NSF, page 9-12](#)
- [Verifying EIGRP NSF, page 9-13](#)

Configuring SSO

You must configure SSO in order to use NSF with any supported protocol. To configure SSO, perform this task:

	Command	Purpose
Step 1	Switch(config)# redundancy	Enters redundancy configuration mode.
Step 2	Switch(config-red)# mode sso	Configures SSO. When this command is entered, the standby supervisor engine is reloaded and begins to work in SSO mode.
Step 3	Switch(config-red)# end	Returns to EXEC mode.
Step 4	Switch# show running-config	Verifies that SSO is enabled.
Step 5	Switch# show redundancy states	Displays the operating redundancy mode.



Note

The **sso** keyword is not supported if the IOS-XE software is running in LAN Base level.

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

```
Switch> enable
Switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)# redundancy
Switch(config-red)# mode sso
Switch(config-red)# end
Switch# 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
Switch#
```

Configuring CEF NSF

The CEF NSF feature operates by default while the networking device is running in SSO mode. No configuration is necessary.

Verifying CEF NSF

To verify that CEF is NSF-capable, enter the **show cef state** command:

```
Switch# 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 NSF



Note

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

To configure BGP for NSF, perform this task (repeat this procedure on each of the BGP NSF peer devices):

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.

	Command	Purpose
Step 2	Switch(config)# router bgp <i>as-number</i>	Enables a BGP routing process, which places the switch in switch configuration mode.
Step 3	Switch(config-router)# bgp graceful-restart	<p>Enables the BGP graceful restart capability, starting BGP NSF.</p> <p>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.</p> <p>Use this command on the restarting switch and all of its peers.</p>

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:

- Step 1** Verify that “bgp graceful-restart” appears in the BGP configuration of the SSO-enabled switch by entering the **show running-config** command:

```
Switch# show running-config
.
.
.
router bgp 120
.
.
.
bgp graceful-restart
  neighbor 10.2.2.2 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:

```
Switch# show ip bgp neighbors
BGP neighbor is 31.31.31.7, remote AS 1, internal link
  BGP version 4, remote router ID 7.7.7.7
  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



Note

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.

To configure OSPF NSF, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# router ospf <i>processID</i>	Enables an OSPF routing process, which places the switch in router configuration mode.
Step 3	Switch(config-router)# nsf	Enables NSF operations for OSPF.

Verifying OSPF NSF

To verify OSPF NSF, you must check that the NSF function is configured on the SSO-enabled networking device. To verify OSPF NSF, follow these steps:

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

```

Switch# show running-config

route ospf 120
log-adjacency-changes
nsf
network 192.168.20.0 0.0.0.255 area 0
network 192.168.30.0 0.0.0.255 area 1
network 192.168.40.0 0.0.0.255 area 2
.
.
.

```

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

```

Switch> show ip ospf
Routing Process "ospf 1" with ID 187.1.1.1
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)

```

```

Supports area transit capability
Router is not originating router-LSAs with maximum metric
Initial SPF schedule delay 5000 msec
Minimum hold time between two consecutive SPFs 10000 msec
Maximum wait time between two consecutive SPFs 10000 msec
Incremental-SPF disabled
Minimum LSA interval 5 secs
Minimum LSA arrival 1000 msec
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msec
Retransmission pacing timer 66 msec
Number of external LSA 0. Checksum Sum 0x000000
Number of opaque AS LSA 0. Checksum Sum 0x000000
Number of DCbitless external and opaque AS LSA 0
Number of DoNotAge external and opaque AS LSA 0
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Number of areas 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

```

Switch#

Configuring EIGRP NSF

To configure EIGRP NSF, perform this task:

	Command	Purpose
Step 1	Switch# configure terminal	Enters global configuration mode.
Step 2	Switch(config)# router eigrp as-number	Enables an EIGRP routing process, which places the switch in router configuration mode.
Step 3	Switch(config-router)# nsf	Enables EIGRP NSF. Use this command on the “restarting” switch and all of its peers.

Verifying EIGRP NSF

To verify EIGRP NSF, you must check that the NSF function is configured on the SSO-enabled networking device. To verify EIGRP NSF, follow these steps:

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

```
Switch# 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:

```
Switch# 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 187.1.1.1
  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
    121.1.1.1       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
    150.1.1.1
  Maximum path: 1
  Routing Information Sources:
    Gateway         Distance      Last Update
    150.1.1.1       20           00:01:03
  Distance: external 20 internal 200 local 200

Switch#
```

Cisco High Availability Features in Cisco IOS XE 3.1.0SG

This section provides a list of High Availability software features that are supported in Cisco IOS XE 3.1.0SG. Links to the feature documentation are included.

Feature guides may contain information about more than one feature. To find information about a specific feature within a feature guide, see the Feature Information table at the end of the guide.

Feature guides document features that are supported on many different software releases and platforms. Your Cisco software release or platform may not support all the features documented in a feature guide. See the Feature Information table at the end of the feature guide for information about which features in that guide are supported in your software release. 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.

The following high availability features were introduced with Cisco IOS XE 3.1.0SG:

- Enhanced High System Availability
- NSF - Graceful Restart (GR) and Non Stop Routing (NSR) for IS-IS
- NSF - OSPF
- NSF/SSO (Nonstop Forwarding with Stateful Switchover)
- SSO - HDLC
- SSO - HSRP
- SSO - Multilink PPP (MLP)
- SSO - PPP

For details on Cisco High Availability, refer to the following URL:

http://www.cisco.com/en/US/products/ps6550/products_ios_technology_home.html