BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

First Published: March 20, 2006
Last Updated: March 20, 2006

The BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) feature enables provider edge (PE) routers to maintain Border Gateway Protocol (BGP) state with customer edge (CE) routers and ensure continuous packet forwarding during a Route Processor (RP) switchover or during a planned In-Service Software Upgrade (ISSU) for a PE router. CE routers do not need to be Nonstop Forwarding (NSF)-capable or NSF-aware to benefit from BGP NSR capabilities on PE routers. Only PE routers need to be upgraded to support BGP NSR—no CE router upgrades are required. BGP NSR with SSO, thus, enables service providers to provide the benefits NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart.

Finding Feature Information in This Module

Your Cisco IOS software release may not support all of the features documented in this module. To reach links to specific feature documentation in this module and to see a list of the releases in which each feature is supported, use the "Feature Information for BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)" section on page 41.

Finding Support Information for Platforms and Cisco IOS Software Images

Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at http://www.cisco.com/go/fn. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Contents

CISCO SYSTEMS

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

Prerequisites for BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

- This document assumes that your network is configured to run BGP.
- This document assumes that Multiprotocol Layer Switching (MPLS) Layer 3 Virtual Private Networks (VPNs) are configured.
- This document assumes that you are familiar NSF and SSO concepts and tasks.

Restrictions for BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

- This feature is supported on Cisco 10000 Series Performance Routing Engines 2 (PRE2s) and Cisco 10000 Series Performance Routing Engines 3 (PRE3s) .

Information About BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

To configure BGP NSR with SS0, you should be familiar with the following concepts:

Overview of BGP NSR with SSO

Prior to the introduction of BGP NSR with SS0 in Cisco IOS Release 12.2(28)SB, BGP required that all neighboring devices participating in BGP NSF be configured to be either NSF-capable or NSF-aware (by configuring the devices to support the BGP graceful restart mechanism). BGP NSF, thus, required that all neighboring devices be upgraded to a version of Cisco IOS software that supports BGP graceful restart. However, in many MPLS VPN deployments, there are situations where PE routers engage in exterior BGP (eBGP) peering sessions with CE routers that do not support BGP graceful restart and cannot be upgraded to a software version that supports BGP graceful restart in the same time frame as the provider (P) routers.

BGP NSR with SSO provides a high availability (HA) solution to service providers whose PE routers engage in eBGP peering relationships with CE routers that do not support BGP graceful restart. BGP NSR works with SSO to synchronize BGP state information between the active and standby RP. SSO minimizes the amount of time a network is unavailable to its users following a switchover. When the

BGP NSR with SSO feature is configured, in the event of an RP switchover, the PE router uses BGP NSR with SSO to maintain BGP state for eBGP peering sessions with CEs that are not NSF-aware (see Figure 1). Additionally, the BGP NSR with SSO feature dynamically detects NSF-aware peers and runs graceful restart with those CE routers. For eBGP peering sessions with NSF-aware peers and for internal BGP (iBGP) sessions with BGP Route Reflectors (RRs) in the service provider core, the PE uses NSF to maintain BGP state. BGP NSR with SSO, thus, enables service providers to provide the benefits of NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart.

Figure 1 BGP NSR with SSO Operations During an RP Switchover

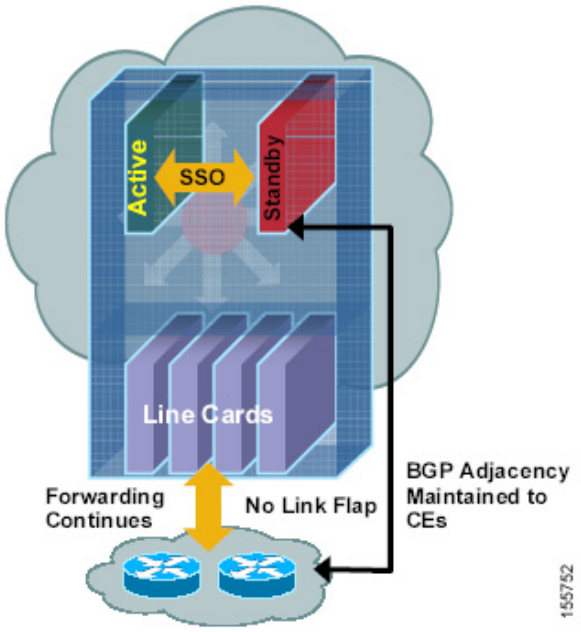

BGP NSR with SSO is supported in BGP peer, BGP peer group, and BGP session template configurations. To configure support for BGP NSR with SSO in BGP peer and BGP peer group configurations, use the **neighbor ha-mode sso** command in address family configuration mode for IPv4 VRF address family BGP peer sessions. To include support for Cisco BGP NSR with SSO in a peer session template, use the **ha-mode sso** command in session-template configuration mode.

Benefits of BGP NSR with SSO

- Minimizes services disruptions—BGP NSR with SSO reduces impact on customer traffic during RP switchovers (scheduled or unscheduled events), extending HA deployments and benefits at the edge.
- Enhances high-availability NSF and SSO deployment at the edge—BGP NSR with SSO allows incremental deployment by upgrading the provider edge with the NSR capability so that customer-facing edge routers are synchronized automatically and no coordination or NSF awareness is needed with the customer side Cisco or third-party customer edge routers. The BGP NSR feature dynamically detects NSF-aware peers and runs graceful restart with those CE routers.
- Provides transparent route convergence—BGP NSR with SSO eliminates route flaps by keeping BGP state on both active and standby RPs and ensures continuous packet forwarding with minimal packet loss during RP failovers.

How to Configure BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

This section contains the following procedures:

- Configuring a PE Router to Support BGP NSR with SSO, page 4 (required)
- Verifying BGP Support for NSR with SSO, page 10 (optional)

Configuring a PE Router to Support BGP NSR with SSO

Perform this task to enable a PE router to maintain BGP state with CE routers and ensure continuous packet forwarding during a RP switchover or during a planned ISSU. BGP NSR with SSO enables service providers to provide the benefits NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart.

BGP NSR with SSO is supported in BGP peer, BGP peer group, and BGP session template configurations. Perform one of the following tasks in this section on a PE router, depending on whether you want to configure support for BGP NSR with SSO in a peer, a peer group, or a session template configuration:

- Configuring a Peer to Support BGP NSR with SSO, page 5
- Configuring a Peer Group to Support BGP NSR with SSO, page 6
- Configuring Support for BGP NSR with SSO in a Peer Session Template, page 8

Prerequisites

- These tasks assume that you are familiar with BGP peer, BGP peer group, and BGP session template concepts. For more information, see the "Configuring a Basic BGP Network" chapter in the *Cisco IOS IP Routing Configuration Guide*, Release 12.4.
- The active and standby RP must be in SSO mode. For information about configuring SSO mode, see the "Configuring SSO" task in the *Stateful Switchover* document.
- Graceful restart should be enabled on the PE router. For more information about configuring graceful restart, see the "BGP Nonstop Forwarding (NSF) Awareness" document.

  Note We recommend that you enable graceful restart on all BGP peers in the provider core that participate in BGP NSF.

- CE routers must support the route refresh capability. For more information, refer to the "BGP" part of the *Cisco IOS IP Configuration Guide*, Release 12.4.

Restrictions

- This feature is supported only on Cisco 10000 Series PRE2s and Cisco 10000 Series PRE3s.

Configuring a Peer to Support BGP NSR with SSO

Perform this task on a PE router if you want to configure a BGP peer to support BGP NSR with SSO.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **bgp graceful-restart** [**restart-time** *seconds*] [**stalepath-time** *seconds*]
5. **address-family ipv4 vrf** *vrf-name*
6. **neighbor** *ip-address* **remote-as** *autonomous-system-number*
7. **neighbor** *ip-address* **ha-mode sso**
8. **neighbor** *ip-address* **activate**
9. **end**
10. **show ip bgp vpnv4 all sso summary**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Router(config)# router bgp 40000`	Enters router configuration mode for the specified routing process.
Step 4	**bgp graceful-restart** [**restart-time** *seconds*] [**stalepath-time** *seconds*] **Example:** `Router(config-router)# bgp graceful-restart`	Enables the BGP graceful restart capability and BGP NSF awareness. • 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 router and all of its peers (NSF-capable and NSF-aware).

	Command or Action	Purpose
Step 5	**address-family ipv4 vrf** *vrf-name* **Example:** `Router(config-router)# address-family ipv4 vrf test`	Enters address family configuration mode for IPv4 VRF address family sessions. • The **vrf** keyword and *vrf-name* argument specify that IPv4 VRF instance information will be exchanged. **Note** Only the syntax necessary for this task is displayed. For more details, see the *Cisco IOS IP Routing: BGP Command Reference*.
Step 6	**neighbor** *ip-address* **remote-as** *autonomous-system-number* **Example:** `Router(config-router-af)# neighbor 192.168.1.1 remote-as 45000`	Adds the IP address of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.
Step 7	**neighbor** *ip-address* **ha-mode sso** **Example:** `Router(config-router-af)# neighbor 192.168.1.1 ha-mode sso`	Configures the neighbor to support BGP NSR with SSO.
Step 8	**neighbor** *ip-address* **activate** **Example:** `Router(config-router-af)# neighbor testgroup activate`	Enables the neighbor to exchange prefixes for the IPv4 address family with the local router. **Note** By default, neighbors that are defined using the **neighbor remote-as** command in router configuration mode exchange only unicast address prefixes.
Step 9	**end** **Example:** `Router(config-router-af)# end`	Exits address family configuration mode and enters privileged EXEC mode.
Step 10	**show ip bgp vpnv4 all sso summary** **Example:** `Router# show ip bgp vpnv4 all sso summary`	(Optional) Displays the number of BGP neighbors that are in SSO mode.

Configuring a Peer Group to Support BGP NSR with SSO

Perform this task on a PE router if you want to configure a BGP peer group to support BGP NSR with SSO.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **bgp graceful-restart** [**restart-time** *seconds*] [**stalepath-time** *seconds*]
5. **address-family ipv4 vrf** *vrf-name*

6. **neighbor** *peer-group-name* **peer-group**
7. **neighbor** *ip-address* **remote-as** *autonomous-system-number*
8. **neighbor** *ip-address* **peer-group** *peer-group-name*
9. **neighbor** *peer-group-name* **ha-mode sso**
10. **neighbor** *peer-group-name* **activate**
11. **end**
12. **show ip bgp vpnv4 all sso summary**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Router(config)# router bgp 40000`	Enters router configuration mode for the specified routing process.
Step 4	**bgp graceful-restart** [**restart-time** *seconds*] [**stalepath-time** *seconds*] **Example:** `Router(config-router)# bgp graceful-restart`	Enables the BGP graceful restart capability and BGP NSF awareness. • 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 router and all of its peers (NSF-capable and NSF-aware).
Step 5	**address-family ipv4 vrf** *vrf-name* **Example:** `Router(config-router)# address-family ipv4 vrf cisco`	Specifies the IPv4 address family and enters address family configuration mode. • The **vrf** keyword and *vrf-name* argument specify that IPv4 VRF instance information will be exchanged. **Note** Only the syntax necessary for this task is displayed. For more details, see the *Cisco IOS IP Routing Protocols Command Reference*.
Step 6	**neighbor** *peer-group-name* **peer-group** **Example:** `Router(config-router-af)# neighbor testgroup peer-group`	Creates a BGP peer group.

	Command or Action	Purpose
Step 7	**neighbor** *ip-address* **remote-as** *autonomous-system-number* **Example:** `Router(config-router-af)# neighbor 192.168.1.1 remote-as 45000`	Adds the IP address of the neighbor in the specified autonomous system to the IPv4 multiprotocol BGP neighbor table of the local router.
Step 8	**neighbor** *ip-address* **peer-group** *peer-group-name* **Example:** `Router(config-router-af)# neighbor 192.168.1.1 peer-group testgroup`	Assigns the IP address of a BGP neighbor to a BGP peer group.
Step 9	**neighbor** *peer-group-name* **ha-mode sso** **Example:** `Router(config-router-af)# neighbor 192.168.1.1 ha-mode sso`	Configures the BGP peer group to support BGP NSR with SSO.
Step 10	**neighbor** *peer-group-name* **activate** **Example:** `Router(config-router-af)# neighbor testgroup activate`	Enables the neighbor to exchange prefixes for the IPv4 address family with the local router.
Step 11	**end** **Example:** `Router(config-router-af)# end`	Exits address family configuration mode and returns to global configuration mode.
Step 12	**show ip bgp vpnv4 all sso summary** **Example:** `Router# show ip bgp vpnv4 all sso summary`	(Optional) Displays the number of BGP neighbors that are in SSO mode.

Configuring Support for BGP NSR with SSO in a Peer Session Template

Perform this task on a PE router if you want to configure support for BGP NSR with SSO in a BGP peer session template.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **template peer-session** *session-template-name*
5. **ha-mode sso**
6. **exit-peer-session**
7. **end**
8. **show ip bgp template peer-session** [*session-template-name*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Router(config)# router bgp 101`	Enters router configuration mode and creates a BGP routing process.
Step 4	**template peer-session** *session-template-name* **Example:** `Router(config-router)# template peer-session CORE1`	Enters session-template configuration mode and creates a peer session template.
Step 5	**ha-mode sso** **Example:** `Router(config-router-stmp)# ha-mode sso`	Configures the neighbor to support BGP NSR with SSO.
Step 6	**exit-peer-session** **Example:** `Router(config-router-stmp)# exit-peer-session`	Exits session-template configuration mode and returns to router configuration mode.
Step 7	**end** **Example:** `Router(config-router)# end`	Exits router configuration mode and returns to privileged EXEC mode.
Step 8	**show ip bgp template peer-session** [*session-template-name*] **Example:** `Router# show ip bgp template peer-session`	(Optional) Displays locally configured peer session templates. • The output can be filtered to display a single peer policy template with the *session-template-name* argument. This command also supports all standard output modifiers.

What to Do Next

After the peer session template is created, the configuration of the peer session template can be inherited by or applied to another peer session template with the **inherit peer-session** or **neighbor inherit peer-session** command.

For more information about configuring peer session templates, see the "Configuring a Basic BGP Network" chapter in the *Cisco IOS IP Routing Configuration Guide*, Release 12.4.

Verifying BGP Support for NSR with SSO

Perform this optional task to verify BGP NSR with SSO support.

SUMMARY STEPS

1. **enable**
2. **show ip bgp vpnv4 all sso summary**
3. **show ip bgp vpnv4 all neighbors**

DETAILED STEPS

Step 1 **enable**

Enables privileged EXEC mode. Enter your password if prompted.

```
Router> enable
```

Step 2 **show ip bgp vpnv4 all sso summary**

This command is used to display the number of BGP neighbors that are in SSO mode.

The following is sample output from the **show ip bgp vpnv4 all sso summary** command:

```
Router# show ip bgp vpnv4 all sso summary

   Stateful switchover support enabled for 40 neighbors
```

Step 3 **show ip bgp vpnv4 all neighbors**

This command displays VPN address information from the BGP table.

The following is sample output from the **show ip bgp vpnv4 all neighbors** command. The "Stateful switchover support" field indicates whether SSO is enabled or disabled. The "SSO Last Disable Reason" field displays information about the last BGP session that lost SSO capability.

```
Router# show ip bgp vpnv4 all neighbors 10.3.3.3

BGP neighbor is 10.3.3.3,  vrf vrf1,  remote AS 3, external link
  Inherits from template 10vrf-session for session parameters
   BGP version 4, remote router ID 10.1.105.12
   BGP state = Established, up for 04:21:39
   Last read 00:00:05, last write 00:00:09, hold time is 30, keepalive interval is 10
seconds
   Configured hold time is 30, keepalive interval is 10 seconds
   Minimum holdtime from neighbor is 0 seconds
   Neighbor capabilities:
     Route refresh: advertised and received(new)
     Address family IPv4 Unicast: advertised and received
     Stateful switchover support enabled
   Message statistics:
     InQ depth is 0
     OutQ depth is 0
                          Sent       Rcvd
     Opens:                  1          1
     Notifications:          0          0
     Updates:                1          4
     Keepalives:          1534       1532
     Route Refresh:          0          0
     Total:               1536       1537
   Default minimum time between advertisement runs is 30 seconds
```

```
  For address family: VPNv4 Unicast
   Translates address family IPv4 Unicast for VRF vrf1
   BGP table version 25161, neighbor version 25161/0
   Output queue size : 0
   Index 7, Offset 0, Mask 0x80
   7 update-group member
   Inherits from template 10vrf-policy
   Overrides the neighbor AS with my AS before sending updates
   Outbound path policy configured
   Route map for outgoing advertisements is Deny-CE-prefixes
                                Sent       Rcvd
   Prefix activity:             ----       ----
     Prefixes Current:            10         50 (Consumes 3400 bytes)
     Prefixes Total:              10         50
     Implicit Withdraw:            0          0
     Explicit Withdraw:            0          0
     Used as bestpath:           n/a          0
     Used as multipath:          n/a          0

                                  Outbound    Inbound
   Local Policy Denied Prefixes:  --------    -------
     route-map:                        150          0
     AS_PATH loop:                     n/a        760
     Total:                            150        760
   Number of NLRIs in the update sent: max 10, min 10

   Address tracking is enabled, the RIB does have a route to 10.3.3.3
   Address tracking requires at least a /24 route to the peer
   Connections established 1; dropped 0
   Last reset never
   Transport(tcp) path-mtu-discovery is enabled
   TCP session must be opened passively
Connection state is ESTAB, I/O status: 1, unread input bytes: 0 Connection is ECN Disabled
Local host: 10.0.21.1, Local port: 179 Foreign host: 10.0.21.3, Foreign port: 51205
Connection tableid (VRF): 1

Enqueued packets for retransmit: 0, input: 0  mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x1625488):
Timer          Starts    Wakeups            Next
Retrans          1746        210             0x0
TimeWait            0          0             0x0
AckHold          1535       1525             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0
Linger              0          0             0x0

iss: 2241977291  snduna: 2242006573  sndnxt: 2242006573     sndwnd:  13097
irs:  821359845  rcvnxt:  821391670  rcvwnd:      14883  delrcvwnd:   1501

SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms Status Flags: passive open, retransmission
timeout, gen tcbs
   0x1000
Option Flags: VRF id set, always push, md5
```

```
Datagrams (max data segment is 4330 bytes):
Rcvd: 3165 (out of order: 0), with data: 1535, total data bytes: 31824
Sent: 3162 (retransmit: 210 fastretransmit: 0),with data: 1537, total data
bytes: 29300
SSO Last Disable Reason: Application Disable (Active)
```

Troubleshooting Tips

To troubleshoot BGP NSR with SSO, use the following commands in privileged EXEC mode, as needed:

- **debug ip bgp sso**—Displays BGP-related SSO events or debugging information for BGP-related interactions between the active RP and the standby RP. This command is useful for monitoring or troubleshooting BGP sessions on a PE router during an RP switchover or during a planned ISSU.
- **debug ip tcp ha**—Displays TCP HA events or debugging information for TCP stack interactions between the active RP and the standby RP. This is command is useful for troubleshooting SSO-aware TCP connections.
- **show tcp**—Displays the status of TCP connections. The display output will display the SSO capability flag and will indicate the reason that the SSO property failed on a TCP connection.
- **show tcp ha connections**—Displays connection-ID-to-TCP mapping data.

Configuration Examples for BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

This section contains the following configuration example:

- Configuring BGP NSR with SSO: Example, page 12

Configuring BGP NSR with SSO: Example

Figure 2 illustrates a sample BGP NSR with SSO network topology, and the configuration examples that follow show configurations from three routers in the topology: the RR1 router, the PE router, and the CE-1 router.

Note The configuration examples omit some of the configuration required for MPLS VPNs because the purpose of these examples is to illustrate the configuration of BGP NSR with SSO.

Figure 2 BGP NSR with SSO Example Topology

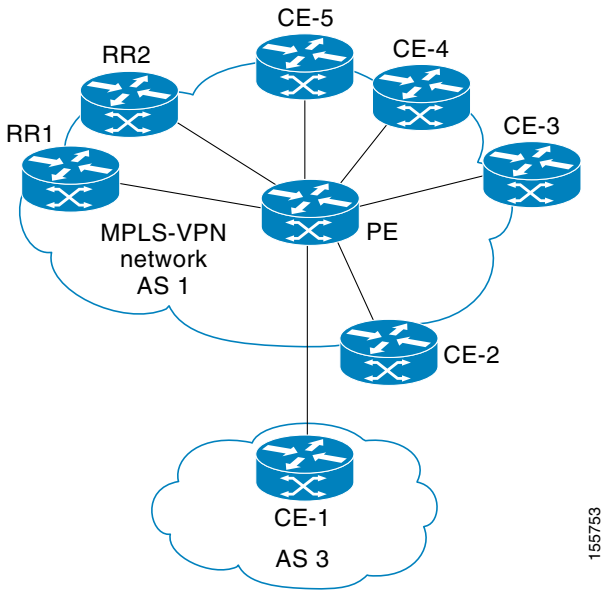

RR1 Configuration

The following example shows the BGP configuration for RR1 in Figure 2. RR1 is configured as a NSF-aware route reflector. In the event of an RP switchover, the PE router uses NSF to maintain the BGP state of the internal peering session with RR1.

```
!
router bgp 1
 no synchronization
 bgp log-neighbor-changes
 bgp graceful-restart restart-time 120
 bgp graceful-restart stalepath-time 360
 bgp graceful-restart
 neighbor 10.2.2.2 remote-as 1
 neighbor 10.2.2.2 update-source Loopback0
 no auto-summary
 !
 address-family vpnv4
 neighbor 10.2.2.2 activate
 neighbor 10.2.2.2 send-community both
 neighbor 10.2.2.2 route-reflector-client
 exit-address-family
 !
```

PE Configuration

The following example shows the BGP NSR with SSO configuration for the PE router in Figure 2. The PE router is configured to support both NSF-awareness and the BGP NSR with SSO capability. In the event of an RP switchover, the PE router uses BGP NSR with SSO to maintain BGP state for the eBGP peering session with the CE-1 router, a CE router in this topology that is not NSF-aware, and uses NSF to maintain BGP state for the iBGP session with RR1. The PE router also detects if any of the other CE routers in the MPLS VPN network are NSF-aware and runs graceful restart with those CE routers.

```
!
router bgp 2
 no synchronization
 bgp log-neighbor-changes
 bgp graceful-restart restart-time 120
 bgp graceful-restart stalepath-time 360
 bgp graceful-restart
 neighbor 10.1.1.1 remote-as 1
 neighbor 10.1.1.1 update-source Loopback0
 no auto-summary
 !
 address-family vpnv4
 neighbor 10.1.1.1 activate
 neighbor 10.1.1.1 send-community both
 exit-address-family
 !
 address-family ipv4 vrf ce-1
 neighbor 10.3.3.3 remote-as 3
 neighbor 10.3.3.3 ha-mode sso
 neighbor 10.3.3.3 activate
 neighbor 10.3.3.3 as-override
 no auto-summary
 no synchronization
 exit-address-family
!
```

CE-1 Configuration

The following example shows the BGP configuration for CE-1 in Figure 2. The CE-1 router is configured as an external peer of the PE router. The CE-1 router is not configured to be NSF-capable or NSF-aware. The CE-1 router, however, does not need to be NSF-capable or NSF-aware to benefit from BGP NSR capabilities on the PE router nor does it need to be upgraded to support BGP NSR.

```
!
router bgp 3
 neighbor 10.2.2.2 remote-as 1
!
```

Additional References

The following sections provide references related to configuring the BGP Support for NSR with SSO feature.

Related Documents

Related Topic	Document Title
BGP concepts and configuration tasks	*Cisco IOS IP Routing: BGP Configuration Guide*
BGP commands: complete command syntax, command mode, command history, command defaults, usage guidelines, and examples	*Cisco IOS IP Routing: BGP Command Reference*
BGP NSF awareness concepts, configuration tasks, and examples	*BGP Nonstop Forwarding (NSF) Awareness*
ISSU concepts, configuration tasks, and examples	*Cisco In Service Software Upgrade Process*
MPLS Layer 3 VPN concepts and configuration tasks	*Cisco IOS Multiprotocol Label Switching Configuration Guide*

Related Topic	Document Title
MPLS Layer 3 VPN commands: complete command syntax, command mode, command history, command defaults, usage guidelines, and examples	*Cisco IOS Multiprotocol Label Switching Command Reference*
NSF and SSO concepts, configuration tasks, and examples	*Cisco Nonstop Forwarding*
SSO concepts, configuration tasks, and examples	*Stateful Switchover*

Standards

Standard	Title
No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature.	—

MIBs

MIB	MIBs Link
No new or modified MIBs are supported by this feature, and support for existing MIBs has not been modified by this feature.	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

RFCs

RFC	Title
draft-ietf-idr-restart-06.txt	*Graceful Restart Mechanism for BGP*

Technical Assistance

Description	Link
The Cisco Technical Support & Documentation website contains thousands of pages of searchable technical content, including links to products, technologies, solutions, technical tips, and tools. Registered Cisco.com users can log in from this page to access even more content.	http://www.cisco.com/techsupport

Command Reference

This section documents new and modified commands.

- debug ip bgp sso
- debug ip tcp ha
- neighbor ha-mode sso
- show ip bgp vpnv4
- show ip bgp vpnv4 all sso summary
- show tcp
- show tcp ha connections

debug ip bgp sso

To display Border Gateway Protocol (BGP)-related stateful switchover (SSO) events or debugging information for BGP-related interactions between the active Route Processor (RP) and the standby RP, use the **debug ip bgp sso** command in privileged EXEC mode. To disable debugging output, use the **no** form of this command.

debug ip bgp sso {**events** | **transactions**} [**detail**]

no debug ip bgp sso {**events** | **transactions**} [**detail**]

Syntax Description

events	Displays BGP-related SSO failures.
transactions	Displays debugging information for failed BGP-related interactions between the active RP and the standby RP
detail	(Optional) Displays detailed debugging information about successful BGP-related SSO operations and successful BGP-related intereactions between the active and the standby RP.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.

Usage Guidelines

The **debug ip bgp sso** command is used to display BGP-related SSO events or debugging information for BGP-related interactions between the active RP and the standby RP. This command is useful for monitoring or troubleshooting BGP sessions on a provider edge (PE) router during an RP switchover or during a planned In-Service Software Upgrade (ISSU).

Examples

The following is sample output from the **debug ip bgp sso** command with the **events** keyword. The following output indicates that the 10.34.32.154 BGP session is no longer SSO capable.

```
*Mar 28 02:29:43.526: BGPSSO: 10.34.32.154 reset SSO and decrement count
```

Tip Use the **show ip bgp vpnv4 all neighbors** command to display the reason that the SSO-capable BGP session has been disabled.

The following is sample output from the **debug ip bgp sso** command with the **transactions** keyword. The following output shows an SSO notification indicating that the SSO capability is pending for 602 BGP neighbors. This notification is generated as the state between the active and standby RP is being synchronized during the bulk synchonization phase of SSO initialization. During this phase, the Transmission Control Blocks (TCBs) must be synchronized with the TCBs on the standby RP before SSO initialization is complete.

```
*Mar 28 02:32:12.102: BGPSSO: tcp sso notify pending for 602 nbrs
```

debug ip tcp ha

To display TCP high availabilty (HA) events or debugging information for TCP stack interactions between the active Route Processor (RP) and the standby RP, use the **debug ip tcp ha** command in privileged EXEC mode. To disable debugging output, use the **no** form of this command.

debug ip tcp ha {**events** | **transactions**} [**detail**]

no debug ip tcp ha {**events** | **transactions**} [**detail**]

Syntax Description

events	Displays TCP HA failures.
transactions	Displays failed TCP stack interactions between the active RP and standby RP.
detail	(Optional) Displays detailed debugging information about successful TCP HA operations and useful informational messages or about successful TCP stack interactions between the active and standby RP.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.

Usage Guidelines

The **debug ip tcp ha** command is used to display TCP stateful switchover (SSO) events or debugging information for TCP stack interactions between the active RP and the standby RP. This is command is useful for troubleshooting SSO-aware TCP connections.

Use the **debug ip tcp ha** command with the **transactions** keyword to display failed TCP stack interactions between the active RP and standby RP. This form of the command displays failed TCP HA messages, RF redundancy-related client-application transactions, IPC client-application transactions, and In-Service Software Upgrade (ISSU) transactions.

Use the **debug ip tcp ha** command with the **transactions** and **detail** keywords to display successful TCP stack interactions between the active and standby RP. This form of the command displays successful TCP HA messages, RF redundancy-related client-application transactions, IPC client-application transactions, and ISSU transactions.

Use the **debug ip tcp ha** command with the **events** keyword to display TCP HA failures. This form of the command displays TCP HA failed encode or decode messages, system resources failures (such as memory allocation failures in the context of TCP HA), failed state changes, and failures that occur when SSO is enabled or disabled.

Use the **debug ip tcp ha** command with the **events** and **detail** keywords to display successful TCP HA operations and useful informational messages. This form of the command displays successful TCP encode or decode messages, state changes, and operations that occur when SSO is enabled or disabled.

Examples

The following is sample output from the **debug ip tcp ha** command with the **transactions** and **detail** keywords. The following output shows packet flow from the active to the standby RP for an established TCP SSO connection:

```
*Feb 19 23:28:23.324: TCPHA: Sending pkt msg, conn_id = 39, seq no = 2727115707
*Feb 19 23:28:23.324: TCPHA: Sending pkt msg, conn_id = 396, seq no = 2959469308
*Feb 19 23:28:23.324: TCPHA: Sending pkt msg, conn_id = 41, seq no = 1270243395
*Feb 19 23:28:23.932: TCPHA: Sending pkt msg, conn_id = 42, seq no = 974255741
*Feb 19 23:28:23.932: TCPHA: Sending pkt msg, conn_id = 475, seq no = 3059612402
*Feb 19 23:28:24.544: TCPHA: Sending dummy pkt to standby; cid=109, size=19

*Feb 19 23:28:42.976: TCPHA: Recd IPC msg len 24, type 3
*Feb 19 23:28:42.976: TCPHA: Recd IPC msg len 24, type 3
*Feb 19 23:28:43.172: TCPHA: Recd IPC msg len 79, type 2
*Feb 19 23:28:43.172: TCPHA: Recd IPC msg len 79, type
```

neighbor ha-mode sso

To configure a Border Gateway Protocol (BGP) neighbor to support BGP Nonstop Routing (NSR) with stateful switchover (SSO), use the **neighbor ha-mode sso** command in the appropriate command mode. To remove the configuration, use the **no** form of this command.

neighbor *ip-address* **ha-mode sso**

no neighbor *ip-address* **ha-mode sso**

Syntax Description

ip-address	IP address of the neighboring router.

Command Default

BGP NSR with SSO support is disabled.

Command Modes

Address family configuration
Session-template configuration

Command History

Release	Modification
12.2(28)SB	This command was introduced.

Usage Guidelines

The **neighbor ha-mode sso** command is used to configure a BGP neighbor to support BGP NSR with SSO. BGP NSR with SSO is disabled by default.

BGP NSR with SSO is supported in BGP peer, BGP peer group, and BGP session template configurations. To configure BGP NSR with SSO in BGP peer and BGP peer group configurations, use the **neighbor ha-mode sso** command in address family configuration mode for IPv4 VRF address family BGP peer sessions. To include support for Cisco BGP NSR with SSO in a peer session template, use the **ha-mode sso** command in session-template configuration mode.

Examples

The following example shows how to configure a BGP neighbor to support SSO:

```
Router(config-router-af)# neighbor 10.3.32.154 ha-mode sso
```

Related Commands

Command	Description
show ip bgp vpnv4	Displays VPN address information from the BGP table.
show ip bgp vpnv4 all sso summary	Displays the number of BGP neighbors that support SSO.

show ip bgp vpnv4

To display Virtual Private Network Version 4 (VPNv4) address information from the Border Gateway Protocol (BGP) table, use the **show ip bgp vpnv4** command in user EXEC or privileged EXEC mode.

show ip bgp vpnv4 {**all** | **rd** *route-distinguisher* | **vrf** *vrf-name*} [**rib-failure**] [*ip-prefix/length* [**longer-prefixes**]] [*network-address* [*mask*] [**longer-prefixes**]] [**cidr-only**] [**community**] [**community-list**] [**dampened-paths**] [**filter-list**] [**flap-statistics**] [**inconsistent-as**] [**neighbors**] [**paths** [*line*]] [**peer-group**] [**quote-regexp**] [**regexp**] [**summary**] [**labels**]

Syntax Description

all	Displays the complete VPNv4 database.
rd *route-distinguisher*	Displays Network Layer Reachability Information (NLRI) prefixes that match the named route distinguisher.
vrf *vrf-name*	Displays NLRI prefixes associated with the named VPN routing and forwarding (VRF) instance.
rib-failure	(Optional) Displays BGP routes that failed to install in the VRF table.
ip-prefix/length	(Optional) The IP prefix address (in dotted decimal format) and the length of the mask (0 to 32). The slash mark must be included.
longer-prefixes	(Optional) Displays the entry, if any, that exactly matches the specified prefix parameter and all entries that match the prefix in a "longest-match" sense. That is, prefixes for which the specified prefix is an initial substring.
network-address	(Optional) The IP address of a network in the BGP routing table.
mask	(Optional) The mask of the network address, in dotted decimal format.
cidr-only	(Optional) Displays only routes that have nonclassful net masks.
community	(Optional) Displays routes that match this community.
community-list	(Optional) Displays routes that match this community list.
dampened-paths	(Optional) Displays paths suppressed because of dampening (BGP route from peer is up and down).
filter-list	(Optional) Displays routes that conform to the filter list.
flap-statistics	(Optional) Displays flap statistics of routes.
inconsistent-as	(Optional) Displays only routes that have inconsistent autonomous systems of origin.
neighbors	(Optional) Displays details about TCP and BGP neighbor connections.
paths	(Optional) Displays path information.
line	(Optional) A regular expression to match the BGP autonomous system paths.
peer-group	(Optional) Displays information about peer groups.
quote-regexp	(Optional) Displays routes that match the autonomous system path regular expression.
regexp	(Optional) Displays routes that match the autonomous system path regular expression.

summary	(Optional) Displays BGP neighbor status.
labels	(Optional) Displays incoming and outgoing BGP labels for each NLRI prefix.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.2(2)T	The output of the **show ip bgp vpnv4 all** *ip-prefix* command was enhanced to display attributes including multipaths and a best path to the specified network.
12.0(21)ST	The **tags** keyword was replaced by the **labels** keyword to conform to the MPLS guidelines. This command was integrated into Cisco IOS Release 12.0(21)ST.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.0(27)S	The output of the **show ip bgp vpnv4 all labels** command was enhanced to display explicit-null label information.
12.3	The **rib-failure** keyword was added for VRFs.
12.2(22)S	The output of the **show ip bgp vpnv4 vrf** *vrf-name* **labels** command was modified so that directly connected VRF networks no longer display as aggregate; no label appears instead.
12.2(25)S	This command was updated to display MPLS VPN nonstop forwarding information.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB and implemented on the Cisco 10000 series router. The display output was modified to indicate whether BGP Nonstop Routing (NSR) with stateful switchover (SSO) is enabled and the reason the last BGP lost SSO capability.

Usage Guidelines

Use this command to display VPNv4 information from the BGP database. The **show ip bgp vpnv4 all** command displays all available VPNv4 information. The **show ip bgp vpnv4 summary** command displays BGP neighbor status. The **show ip bgp vpnv4 all labels** command displays explicit-null label information.

Examples

The following example shows output for all available VPNv4 information in a BGP routing table:

```
Router# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 10.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i10.6.6.6/32       10.0.0.21                11    100      0 ?
*> 10.7.7.7/32       10.150.0.2               11         32768 ?
*>i10.69.0.0/30      10.0.0.21                 0    100      0 ?
*> 10.150.0.0/24     0.0.0.0                   0         32768 ?
```

Table 1 describes the significant fields shown in the display.

Table 1 show ip bgp vpnv4 all Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows how to display a table of labels for NLRI prefixes that have a route distinguisher value of 100:1.

```
Router# show ip bgp vpnv4 rd 100:1 labels

Network            Next Hop       In label/Out label
Route Distinguisher: 100:1 (vrf1)
   10.0.0.0         10.20.0.60      34/nolabel
   10.0.0.0         10.20.0.60      35/nolabel
   10.0.0.0         10.20.0.60      26/nolabel
                    10.20.0.60      26/nolabel
   10.0.0.0         10.15.0.15      nolabel/26
```

Table 2 describes the significant fields shown in the display.

Table 2 show ip bgp vpnv4 rd labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Specifies the BGP next hop address.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next hop router.

The following example shows VPNv4 routing entries for the VRF named vpn1:

```
Router# show ip bgp vpnv4 vrf vpn1

BGP table version is 18, local router ID is 10.14.14.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP,? - incomplete

Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:101 (default for vrf vpn1)
*>i10.6.6.6/32       10.0.0.21                11    100      0 ?
*> 10.7.7.7/32       10.150.0.2               11         32768 ?
```

```
*>i10.69.0.0/30      10.0.0.21               0    100      0 ?
*> 10.150.0.0/24      0.0.0.0                0         32768 ?
*> 10.0.0.1/32       10.150.0.2             11         32768 ?
*>i10.0.0.3/32       10.0.0.21              11    100      0 ?
```

Table 3 describes the significant fields shown in the display.

Table 3 show ip bgp vpnv4 vrf Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
Metric	Displays the BGP metric.
LocPrf	Displays the local preference.
Weight	Displays the BGP weight.
Path	Displays the BGP path per route.

The following example shows attributes for network 10.22.22.0 that include multipaths and a best path:

```
Router# show ip bgp vpnv4 all 10.22.22.0

BGP routing table entry for 10:1:10.22.22.0/24, version 50
Paths:(6 available, best #1)
Multipath:iBGP
  Advertised to non peer-group peers:
  10.1.12.12
  22
    10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath, best
      Extended Community:RT:100:1
      Originator:10.0.0.8, Cluster list:10.1.1.44
  22
    10.22.1.9 (metric 11) from 10.11.1.2 (10.0.0.9)
      Origin IGP, metric 0, localpref 100, valid, internal, multipath
      Extended Community:RT:100:1
      Originator:10.0.0.9, Cluster list:10.1.1.22
```

Table 4 describes the significant fields shown in the display.

Table 4 show ip bgp vpnv4 all network-address Field Descriptions

Field	Description
BGP routing table entry for... version	Internal version number of the table. This number is incremented whenever the table changes.
Paths	Number of autonomous system paths to the specified network. If multiple paths exist, one of the multipaths is designated the best path.
Multipath	Indicates the maximum paths configured (iBGP or eBGP).
Advertised to non peer-group peers	IP address of the BGP peers to which the specified route is advertised.
10.22.7.8 (metric 11) from 10.11.3.4 (10.0.0.8)	Indicates the next hop address and the address of the gateway that sent the update.

Table 4 show ip bgp vpnv4 all network-address Field Descriptions (continued)

Field	Description
Origin	Indicates the origin of the entry. It can be one of the following values: • IGP—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. • incomplete—Entry originated from other than an IGP or Exterior Gateway Protocol (EGP) and was advertised with the **redistribute** router configuration command. • EGP—Entry originated from an EGP.
metric	If shown, the value of the interautonomous system metric.
localpref	Local preference value as set with the **set local-preference route-map** configuration command. The default value is 100.
valid	Indicates that the route is usable and has a valid set of attributes.
internal/external	The field is *internal* if the path is learned via iBGP. The field is *external* if the path is learned via eBGP.
multipath	One of multiple paths to the specified network.
best	If multiple paths exist, one of the multipaths is designated the best path and this path is advertised to neighbors.
Extended Community	Route Target value associated with the specified route.
Originator	The router ID of the router from which the route originated when route reflector is used.
Cluster list	The router ID of all the route reflectors that the specified route has passed through.

The following example shows routes that BGP could not install in the VRF table:

```
Router# show ip bgp vpnv4 vrf xyz rib-failure

Network            Next Hop                      RIB-failure   RIB-NH Matches
Route Distinguisher: 2:2 (default for vrf bar)
10.1.1.2/32        10.100.100.100      Higher admin distance              No
10.111.111.112/32  10.9.9.9            Higher admin distance             Yes
```

Table 5 describes the significant fields shown in the display.

Table 5 show ip bgp vpnv4 vrf rib-failure Field Descriptions

Field	Description
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.

Table 5 show ip bgp vpnv4 vrf rib-failure Field Descriptions (continued)

Field	Description
RIB-failure	Cause of the Routing Information Base (RIB) failure. Higher admin distance means that a route with a better (lower) administrative distance, such as a static route, already exists in the IP routing table.
RIB-NH Matches	Route status that applies only when Higher admin distance appears in the RIB-failure column and the **bgp suppress-inactive** command is configured for the address family being used. There are three choices: • Yes—Means that the route in the RIB has the same next hop as the BGP route or that the next hop recurses down to the same adjacency as the BGP next hop. • No—Means that the next hop in the RIB recurses down differently from the next hop of the BGP route. • n/a—Means that the **bgp suppress-inactive** command is not configured for the address family being used.

The following example shows the information displayed on the active and standby route processors when they are configured for MPLS VPN nonstop forwarding.

Active Route Processor

```
Router# show ip bgp vpnv4 all labels

Network          Next Hop    In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0     16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0     17/aggregate(vpn1)
Route Distinguisher: 609:1 (vpn0)
10.13.13.13/32   0.0.0.0     18/aggregate(vpn0)
```

```
Router# show ip bgp vpnv4 vrf vpn1 labels

Network          Next Hop    In label/Out label
Route Distinguisher: 100:1 (vpn1)
10.12.12.12/32   0.0.0.0     16/aggregate(vpn1)
10.0.0.0/8       0.0.0.0     17/aggregate(vpn1)
```

Standby Route Processor

```
Router# show ip bgp vpnv4 all labels

Network        Masklen    In label
Route Distinguisher: 100:1
10.12.12.12    /32        16
10.0.0.0       /8         17
Route Distinguisher: 609:1
10.13.13.13    /32        18
```

```
Router# show ip bgp vpnv4 vrf vpn1
labels

Network         Masklen    In label
Route Distinguisher: 100:1
10.12.12.12     /32        16
10.0.0.0        /8         17
```

Table 6 describes the significant fields shown in the display.

Table 6 show ip bgp vpn4 labels Field Descriptions

Field	Description
Network	The network address from the BGP table.
Next Hop	The BGP next hop address.
In label	The label (if any) assigned by this router.
Out label	The label assigned by the BGP next hop router.
Masklen	The mask length of the network address.

The following example displays output, including the explicit-null label, from the **show ip bgp vpnv4 all labels** command on a CSC-PE router:

```
Router# show ip bgp vpnv4 all labels

   Network          Next Hop      In label/Out label
Route Distinguisher: 100:1 (v1)
   10.0.0.0/24      10.0.0.0        19/aggregate(v1)
   10.0.0.1/32      10.0.0.0        20/nolabel
   10.1.1.1/32      10.0.0.0        21/aggregate(v1)
   10.10.10.10/32   10.0.0.1        25/exp-null
   10.168.100.100/32
                    10.0.0.1        23/exp-null
   10.168.101.101/32
                    10.0.0.1        22/exp-null
```

Table 7 describes the significant fields shown in the display.

Table 7 show ip bgp vpnv4 all labels Field Descriptions

Field	Description
Network	Displays the network address from the BGP table.
Next Hop	Displays the address of the BGP next hop.
In label	Displays the label (if any) assigned by this router.
Out label	Displays the label assigned by the BGP next hop router.
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.

The following example displays separate router IDs for each VRF in the output from an image in Cisco IOS Release 12.2(33)SRA and later releases with the Per-VRF Assignment of BGP Router ID feature configured. The router ID is shown next to the VRF name.

```
Router# show ip bgp vpnv4 all

BGP table version is 5, local router ID is 172.17.1.99
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf vrf_trans) VRF Router ID 10.99.1.2
*> 192.168.4.0      0.0.0.0                  0         32768 ?
Route Distinguisher: 42:1 (default for vrf vrf_user) VRF Router ID 10.99.1.1
*> 192.168.5.0      0.0.0.0                  0         32768 ?
```

Table 8 describes the significant fields shown in the display.

Table 8 show ip bgp vpnv4 all (VRF Router ID) Field Descriptions

Field	Description
Route Distinguisher	Displays an 8-byte value added to an IPv4 prefix to create a VPN IPv4 prefix.
vrf	Name of the VRF.
VRF Router ID	Router ID for the VRF.

Related Commands

Command	Description
show ip vrf	Displays the set of defined VRFs and associated interfaces.

show ip bgp vpnv4 all sso summary

To display information about Border Gateway Protocol (BGP) peers that support BGP Nonstop Routing (NSR) with stateful switchover (SSO), use the **show ip bgp vpn4 sso summary** command in privileged EXEC mode.

show ip bgp vpnv4 all sso summary

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.

Usage Guidelines

The **show ip bgp vpnv4 all sso summary** command is used to display the number of BGP neighbors that are in SSO mode.

Examples

The following is sample output from the **show ip bgp vpnv4 all sso summary** command:

```
Router# show ip bgp vpnv4 all sso summary

   Stateful switchover support enabled for 40 neighbors
```

Table 9 describes the significant fields shown in the display.

Table 9 show ip bgp vpnv4 all sso summary Field Descriptions

Field	Description
Stateful Switchover support enabled for	Indicates the number of BGP neighbors that are in SSO mode.

Related Commands

Command	Description
neighbor ha-mode sso	Configures a BGP neighbor to support SSO.

show tcp

To display the status of Transmission Control Protocol (TCP) connections when Cisco IOS or Cisco IOS Software Modularity images re running, use the **show tcp** command in user EXEC or privileged EXEC mode.

show tcp [*line-number*] [**tcb** *address*]

Syntax Description

line-number	(Optional) Absolute line number of the line for which you want to display Telnet connection status.
tcb	(Optional) Specifies the transmission control block (TCB) of the ECN-enabled connection that you want to display.
address	(Optional) TCB hexadecimal address. The valid range is from 0x0 to 0xFFFFFFFF.

Command Modes

User EXEC
Privileged EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.3(7)T	The **tcb** keyword and *address* argument were added.
12.4(2)T	The output is enhanced to display status and option flags.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB. The display output was modified to include the SSO capability flag and to indicate the reason that the SSO property failed on a TCP connection.
12.2(18)SXF4	This command was integrated into Cisco IOS Release 12.2(18)SXF4 to support Software Modularity images.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.

Examples

Example output varies between Cisco IOS software images and Cisco IOS Software Modularity software images. To view the appropriate output, choose one of the following sections:

- Cisco IOS Software
- Cisco IOS Software Modularity

Cisco IOS Software

The following is sample output that displays the status and option flags:

```
Router# show tcp
.
.
.
Status Flags: passive open, active open, retransmission timeout, app closed

Option Flags: vrf id set
```

```
IP Precedence value: 6
.
.
.
SRTT: 273 ms, RTTO: 490 ms, RTV: 217 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
 Status Flags: active open, retransmission timeout
 Option Flags: vrf id set
 IP Precedence value: 6
```

Table 10 contains the types of flags, all possible command output enhancements, and descriptions. See Table 11 through Table 15 for descriptions of the other fields in the sample output.

Table 10 Type of Flags, All Possible Output Enhancements, and Descriptions

Type of Flag	Output Enhancement	Description
Status		
	Passive open	Set if passive open was done.
	Active open	Set if active open was done.
	Retransmission timeout	Set if retransmission timeout aborts.
	Net output pending	Output to network is pending.
	Wait for FIN	Wait for FIN to be acknowledged.
	App closed	Application has closed the TCB.
	Sync listen	Listen and establish a handshake.
	Gen tcbs	TCBs are generated as passive listener.
	Path mtu discovery	Path maximum transmission unit (MTU) discovery is enabled.
	Half closed	TCB is half closed.
	Timestamp echo present	Echo segment is present.
	Stopped reading	Read half is shut down.
Option		
	VRF id set	Set if connection has a VRF table identifier.
	Idle user	Set if the connection is idle.
	Sending urgent data	Set if urgent data is being sent.
	Keepalive running	Set if keepalive timer is running, or if an Explicit Congestion Notification (ECN)-enabled connection, or a TCB address bind is in effect.
	Nagle	Set if performing the Nagle algorithm.
	Always push	All packets and full-sized segments (internal use) are pushed.
	Path mtu capable	Path MTU discovery is configured.
	MD5	Message digest 5 (MD) messages are generated.
	Urgent data removed	Urgent data is removed.
	SACK option permitted	Peer permits a selective acknowledgment (SACK) option.

Table 10 Type of Flags, All Possible Output Enhancements, and Descriptions (continued)

Type of Flag	Output Enhancement	Description
	Timestamp option used	Time-stamp option is in use.
	Reuse local address	Local address can be reused.
	Non-blocking reads	Nonblocking TCP is read.
	Non-blocking writes	Nonblocking TCP is written.
	No delayed ACK	No TCP delayed acknowledgment is sent.
	Win-scale	Peer permits window scaling.
	Linger option set	The linger-on close option is set.

The following is sample output from the **show tcp** command:

```
Router# show tcp

tty0, connection 1 to host cider
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 172.31.232.17, Local port: 11184
Foreign host: 172.31.1.137, Foreign port: 23

Enqueued packets for retransmit: 0, input: 0, saved: 0

Event Timers (current time is 67341276):
Timer:       Retrans    TimeWait    AckHold    SendWnd  KeepAlive
Starts:           30           0         32          0          0
Wakeups:           1           0         14          0          0
Next:              0           0          0          0          0

iss:   67317172  snduna:   67317228  sndnxt:   67317228     sndwnd:   4096
irs: 1064896000  rcvnxt: 1064897597  rcvwnd:       2144  delrcvwnd:      0

SRTT: 317 ms, RTTO: 900 ms, RTV: 133 ms, KRTT: 0 ms
minRTT: 4 ms, maxRTT: 300 ms, ACK hold: 300 ms
Flags: higher precedence, idle user, retransmission timeout
Datagrams (max data segment is 536 bytes):
Rcvd: 41 (out of order: 0), with data: 34, total data bytes: 1596
Sent: 57 (retransmit: 1), with data: 35, total data bytes: 55
```

Table 11 describes the first five lines of output shown in the above display.

Table 11 show tcp Field Descriptions—First Section of Output

Field	Description
tty	Identifying number of the line.
connection	Identifying number of the TCP connection.
to host	Name of the remote host to which the connection has been made.

Table 11 show tcp Field Descriptions—First Section of Output (continued)

Field	Description
Connection state is	A connection progresses through a series of states during its lifetime. The states that follow are shown in the order in which a connection progresses through them. • LISTEN—Waiting for a connection request from any remote TCP and port. • SYNSENT—Waiting for a matching connection request after having sent a connection request. • SYNRCVD—Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTAB—Indicates an open connection; data received can be delivered to the user. This is the normal state for the data transfer phase of the connection. • FINWAIT1—Waiting for a connection termination request from the remote TCP or an acknowledgment of the connection termination request previously sent. • FINWAIT2—Waiting for a connection termination request from the remote TCP host. • CLOSEWAIT—Waiting for a connection termination request from the local user. • CLOSING—Waiting for a connection termination request acknowledgment from the remote TCP host. • LASTACK—Waiting for an acknowledgment of the connection termination request previously sent to the remote TCP host. • TIMEWAIT—Waiting for enough time to pass to be sure that the remote TCP host has received the acknowledgment of its connection termination request. • CLOSED—Indicates no connection state at all. • For more information about TCBs, see RFC 793, *Transmission Control Protocol Functional Specification.*
I/O status	Number that describes the current internal status of the connection.
unread input bytes	Number of bytes that the lower-level TCP processes have read but that the higher-level TCP processes have not yet processed.
Local host	IP address of the network server.
Local port	Local port number, as derived from the following equation: *line-number* + (512 * *random-number*). (The line number uses the lower nine bits; the other bits are random.)
Foreign host	IP address of the remote host to which the TCP connection has been made.
Foreign port	Destination port for the remote host.

Table 11 show tcp Field Descriptions—First Section of Output (continued)

Field	Description
Enqueued packets for retransmit	Number of packets that are waiting on the retransmit queue. These are packets on this TCP connection that have been sent but that have not yet been acknowledged by the remote TCP host.
input	Number of packets that are waiting on the input queue to be read by the user.
saved	Number of received out-of-order packets that are waiting for all packets in the datagram to be received before they enter the input queue. For example, if packets 1, 2, 4, 5, and 6 have been received, packets 1 and 2 would enter the input queue, and packets 4, 5, and 6 would enter the saved queue.

Note Use the **show tcp brief** command to display information about the ECN-enabled connections.

The following line of output shows the current elapsed time according to the system clock of the local host. The time shown is the number of milliseconds since the system started.

```
Event Timers (current time is 67341276):
```

The following lines of output display the number of times that various local TCP timeout values were reached during this connection. In this example, the local host re-sent data 30 times because it received no response from the remote host, and it sent an acknowledgment many more times because there was no data.

```
Timer:        Retrans    TimeWait     AckHold     SendWnd     Keepalive     GiveUp     PmtuAger
Starts:            30           0          32           0             0          0            0
Wakeups:            1           0          14           0             0          0            0
Next:               0           0           0           0             0          0            0
```

Table 12 describes the fields in the above lines of output.

Table 12 show tcp Field Descriptions—Second Section of Output

Field	Description
Timer	Names of the timer types in the output.
Starts	Number of times that the timer has been triggered during this connection.
Wakeups	Number of keepalives sent without receiving any response. (This field is reset to zero when a response is received.)
Next	System clock setting that triggers a timer for the next time an event (for example, TimeWait, AckHold, SendWnd, etc.) occurs.
Retrans	Retransmission timer is used to time TCP packets that have not been acknowledged and that are waiting for retransmission.
TimeWait	A time-wait timer ensures that the remote system receives a request to disconnect a session.
AckHold	An acknowledgment timer delays the sending of acknowledgments to the remote TCP in an attempt to reduce network use.

Table 12 show tcp Field Descriptions—Second Section of Output (continued)

Field	Description
SendWnd	A send-window timer ensures that there is no closed window due to a lost TCP acknowledgment.
KeepAlive	A keepalive timer controls the transmission of test messages to the remote device to ensure that the link has not been broken without the knowledge of the local device.
GiveUp	A give-up timer determines the amount of time a local host will wait for an acknowledgment (or other appropriate reply) of a transmitted message after the the maximum number of retransmissions has been reached. If the timer expires, the local host gives up retransmission attempts and declares the connection dead.
PmtuAger	A path MTU (PMTU) age timer is an interval that displays how often TCP estimates the PMTU with a larger maximum segment size (MSS). When the age timer is used, TCP path MTU becomes a dynamic process. If the MSS is smaller than what the peer connection can manage, a larger MSS is tried every time the age timer expires. The discovery process stops when the send MSS is as large as the peer negotiated or the timer has been manually disabled by being set to infinite.

The following lines of output display the sequence numbers that TCP uses to ensure sequenced, reliable transport of data. The local host and remote host each use these sequence numbers for flow control and to acknowledge receipt of datagrams.

```
iss:   67317172  snduna:   67317228  sndnxt:   67317228     sndwnd:   4096
irs: 1064896000  rcvnxt: 1064897597  rcvwnd:       2144  delrcvwnd:      0
```

Table 13 describes the fields shown in the display above.

Table 13 show tcp Field Descriptions—Sequence Numbers

Field	Description
iss	Initial send sequence number.
snduna	Last send sequence number that the local host sent but for which it has not received an acknowledgment.
sndnxt	Sequence number that the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number that the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window—data that the local host has read from the connection but has not yet subtracted from the receive window that the host has advertised to the remote host. The value in this field gradually increases until it is larger than a full-sized packet, at which point it is applied to the rcvwnd field.

The following lines of output display values that the local host uses to keep track of transmission times so that TCP can adjust to the network that it is using.

```
SRTT: 317 ms, RTTO: 900 ms, RTV: 133 ms, KRTT: 0 ms
minRTT: 4 ms, maxRTT: 300 ms, ACK hold: 300 ms
Flags: higher precedence, idle user, retransmission timeout
```

Table 14 describes the significant fields shown in the output above.

Table 14 show tcp Field Descriptions—Line Beginning with "SRTT"

Field	Description
SRTT	A calculated smoothed round-trip timeout.
RTTO	Round-trip timeout.
RTV	Variance of the round-trip time.
KRTT	New round-trip timeout (using the Karn algorithm). This field separately tracks the round-trip time of packets that have been re-sent.
minRTT	Smallest recorded round-trip timeout (hard-wire value used for calculation).
maxRTT	Largest recorded round-trip timeout.
ACK hold	Time for which the local host will delay an acknowledgment in order to add data to it.
Flags	Properties of the connection.

Note For more information on the above fields, see *Round Trip Time Estimation*, P. Karn & C. Partridge, ACM SIGCOMM-87, August 1987.

The following lines of output display the number of datagrams that are transported with data.

```
Datagrams (max data segment is 536 bytes):
Rcvd: 41 (out of order: 0), with data: 34, total data bytes: 1596
Sent: 57 (retransmit: 1), with data: 35, total data bytes: 55
```

Table 15 describes the significant fields shown in the last lines of the **show tcp** command output.

Table 15 show tcp Field Descriptions—Last Section of Output

Field	Description
Rcvd	Number of datagrams that the local host has received during this connection (and the number of these datagrams that were out of order).
with data	Number of these datagrams that contained data.
total data bytes	Total number of bytes of data in these datagrams.
Sent	Number of datagrams that the local host sent during this connection (and the number of these datagrams that needed to be re-sent).
with data	Number of these datagrams that contained data.
total data bytes	Total number of bytes of data in these datagrams.

The following is sample output from the **show tcp tcb** command that displays detailed information by hexadecimal address about an ECN-enabled connection:

```
Router# show tcp tcb 0x62CD2BB8

Connection state is LISTEN, I/O status: 1, unread input bytes: 0
Connection is ECN enabled
Local host: 10.10.10.1, Local port: 179
Foreign host: 10.10.10.2, Foreign port: 12000

Enqueued packets for retransmit: 0, input: 0 mis-ordered: 0 (0 bytes)

Event Timers (current time is 0x4F31940):
Timer          Starts    Wakeups            Next
Retrans             0          0             0x0
TimeWait            0          0             0x0
AckHold             0          0             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss:          0 snduna:          0 sndnxt:          0     sndwnd:      0
irs:          0 rcvnxt:          0 rcvwnd:       4128  delrcvwnd:      0

SRTT: 0 ms, RTTO: 2000 ms, RTV: 2000 ms, KRTT: 0 ms
minRTT: 60000 ms, maxRTT: 0 ms, ACK hold: 200 ms
Flags: passive open, higher precedence, retransmission timeout

TCB is waiting for TCP Process (67)

Datagrams (max data segment is 516 bytes):
Rcvd: 6 (out of order: 0), with data: 0, total data bytes: 0
Sent: 0 (retransmit: 0, fastretransmit: 0), with data: 0, total data
bytes: 0
```

Cisco IOS Software Modularity

The following is sample output from the **show tcp tcb** command from a Software Modularity image:

```
Router# show tcp tcb 0x1059C10

Connection state is ESTAB, I/O status: 0, unread input bytes: 0
Local host: 10.4.2.32, Local port: 23
Foreign host: 10.4.2.39, Foreign port: 11000
VRF table id is: 0

Current send queue size: 0 (max 65536)
Current receive queue size: 0 (max 32768)  mis-ordered: 0 bytes

Event Timers (current time is 0xB9ACB9):
Timer          Starts    Wakeups            Next(msec)
Retrans             6          0                 0
SendWnd             0          0                 0
TimeWait            0          0                 0
AckHold             8          4                 0
KeepAlive          11          0           7199992
PmtuAger            0          0                 0
GiveUp              0          0                 0
Throttle            0          0                 0
```

```
irs:    1633857851  rcvnxt: 1633857890  rcvadv: 1633890620  rcvwnd:  32730
iss:    4231531315  snduna: 4231531392  sndnxt: 4231531392  sndwnd:   4052
sndmax: 4231531392  sndcwnd:     10220

SRTT: 84 ms,  RTTO: 650 ms,  RTV: 69 ms,  KRTT: 0 ms
minRTT: 0 ms,  maxRTT: 200 ms, ACK hold: 200 ms

Keepalive time: 7200 sec, SYN wait time: 75 sec
Giveup time: 0 ms, Retransmission retries: 0, Retransmit forever: FALSE

State flags: none

Feature flags: Nagle

Request flags: none
Window scales: rcv 0, snd 0, request rcv 0, request snd 0
Timestamp option: recent 0, recent age 0, last ACK sent          0

Datagrams (in bytes): MSS 1460, peer MSS 1460, min MSS 1460, max MSS 1460
Rcvd: 14 (out of order: 0), with data: 10, total data bytes: 38
Sent: 10 (retransmit: 0, fastretransmit: 0), with data: 5, total data bytes: 76

Header prediction hit rate: 72 %

Socket states: SS_ISCONNECTED, SS_PRIV

Read buffer flags: SB_WAIT, SB_SEL, SB_DEL_WAKEUP
Read notifications: 4

Write buffer flags: SB_DEL_WAKEUP
Write notifications: 0
Socket status: 0
```

Related Commands

Command	Description
show tcp brief	Displays a concise description of TCP connection endpoints.

show tcp ha connections

To display connection-ID-to-TCP mapping data, use the **show tcp ha connections** command in privileged EXEC mode.

show tcp ha connections

Syntax Description

This command has no arguments or keywords.

Command Modes

Privileged EXEC

Command History

Release	Modification
12.2(28)SB	This command was introduced.

Usage Guidelines

The **show tcp ha connections** command is used to display connection-ID-to-TCP mapping data.

Examples

The following is sample output from the **show tcp ha connections** command:

```
Router# show tcp ha connections

SSO enabled for 40 connections
TCB       Local Address           Foreign Address        (state)    Conn Id
71EACE60  10.0.56.1.179           10.0.56.3.58671         ESTAB      37
71EA9320  10.0.53.1.179           10.0.53.3.58659         ESTAB      34
71EA35F8  10.0.41.1.179           10.0.41.3.58650         ESTAB      22
71A21FE0  10.0.39.1.179           10.0.39.3.58641         ESTAB      20
71EAA6E0  10.0.54.1.179           10.0.54.3.58663         ESTAB      35
71EA2238  10.0.40.1.179           10.0.40.3.58646         ESTAB      21
71EABAA0  10.0.55.1.179           10.0.55.3.58667         ESTAB      36
71EAE710  10.0.28.1.179           10.0.28.3.58676         ESTAB      9
71EA2728  10.0.50.1.179           10.0.50.3.58647         ESTAB      31
720541D8  10.0.49.1.179           10.0.49.3.58642         ESTAB      30
71EAA1F0  10.0.44.1.179           10.0.44.3.58662         ESTAB      25
2180B3A8  10.0.33.1.179           10.0.33.3.58657         ESTAB      14
71EAB5B0  10.0.45.1.179           10.0.45.3.58666         ESTAB      26
21809FE8  10.0.32.1.179           10.0.32.3.58653         ESTAB      13
71EA8E30  10.0.43.1.179           10.0.43.3.58658         ESTAB      24
71EAD350  10.0.27.1.179           10.0.27.3.58672         ESTAB      8
2180A9C8  10.0.52.1.179           10.0.52.3.58655         ESTAB      33
2180A4D8  10.0.42.1.179           10.0.42.3.58654         ESTAB      23
71EABF90  10.0.26.1.179           10.0.26.3.58668         ESTAB      7
71EA3AE8  10.0.51.1.179           10.0.51.3.58651         ESTAB      32
720546C8  10.0.59.1.179           10.0.59.3.58643         ESTAB      40
```

Table 16 describes the significant fields shown in the display.

Table 16 show tcp ha connections Field Descriptions

Field	Description
SSO enabled for	Displays the number of TCP connections that support BGP Nonstop Routing (NSR) with SSO.
TCB	An internal identifier for the endpoint.
Local Address	The local IP address and port.
Foreign Address	The foreign IP address and port (at the opposite end of the connection).
(state)	TCP connection state. A connection progresses through a series of states during its lifetime. The states that follow are shown in the order in which a connection progresses through them. • LISTEN—Waiting for a connection request from any remote TCP and port. • SYNSENT—Waiting for a matching connection request after having sent a connection request. • SYNRCVD—Waiting for a confirming connection request acknowledgment after having both received and sent a connection request. • ESTAB—Indicates an open connection; data received can be delivered to the user. This is the normal state for the data transfer phase of the connection. • FINWAIT1—Waiting for a connection termination request from the remote TCP or an acknowledgment of the connection termination request previously sent.
Conn id	Identifying number of the TCP connection.

Feature Information for BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)

Table 17 lists the release history for this feature.

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Cisco IOS software images are specific to a Cisco IOS software release, a feature set, and a platform. Use Cisco Feature Navigator to find information about platform support and Cisco IOS software image support. Access Cisco Feature Navigator at http://www.cisco.com/go/fn. You must have an account on Cisco.com. If you do not have an account or have forgotten your username or password, click **Cancel** at the login dialog box and follow the instructions that appear.

Note Table 17 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

Table 17 Feature Information for BGP Support for NSR with SSO

Feature Name	Releases	Feature Information
BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO)	12.2(28)SB	The BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) enables PE routers to maintain BGP state with CE routers and ensure continuous packet forwarding during an RP switchover or during a planned ISSU for a PE router. CE routers do not need to be NSF-capable or NSF-aware to benefit from BGP NSR capabilities on PE routers. Only PE routers need to be upgraded to support BGP NSR—no CE router upgrades are required. BGP NSR with SSO, thus, enables service providers to provide the benefits NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart. In 12.2(28)SB, this feature was introduced on the Cisco 10000 series router. The following commands were introduced or modified by this feature: **debug ip bgp sso**, **debug ip tcp ha**, **neighbor ha-mode sso**, **show ip bgp vpnv4**, **show ip bgp vpnv4 all sso summary**, **show tcp**, **show tcp ha connections**.