neighbor timers

To set the timers for a specific BGP peer or peer group, use the **neighbor timers** command in address family or router configuration mode. To clear the timers for a specific BGP peer or peer group, use the **no** form of this command.

neighbor [*ip-address* | *peer-group-name*] **timers** *keepalive holdtime* [*min-holdtime*]

no neighbor [*ip-address* | *peer-group-name*] **timers**

Syntax Description

ip-address	(Optional) A BGP peer or peer group IP address.
peer-group-name	(Optional) Name of the BGP peer group.
keepalive	Frequency (in seconds) with which the Cisco IOS software sends *keepalive* messages to its peer. The default is 60 seconds.The range is from 0 to 65535.
holdtime	Interval (in seconds) after not receiving a *keepalive* message that the software declares a peer dead. The default is 180 seconds. The range is from 0 to 65535.
min-holdtime	(Optional) Interval (in seconds) specifying the minimum acceptable hold-time from a BGP neighbor. The minimum acceptable hold-time must be less than, or equal to, the interval specified in the *holdtime* argument. The range is from 0 to 65535.

Defaults

keepalive: 60 seconds
holdtime: 180 seconds

Command Modes

Address family configuration (config-router-af)
Router configuration (config-router)

Command History

Release	Modification
12.0	This command was introduced.
12.0(26)S	The *min-holdtime* argument was added.
12.3(7)T	The *min-holdtime* argument was added.
12.2(22)S	The *min-holdtime* argument was added.
12.2(27)SBC	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2(33)SRA	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	The *min-holdtime* argument was added and this command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

The timers configured for a specific neighbor or peer group override the timers configured for all BGP neighbors using the **timers bgp** command.

When configuring the *holdtime* argument for a value of less than twenty seconds, the following warning is displayed:

```
% Warning: A hold time of less than 20 seconds increases the chances of peer flapping
```

If the minimum acceptable hold-time interval is greater than the specified hold-time, a notification is displayed:

```
% Minimum acceptable hold time should be less than or equal to the configured hold time
```

Note When the minimum acceptable hold-time is configured on a BGP router, a remote BGP peer session is established only if the remote peer is advertising a hold-time that is equal to, or greater than, the minimum acceptable hold-time interval. If the minimum acceptable hold-time interval is greater than the configured hold-time, the next time the remote session tries to establish, it will fail and the local router will send a notification stating “unacceptable hold time.”

Examples

The following example changes the keepalive timer to 70 seconds and the hold-time timer to 210 seconds for the BGP peer 192.168.47.0:

```
router bgp 109
 neighbor 192.168.47.0 timers 70 210
```

The following example changes the keepalive timer to 70 seconds, the hold-time timer to 130 seconds, and the minimum hold-time interval to 100 seconds for the BGP peer 192.168.1.2:

```
router bgp 45000
 neighbor 192.168.1.2 timers 70 130 100
```

neighbor transport

To enable a TCP transport session option for a Border Gateway Protocol (BGP) session, use the **neighbor transport** command in router or address family configuration mode. To disable a TCP transport session option for a BGP session, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **transport** {**connection-mode** {**active** | **passive**} | **path-mtu-discovery** [**disable**] | **multi-session** | **single-session**}

no neighbor {*ip-address* | *peer-group-name*} **transport** {**connection-mode** | **path-mtu-discovery** | **multi-session** | **single-session**}

Syntax Description

ip-address	IP address of the BGP neighbor.
peer-group-name	Name of a BGP peer group.
connection-mode	Specifies the type of connection (active or passive).
active	Specifies an active connection.
passive	Specifies a passive connection.
path-mtu-discovery	Enables TCP transport path maximum transmission unit (MTU) discovery. TCP path MTU discovery is enabled by default.
multi-session	Enables a separate TCP transport session for each address family.
single-session	Enables all address families to use a single TCP transport session.
disable	Disables TCP path MTU discovery.

Command Default

If this command is not configured, TCP path MTU discovery is enabled by default, but no other TCP transport session options are enabled.

Command Modes

Router configuration (config-router)
Address family configuration (config-router-af)

Command History

Release	Modification
12.4	This command was introduced.
12.2(33)SRA	This command was modified. The **path-mtu-discovery** keyword was added.
12.2(33)SRB	This command was modified. The **multi-session**, **single-session**, and **disable** keywords were added.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
12.4(20)T	This command was modified. The **path-mtu-discovery** keyword was added.

Usage Guidelines

This command is used to specify various transport options. An active or passive transport connection can be specified for a BGP session. TCP transport path MTU discovery can be enabled to allow a BGP session to take advantage of larger MTU links. Use the **show ip bgp neighbors** command to determine whether TCP path MTU discovery is enabled.

In Cisco IOS Release 12.2(33)SRB and later releases, options can be specified for the transport of address family traffic using a single TCP session or to enable a separate TCP session for each address family. Multiple TCP sessions are used to support Multi-Topology Routing (MTR), and the single session option is available for backwards compatibility for non-MTR configurations and for scalability purposes.

In Cisco IOS Release 12.2(33)SRB and later releases, the ability to disable TCP path MTU discovery, for a single neighbor or for an inheriting peer or peer group, was added. If you use the **disable** keyword to disable discovery, discovery is also disabled on any peer or peer group that inherits the template in which you disabled discovery.

The following example shows how to configure the TCP transport connection to be active for a single internal BGP (iBGP) neighbor:

```
router bgp 45000
 neighbor 172.16.1.2 remote-as 45000
 neighbor 172.16.1.2 activate
 neighbor 172.16.1.2 transport connection-mode active
 end
```

The following example shows how to configure the TCP transport connection to be passive for a single external BGP (eBGP) neighbor:

```
router bgp 45000
 neighbor 192.168.1.2 remote-as 40000
 neighbor 192.168.1.2 activate
 neighbor 192.168.1.2 transport connection-mode passive
 end
```

The following example shows how to disable TCP path MTU discovery for a single BGP neighbor:

```
router bgp 45000
 neighbor 172.16.1.2 remote-as 45000
 neighbor 172.16.1.2 activate
 no neighbor 172.16.1.2 transport path-mtu-discovery
 end
```

The following example shows how to reenable TCP path MTU discovery for a single BGP neighbor, if TCP path MTU discovery is disabled:

```
router bgp 45000
 neighbor 172.16.1.2 remote-as 45000
 neighbor 172.16.1.2 activate
 neighbor 172.16.1.2 transport path-mtu-discovery
 end
```

The following example shows how to enable a separate TCP session for each address family for an MTR topology configuration:

```
router bgp 45000
 scope global
  neighbor 172.16.1.2 remote-as 45000
  neighbor 172.16.1.2 transport multi-session
  address-family ipv4
   topology VIDEO
   bgp tid 100
   neighbor 172.16.1.2 activate
   end
```

The following example shows how to disable TCP path MTU discovery and verify that it is disabled:

```
router bgp 100
 bgp log-neighbor-changes
 timers bgp 0 0
 redistribute static
 neighbor 10.4.4.4 remote-as 100
 neighbor 10.4.4.4 update-source Loopback 0
!end

Router# show ip bgp neighbors 10.4.4.4 | include path

    Used as bestpath:             n/a          0
    Used as multipath:            n/a          0
  Transport(tcp) path-mtu-discovery is enabled
Option Flags: nagle, path mtu capable
Router#

Router# configure terminal
Router(config)# router bgp 100

Router(config-router)# neighbors 10.4.4.4 transport path-mtu-discovery disable
Router(config-router)# end

Router# show ip bgp neighbor 10.4.4.4 | include path

    Used as bestpath:             n/a          0
    Used as multipath:            n/a          0
  Transport(tcp) path-mtu-discovery is disabled
```

Related Commands

Command	Description
bgp tid	Configures BGP to accept routes with a specified topology ID.
bgp transport	Enables transport session parameters globally for all BGP neighbor sessions.
scope	Defines the scope for a BGP routing session and enters router scope configuration mode.
show ip bgp neighbors	Displays information about BGP and TCP connections to neighbors.
topology (BGP)	Configures a process to route IP traffic under the specified topology instance.

neighbor ttl-security

To secure a Border Gateway Protocol (BGP) peering session and to configure the maximum number of hops that separate two external BGP (eBGP) peers, use the **neighbor ttl-security** command in address-family or router configuration mode. To disable this feature, use the **no** form of this command.

neighbor *neighbor-address* **ttl-security hops** *hop-count*

no neighbor *neighbor-address* **ttl-security hops** *hop-count*

Syntax Description

neighbor-address	IP address of the neighbor.
hops *hop-count*	Number of hops that separate the eBGP peers. The TTL value is calculated by the router from the configured *hop-count* argument. The value for the *hop-count* argument is a number between 1 and 254.

Defaults

No default behavior or values

Command Modes

Address-family configuration
Router configuration

Command History

Release	Modification
12.0(27)S	This command was introduced.
12.3(7)T	This command was integrated into Cisco IOS Release 12.3(7)T.
12.2(25)S	This command was integrated into Cisco IOS Release 12.2(25)S.
12.2(18)SXE	This command was integrated into Cisco IOS Release 12.3(7)T.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.

Usage Guidelines

The **neighbor ttl-security** command provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers.

This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network.

This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL

value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary.

To maximize the effectiveness of this feature, the *hop-count* value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session.

The following restrictions apply to the configuration of this command:

- This feature is not supported for internal BGP (iBGP) peers or iBGP peer groups.
- The **neighbor ttl-security** command cannot be configured for a peer that is already configured with the **neighbor ebgp-multihop** command. The configuration of these commands is mutually exclusive, and only one of these commands is needed to enable a multihop eBGP peering session. An error message will be displayed in the console if you attempt to configure both commands for the same peering session.
- The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack.
- This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.

Examples

The following example sets the hop count to 2 for a directly connected neighbor. Because the *hop-count* argument is set to 2, BGP will accept only IP packets with a TTL count in the header that is equal to or greater than 253. If a packet is received with any other TTL value in the IP packet header, the packet will be silently discarded.

```
neighbor 10.0.0.1 ttl-security hops 2
```

Related Commands

Command	Description
neighbor ebgp-multihop	Accepts or initiates BGP connections to external peers residing on networks that are not directly connected.
show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors.

neighbor unsuppress-map

To selectively advertise routes previously suppressed by the **aggregate-address** command, use the **neighbor unsuppress-map** command in address family or router configuration mode. To restore the system to the default condition, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **unsuppress-map** *route-map-name*

no neighbor {*ip-address* | *peer-group-name*} **unsuppress-map** *route-map-name*

Syntax Description

ip-address	IP address of the BGP-speaking neighbor.
peer-group-name	Name of a BGP peer group.
route-map-name	Name of a route map.

Command Default

No routes are unsuppressed.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
12.0(5)T	This command was introduced.
12.0(5)T	Address family configuration mode was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use of the **neighbor unsuppress-map** command allows specified suppressed routes to be advertised.

Examples

The following BGP router configuration shows that routes specified by a route map named map1 are suppressed:

```
access-list 3 deny 172.16.16.6
access-list 3 permit any
route-map map1 permit 10
match ip address 3
!
router bgp 65000
network 172.16.0.0
neighbor 192.168.1.2 remote-as 40000
aggregate-address 172.0.0.0 255.0.0.0 suppress-map map1
neighbor 192.168.1.2 unsuppress-map map1
neighbor 192.168.1.2 activate
```

The following example shows the routes specified by internal-map being unsuppressed for neighbor 172.16.16.6:

```
router bgp 100
address-family ipv4 multicast
 network 172.16.0.0
 neighbor 172.16.16.6 unsuppress-map internal-map
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IPv4 address prefixes.
address-family vpnv4	Places the routing in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPNv4 address prefixes.
aggregate-address	Creates an aggregate entry in a BGP routing table.
neighbor route-map	Applies a route map to inbound or outbound routes.

neighbor update-source

To have the Cisco IOS software allow Border Gateway Protocol (BGP) sessions to use any operational interface for TCP connections, use the **neighbor update-source** command in router configuration mode. To restore the interface assignment to the closest interface, which is called the best local address, use the **no** form of this command.

neighbor {*ip-address* | *ipv6-address*[**%**] | *peer-group-name*} **update-source** *interface-type interface-number*

no neighbor {*ip-address* | *ipv6-address*[**%**] | *peer-group-name*} **update-source** *interface-type interface-number*

Syntax Description

ip-address	IPv4 address of the BGP-speaking neighbor.
ipv6-address	IPv6 address of the BGP-speaking neighbor.
%	(Optional) IPv6 link-local address identifier. This keyword needs to be added whenever a link-local IPv6 address is used outside the context of its interface.
peer-group-name	Name of a BGP peer group.
interface-type	Interface type.
interface-number	Interface number.

Command Default

Best local address

Command Modes

Router configuration (config-router)

Command History

Release	Modification
10.0	This command was introduced.
12.2(4)T	The *ipv6-address* argument was added.
12.0(21)ST	This command was integrated into Cisco IOS Release 12.0(21)ST.
12.0(22)S	This command was integrated into Cisco IOS Release 12.0(22)S.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(25)SG	This command was integrated into Cisco IOS Release 12.2(25)SG.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	The **%** keyword was added.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
Cisco IOS XE Release 2.1	This command was introduced on Cisco ASR 1000 series routers.

Usage Guidelines

This command can work in conjunction with the loopback interface feature described in the "Interface Configuration Overview" chapter of the *Cisco IOS Interface and Hardware Component Configuration Guide*.

If you specify a BGP peer group by using the *peer-group-name* argument, all the members of the peer group will inherit the characteristic configured with this command.

The **neighbor update-source** command must be used to enable IPv6 link-local peering for internal or external BGP sessions.

The **%** keyword is used whenever link-local IPv6 addresses are used outside the context of their interfaces and for these link-local IPv6 addresses you must specify the interface they are on. The syntax becomes <IPv6 local-link address>%<interface name>, for example, FE80::1%Ethernet1/0. Note that the interface type and number must not contain any spaces, and be used in full-length form because name shortening is not supported in this situation. The % keyword and subsequent interface syntax is not used for non-link-local IPv6 addresses.

Examples

The following example sources BGP TCP connections for the specified neighbor with the IP address of the loopback interface rather than the best local address:

```
router bgp 65000
 network 172.16.0.0
 neighbor 172.16.2.3 remote-as 110
 neighbor 172.16.2.3 update-source Loopback0
```

The following example sources IPv6 BGP TCP connections for the specified neighbor in autonomous system 65000 with the global IPv6 address of loopback interface 0 and the specified neighbor in autonomous system 65400 with the link-local IPv6 address of Fast Ethernet interface 0/0. Note that the link-local IPv6 address of FE80::2 is on Ethernet interface 1/0.

```
router bgp 65000
 neighbor 3ffe::3 remote-as 65000
 neighbor 3ffe::3 update-source Loopback0
 neighbor fe80::2%Ethernet1/0 remote-as 65400
 neighbor fe80::2%Ethernet1/0 update-source FastEthernet 0/0
 address-family ipv6
  neighbor 3ffe::3 activate
  neighbor fe80::2%Ethernet1/0 activate
  exit-address-family
```

Related Commands

Command	Description
neighbor activate	Enables the exchange of information with a BGP neighboring router.
neighbor remote-as	Adds an entry to the BGP or multiprotocol BGP neighbor table.

neighbor version

To configure the Cisco IOS software to accept only a particular BGP version, use the **neighbor version** command in router configuration mode. To use the default version level of a neighbor, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **version** *number*

no neighbor {*ip-address* | *peer-group-name*} **version** *number*

Syntax Description

ip-address	IP address of the BGP-speaking neighbor.
peer-group-name	Name of a BGP peer group.
number	BGP version number. The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Defaults

BGP Version 4

Command Modes

Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Entering this command disables dynamic version negotiation.

Note The Cisco implementation of BGP in Cisco IOS Release 12.0(5)T or earlier releases supports BGP Versions 2, 3, and 4, with dynamic negotiation down to Version 2 if a neighbor does not accept BGP Version 4 (the default version).

The Cisco implementation of BGP in Cisco IOS Release 12.0(6)T or later releases supports BGP Version 4 only and does not support dynamic negotiation down to Version 2.

If you specify a BGP peer group by using the *peer-group-name* argument, all the members of the peer group will inherit the characteristic configured with this command.

Examples

The following example locks down to Version 4 of the BGP protocol:

```
router bgp 109
 neighbor 172.16.27.2 version 4
```

Related Commands

Command	Description
neighbor remote-as	Creates a BGP peer group.

neighbor weight

To assign a weight to a neighbor connection, use the **neighbor weight** command in address family or router configuration mode. To remove a weight assignment, use the **no** form of this command.

neighbor {*ip-address* | *peer-group-name*} **weight** *number*

no neighbor {*ip-address* | *peer-group-name*} **weight** *number*

Syntax Description

ip-address	IP address of the neighbor.
peer-group-name	Name of a BGP peer group.
number	Weight to assign. Acceptable values are from 0 to 65535.

Defaults

Routes learned through another BGP peer have a default weight of 0 and routes sourced by the local router have a default weight of 32768.

Command Modes

Address family
Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0(7)T	Address family configuration mode was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

All routes learned from this neighbor will have the assigned weight initially. The route with the highest weight will be chosen as the preferred route when multiple routes are available to a particular network.

The weights assigned with the **set weight** route-map command override the weights assigned using the **neighbor weight** command.

Note For weight changes to take effect, use of the **clear ip bgp peer-group *** command may be necessary.

If you specify a BGP peer group by using the *peer-group-name* argument, all the members of the peer group will inherit the characteristic configured with this command.

Examples

The following router configuration mode example sets the weight of all routes learned via 172.16.12.1 to 50:

```
router bgp 109
 neighbor 172.16.12.1 weight 50
```

The following address family configuration mode example sets the weight of all routes learned via 172.16.12.1 to 50:

```
router bgp 109
address-family ipv4 multicast
 neighbor 172.16.12.1 weight 50
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IP Version 4 address prefixes.
address-family vpnv4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard Virtual Private Network (VPN) Version 4 address prefixes.
neighbor distribute-list	Distributes BGP neighbor information as specified in an access list.
neighbor filter-list	Sets up a BGP filter.
neighbor remote-as	Creates a BGP peer group.

network (BGP and multiprotocol BGP)

To specify the networks to be advertised by the Border Gateway Protocol (BGP) and multiprotocol BGP routing processes, use the **network** command in address family or router configuration mode. To remove an entry from the routing table, use the **no** form of this command.

network {*network-number* [**mask** *network-mask*] | *nsap-prefix*} [**route-map** *map-tag*]

no network {*network-number* [**mask** *network-mask*] | *nsap-prefix*} [**route-map** *map-tag*]

Syntax Description

network-number	Network that BGP or multiprotocol BGP will advertise.
mask *network-mask*	(Optional) Network or subnetwork mask with mask address.
nsap-prefix	Network service access point (NSAP) prefix of the Connectionless Network Service (CLNS) network that BGP or multiprotocol BGP will advertise. This argument is used only under NSAP address family configuration mode.
route-map *map-tag*	(Optional) Identifier of a configured route map. The route map should be examined to filter the networks to be advertised. If not specified, all networks are advertised. If the keyword is specified, but no route map tags are listed, no networks will be advertised.

Command Default

No networks are specified.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0	The limit of 200 network commands per BGP router was removed.
11.1(20)CC	The **nlri unicast**, **nlri multicast**, and **nlri unicast multicast** keywords were added.
12.0(7)T	The **nlri unicast**, **nlri multicast**, and **nlri unicast multicast** keywords were removed. Address family configuration mode was added.
12.2(8)T	The *nsap-prefix* argument was added to address family configuration mode.
12.2(25)SG	This command was integrated into Cisco IOS Release 12.2(25)SG.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

BGP and multiprotocol BGP networks can be learned from connected routes, from dynamic routing, and from static route sources.

The maximum number of **network** commands you can use is determined by the resources of the router, such as the configured NVRAM or RAM.

Examples

The following example sets up network 10.108.0.0 to be included in the BGP updates:

```
router bgp 65100
 network 10.108.0.0
```

The following example sets up network 10.108.0.0 to be included in the multiprotocol BGP updates:

```
router bgp 64800
address family ipv4 multicast
 network 10.108.0.0
```

The following example advertises NSAP prefix 49.6001 in the multiprotocol BGP updates:

```
router bgp 64500
 address-family nsap
  network 49.6001
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Enters the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IP Version 4 address prefixes.
address-family vpnv4	Enters the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPNv4 address prefixes.
default-information originate (BGP)	Allows the redistribution of network 0.0.0.0 into BGP.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another.
router bgp	Configures the BGP routing process.

network backdoor

To specify a backdoor route to a BGP-learned prefix that provides better information about the network, use the **network backdoor** command in address family or router configuration mode. To remove an address from the list, use the **no** form of this command.

network *ip-address* **backdoor**

no network *ip-address* **backdoor**

Syntax Description

ip-address	IP address of the network to which you want a backdoor route.

Defaults

No network is marked as having a back door.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
10.0	This command was introduced.
12.0(7)T	Address family configuration mode was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

A backdoor network is assigned an administrative distance of 200. The objective is to make Interior Gateway Protocol (IGP) learned routes preferred. A backdoor network is treated as a local network, except that it is not advertised. A network that is marked as a back door is not sourced by the local router, but should be learned from external neighbors. The BGP best path selection algorithm does not change when a network is configured as a back door.

Examples

The following address family configuration example configures network 10.108.0.0 as a local network and network 192.168.7.0 as a backdoor network:

```
router bgp 109
address-family ipv4 multicast
 network 10.108.0.0
 network 192.168.7.0 backdoor
```

The following router configuration example configures network 10.108.0.0 as a local network and network 192.168.7.0 as a backdoor network:

```
router bgp 109
 network 10.108.0.0
 network 192.168.7.0 backdoor
```

Related Commands

Command	Description
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IP Version 4 address prefixes.
address-family vpnv4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPN Version 4 address prefixes.
distance bgp	Allows the use of external, internal, and local administrative distances that could be a better route to a node.
network (BGP and multiprotocol BGP)	Specifies networks to be advertised by the BGP and multiprotocol BGP routing processes.
router bgp	Assigns an absolute weight to a BGP network.

prefix-length-size

To specify the length (in bytes) of the prefix length field of prefixes being advertised to neighbors, use the **prefix-length-size** command in L2VPN VPLS address-family configuration mode. To restore the default value, use the **no** form of this command.

prefix-length-size {1|2}

no prefix-length-size

Syntax Description

1|2	Specifies the length in bytes of the prefix length field (either 1 byte or 2 bytes).

Command Default

1 byte

Command Modes

L2VPN VPLS address-family configuration (config-router-af)

Command History

Release	Modification
12.2(33)SRD	This command was introduced.

Usage Guidelines

You might need to configure this command for interoperability with Juniper's JunOS. If the neighbor is a Juniper JunOS router, change the prefix length size to 2 bytes.

The size of the prefix length field is either 1 or 2 bits or bytes, depending on the address family of the prefix, as follows:

Address Family	Prefix Length Measured In	Prefix Length
All other address families	Bits	1
L2VPN AF (old Cisco)	Bits	1
L2VPN AF (AB76)	Bits (default, no prefix-length-size)	1
L2VPN AF (AB76)	Bytes (with prefix-length-size 2)	2
L2VPN AF (JUNOS)	Bytes	2

Examples

The following example configures the prefix length size to 2 bytes for L2VPN VPLS prefixes advertised to neighbors:

```
router bgp 1600
 address-family l2vpn vpls
 prefix-length-size 2
 neighbor 100.16.11.10 activate
 exit-address-family
```

Related Commands

Command	Description
address-family l2vpn vpls	Enters address family configuration mode to configure a routing session using Layer 2 Virtual Private Network (L2VPN) endpoint provisioning address information.
neighbor prefix-length-size	Configures the size of the prefix-length field for prefixes advertised to a neighbor.

redistribute (BGP to ISO IS-IS)

To redistribute routes from a Border Gateway Protocol (BGP) autonomous system into an International Organization for Standardization (ISO) Intermediate System-to-Intermediate System (IS-IS) routing process, use the **redistribute** command in router configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition where the software does not redistribute routes, use the **no** form of this command.

redistribute *protocol autonomous-system-number* [*route-type*] [**route-map** *map-tag*]

no redistribute *protocol autonomous-system-number* [*route-type*] [**route-map** *map-tag*]

Syntax Description

protocol	Source protocol from which routes are being redistributed. It must be the **bgp** keyword. The **bgp** keyword is used to redistribute dynamic routes.
autonomous-system-number	The autonomous system number of the BGP routing process. The range of values for this argument is any valid autonomous system number from 1 to 65535. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command.
route-type	(Optional) The type of route to be redistributed. It can be one of the following keywords: **clns** or **ip**. The default is **ip.** The **clns** keyword is used to redistribute BGP routes with network service access point (NSAP) addresses into IS-IS. The **ip** keyword is used to redistribute BGP routes with IP addresses into IS-IS.
route-map *map-tag*	(Optional) Identifier of a configured route map. The route map should be examined to filter the importation of routes from this source routing protocol to IS-IS. If not specified, all routes are redistributed. If the keyword is specified, but no route map tags are listed, no routes will be imported.

Command Default

Route redistribution is disabled.

Command Modes

Router configuration (config-router)

Command History

Release	Modification
12.2(8)T	This command was modified. The **clns** keyword was added.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
12.0(32)S12	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added. Support for changing autonomous system number of the BGP routing process was removed.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

The **clns** keyword must be specified to redistribute NSAP prefix routes from BGP into an ISO IS-IS routing process. This version of the **redistribute** command is used only under router configuration mode for IS-IS processes.

In redistribution from IGP (for example, ISIS, OSPF, RIP, or EIGRP) to BGP, the support for changing the autonomous system numbers of BGP from one to another is removed.

Examples

The following example configures NSAP prefix routes from BGP autonomous system 64500 to be redistributed into the IS-IS routing process called osi-proc-17:

```
router isis osi-proc-17
 redistribute bgp 64500 clns
```

In the following example the autonomous system BGP is modified from 200 to 300, this is not supported.

```
Router#config terminal
Router(config-if)#router eigrp 101
Router(config-router)#redistribute bgp 200
Router(config-router)#redistribute bgp 300
Cannot configure or redistribute to BGP AS 300
Please do "no router bgp 200" first
```

Remove support for autonomous system number 200 before configuring number 300.

```
Router(config)#no router bgp 200
Router(config-router)#redistribute bgp 300
```

Related Commands

Command	Description
network (BGP and multiprotocol BGP)	Specifies the list of networks for the BGP routing process.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another.
router bgp	Configures the BGP routing process.
show route-map	Displays all route maps configured or only the one specified.

redistribute (IP)

To redistribute routes from one routing domain into another routing domain, use the **redistribute** command in the appropriate configuration mode. To disable redistribution, use the **no** form of this command.

redistribute *protocol* [*process-id*] {**level-1** | **level-1-2** | **level-2**} [*autonomous-system-number*] [**metric** {*metric-value* | **transparent**}] [**metric-type** *type-value*] [**match** {**internal** | **external 1** | **external 2**}] [**tag** *tag-value*] [**route-map** *map-tag*] [**subnets**] [**nssa-only**]

no redistribute *protocol* [*process-id*] {**level-1** | **level-1-2** | **level-2**} [*autonomous-system-number*] [**metric** {*metric-value* | **transparent**}] [**metric-type** *type-value*] [**match** {**internal** | **external 1** | **external 2**}] [**tag** *tag-value*] [**route-map** *map-tag*] [**subnets**] [**nssa-only**]

Syntax Description

protocol	Source protocol from which routes are being redistributed. It can be one of the following keywords: **bgp**, **connected**, **eigrp**, **isis**, **mobile**, **ospf**, **static** [**ip**], or **rip**. The **static** [**ip**] keyword is used to redistribute IP static routes. The optional **ip** keyword is used when redistributing into the Intermediate System-to-Intermediate System (IS-IS) protocol. The **connected** keyword refers to routes that are established automatically by virtue of having enabled IP on an interface. For routing protocols such as Open Shortest Path First (OSPF) and IS-IS, these routes will be redistributed as external to the autonomous system.
process-id	(Optional) For the **bgp** or **eigrp** keyword, this is an autonomous system number, which is a 16-bit decimal number. For the **isis** keyword, this is an optional *tag* value that defines a meaningful name for a routing process. You can specify only one IS-IS process per router. Creating a name for a routing process means that you use names when configuring routing. For the **ospf** keyword, this is an appropriate OSPF process ID from which routes are to be redistributed. This identifies the routing process. This value takes the form of a nonzero decimal number. For the **rip** keyword, no *process-id* value is needed. By default, no process ID is defined.
level-1	Specifies that, for IS-IS, Level 1 routes are redistributed into other IP routing protocols independently.
level-1-2	Specifies that, for IS-IS, both Level 1 and Level 2 routes are redistributed into other IP routing protocols.
level-2	Specifies that, for IS-IS, Level 2 routes are redistributed into other IP routing protocols independently.

autonomous-system-number	(Optional) Autonomous system number for the redistributed route. Number in the range from 1 to 65535. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command.
metric *metric-value*	(Optional) When redistributing from one OSPF process to another OSPF process on the same router, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified. The default value is 0.
metric transparent	(Optional) Causes RIP to use the routing table metric for redistributed routes as the RIP metric.
metric-type *type-value*	(Optional) For OSPF, specifies the external link type associated with the default route advertised into the OSPF routing domain. It can be one of two values: • **1**—Type 1 external route • **2**—Type 2 external route If a **metric-type** is not specified, the Cisco IOS software adopts a Type 2 external route. For IS-IS, it can be one of two values: • **internal**—IS-IS metric that is < 63. • **external**—IS-IS metric that is > 64 < 128. The default is **internal**.
match {**internal** | **external 1** | **external 2**}	(Optional) For the criteria by which OSPF routes are redistributed into other routing domains. It can be one of the following: • **internal**—Routes that are internal to a specific autonomous system. • **external 1**—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external route. • **external 2**—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external route. The default is internal and external 1.

tag *tag-value*	(Optional) Specifies the 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between Autonomous System Boundary Routers (ASBRs). If none is specified, then the remote autonomous system number is used for routes from Border Gateway Protocol (BGP) and Exterior Gateway Protocol (EGP); for other protocols, zero (0) is used.
route-map	(Optional) Specifies the route map that should be interrogated to filter the importation of routes from this source routing protocol to the current routing protocol. If not specified, all routes are redistributed. If this keyword is specified, but no route map tags are listed, no routes will be imported.
map-tag	(Optional) Identifier of a configured route map.
subnets	(Optional) For redistributing routes into OSPF, the scope of redistribution for the specified protocol. By default, no subnets are defined.
nssa-only	(Optional) Sets the nssa-only attribute for all routes redistributed into OSPF.

Command Default

Route redistribution is disabled.

Command Modes

Router configuration (config-router)
Address family configuration (config-af)
Address family topology configuration (config-router-af-topology)

Command History

Release	Modification
10.0	This command was introduced.
12.0(5)T	This command was modified. Address family configuration mode was added.
12.0(22)S	This command was modified. Address family support under EIGRP was added.
12.2(15)T	This command was modified. Address family support under EIGRP was added.
12.2(18)S	This command was modified. Address family support under EIGRP was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	This command was modified. This command was made available in router address family topology configuration mode.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.0(32)S12	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Release	Modification
12.4(24)T	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is asplain.
Cisco IOS Release 15.0(1)M	This command was modified. The **nssa-only** keyword was added.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

Changing or disabling any keyword will not affect the state of other keywords.

A router receiving a link-state protocol with an internal metric will consider the cost of the route from itself to the redistributing router plus the advertised cost to reach the destination. An external metric only considers the advertised metric to reach the destination.

Routes learned from IP routing protocols can be redistributed at Level 1 into an attached area or at Level 2. The **level-1-2** keyword allows both Level 1 and Level 2 routes in a single command.

Redistributed routing information must be filtered by the **distribute-list out** router configuration command. This guideline ensures that only those routes intended by the administrator are passed along to the receiving routing protocol.

Whenever you use the **redistribute** or the **default-information** router configuration commands to redistribute routes into an OSPF routing domain, the router automatically becomes an ASBR. However, an ASBR does not, by default, generate a *default route* into the OSPF routing domain.

When routes are redistributed into OSPF from protocols other than OSPF or BGP, and no metric has been specified with the **metric-type** keyword and *type-value* argument, OSPF will use 20 as the default metric. When routes are redistributed into OSPF from BGP, OSPF will use 1 as the default metric. When routes are redistributed from one OSPF process to another OSPF process, Autonomous system (AS) external and not-so-stubby-area (NSSA) routes will use 20 as the default metric. When intra-area and inter-area routes are redistributed between OSPF processes, the internal OSPF metric from the redistribution source process is advertised as the external metric in the redistribution destination process. (This is the only case in which the routing table metric will be preserved when routes are redistributed into OSPF.)

When routes are redistributed into OSPF, only routes that are not subnetted are redistributed if the **subnets** keyword is not specified.

On a router internal to an NSSA area, the **nssa-only** keyword causes the originated type-7 NSSA LSAs to have their propagate (P) bit set to zero, which prevents area border routers from translating these LSAs into type-5 external LSAs. On an area border router that is connected to a NSSA and normal areas, the **nssa-only** keyword causes the routes to be redistributed only into the NSSA areas.

Routes configured with the **connected** keyword affected by this **redistribute** command are the routes not specified by the **network** router configuration command.

You cannot use the **default-metric** command to affect the metric used to advertise connected routes.

Note The **metric** value specified in the **redistribute** command supersedes the **metric** value specified using the **default-metric** command.

Default redistribution of IGPs or EGP into BGP is not allowed unless the **default-information originate** router configuration command is specified.

Using the no Form of the redistribute Command

Removing options that you have configured for the **redistribute** command requires careful use of the **no** form of the **redistribute** command to ensure that you obtain the result that you are expecting. See the "Examples" section for more information.

Release 12.2(33)SRB

If you plan to configure the Multi-Topology Routing (MTR) feature, you need to enter the **redistribute** command in router address family topology configuration mode in order for this OSPF router configuration command to become topology-aware.

4-Byte Autonomous System Number Support

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2, for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Examples

The following example shows how OSPF routes are redistributed into a BGP domain:

```
Router(config)# router bgp 109
Router(config-router)# redistribute ospf
```

The following example causes EIGRP routes to be redistributed into an OSPF domain:

```
Router(config)# router ospf 110
Router(config-router)# redistribute eigrp
```

The following example causes the specified EIGRP process routes to be redistributed into an OSPF domain. The EIGRP-derived metric will be remapped to 100 and RIP routes to 200.

```
Router(config)# router ospf 109
Router(config-router)# redistribute eigrp 108 metric 100 subnets
Router(config-router)# redistribute rip metric 200 subnets
```

The following example configures BGP routes to be redistributed into IS-IS. The link-state cost is specified as 5, and the metric type will be set to external, indicating that it has lower priority than internal metrics.

```
Router(config)# router isis
Router(config-router)# redistribute bgp 120 metric 5 metric-type external
```

In the following example, network 172.16.0.0 will appear as an external link-state advertisement (LSA) in OSPF 1 with a cost of 100 (the cost is preserved):

```
Router(config)# interface ethernet 0
Router(config-if)# ip address 172.16.0.1 255.0.0.0
Router(config)# ip ospf cost 100
Router(config)# interface ethernet 1
Router(config-if)# ip address 10.0.0.1 255.0.0.0
!
Router(config)# router ospf 1
Router(config-router)# network 10.0.0.0 0.255.255.255 area 0
Router(config-router)# redistribute ospf 2 subnet
Router(config)# router ospf 2
Router(config-router)# network 172.16.0.0 0.255.255.255 area 0
```

The following example shows how BGP routes are redistributed into OSPF and assigned the local 4-byte autonomous system number in asplain format. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
Router(config)# router ospf 2
Router(config-router)# redistribute bgp 65538
```

The following example removes the **connected metric 1000 subnets** options from the **redistribute connected metric 1000 subnets** command and leaves the **redistribute connected** command in the configuration:

```
Router(config-router)# no redistribute connected metric 1000 subnets
```

The following example removes the **metric 1000** options from the **redistribute connected metric 1000 subnets** command and leaves the **redistribute connected subnets** command in the configuration:

```
Router(config-router)# no redistribute connected metric 1000
```

The following example removes the **subnets** options from the **redistribute connected metric 1000 subnets** command and leaves the **redistribute connected metric 1000** command in the configuration:

```
Router(config-router)# no redistribute connected subnets
```

The following example removes the **redistribute connected** command, and any of the options that were configured for the **redistribute connected** command, from the configuration:

```
Router(config-router)# no redistribute connected
```

The following example shows how EIGRP routes are redistributed into an EIGRP process in a named EIGRP configuration:

```
Router(config)# router eigrp virtual-name
Router(config-router)# address-family ipv4 autonomous-system 1
Router(config-router-af)# topology base
Router(config-router-af-topology)# redistribute eigrp 6473 metric 1 1 1 1 1
```

Related Commands

Command	Description
address-family (EIGRP)	Enters address-family configuration mode to configure an EIGRP routing instance.
address-family ipv4 (BGP)	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard IPv4 address prefixes.

Command	Description
address-family vpnv4	Places the router in address family configuration mode for configuring routing sessions such as BGP, RIP, or static routing sessions that use standard VPNv4 address prefixes.
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
default-information originate (BGP)	Allows the redistribution of network 0.0.0.0 into BGP.
default-information originate (IS-IS)	Generates a default route into an IS-IS routing domain.
default-information originate (OSPF)	Generates a default route into an OSPF routing domain.
distribute-list out (IP)	Suppresses networks from being advertised in updates.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
router bgp	Configures the BGP routing process.
router eigrp	Configures the EIGRP address-family process.
show route-map	Displays all route maps configured or only the one specified.
topology (EIGRP)	Configures an EIGRP process to route IP traffic under the specified topology instance and enters address family topology configuration mode.

redistribute (ISO IS-IS to BGP)

To redistribute routes from an International Organization for Standardization (ISO) Intermediate System-to-Intermediate System (IS-IS) routing process into a Border Gateway Protocol (BGP) autonomous system, use the **redistribute** command in address family or router configuration mode. To remove the **redistribute** command from the configuration file and restore the system to its default condition where the software does not redistribute routes, use the **no** form of this command.

redistribute *protocol* [*process-id*] [*route-type*] [**route-map** [*map-tag*]]

no redistribute *protocol* [*process-id*] [*route-type*] [**route-map** [*map-tag*]]

Syntax Description

protocol	Source protocol from which routes are being redistributed. It can be one of the following keywords: **isis** or **static**. The **isis** keyword is used to redistribute dynamic routes. The **static** keyword is used to redistribute static routes.
process-id	(Optional) When IS-IS is used as a source protocol, this argument defines a meaningful name for a routing process. The *process-id* argument identifies from which IS-IS routing process routes will be redistributed. Routes can be redistributed only from IS-IS routing processes that involve Level 2 routes, including IS-IS Level 1-2 and Level 2 routing processes. The *process-id* argument is not used when the protocol keyword is **static**.
route-type	(Optional) The type of route to be redistributed. It can be one of the following keywords: **clns** or **ip**. The default is **ip**. The **clns** keyword is used to redistribute Connectionless Network Service (CLNS) routes with network service access point (NSAP) addresses into BGP. The **ip** keyword is used to redistribute IS-IS routes with IP addresses into BGP.
route-map *map-tag*	(Optional) Identifier of a configured route map. The route map should be examined to filter the importation of routes from this source routing protocol to BGP. If no route map is specified, all routes are redistributed. If the **route-map** keyword is specified, but no *map-tag* value is entered, no routes will be imported.

Command Default

Route redistribution is disabled.

route-type: **ip**

route-map *map-tag*: If the **route-map** argument is not entered, all routes are redistributed; if no *map-tag* value is entered, no routes are imported.

Command Modes

Address family configuration (Cisco IOS 12.3(8)T and later releases)
Router configuration (T-releases after Cisco IOS 12.3(8)T)

Command History

Release	Modification
12.2(8)T	The **clns** keyword was added.
12.3(8)T	Beginning with Cisco IOS Release 12.3(8)T this version of the **redistribute** command should be entered under address family mode rather than router configuration mode.
12.2(33)SRB	This command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **clns** keyword must be specified to redistribute NSAP prefix routes from an ISO IS-IS routing process into BGP. Beginning with Cisco IOS Release 12.3(8)T, this version of the **redistribute** command is entered only in address family configuration mode for BGP processes.

Examples

Cisco IOS Releases Prior to Release 12.3(8)T

The following example configures CLNS NSAP routes from the IS-IS routing process called osi-proc-6 to be redistributed into BGP:

```
Router(config)# router bgp 64352
Router(config-router)# redistribute isis osi-proc-6 clns
```

Cisco IOS Releases 12.3(8)T and Later Releases

The following example configures CLNS NSAP routes from the IS-IS routing process called osi-proc-15 to be redistributed into BGP:

```
Router(config)# router bgp 404
Router(config-router)# address-family nsap
Router(config-router-af)# redistribute isis osi-proc-15 clns
```

Related Commands

Command	Description
network (BGP and multiprotocol BGP)	Specifies the list of networks for the BGP routing process.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another.
show route-map	Displays all route maps configured or only the one specified.

redistribute dvmrp

To configure redistribution of Distance Vector Multicast Routing Protocol (DVMRP) routes into multiprotocol BGP, use the **redistribute dvmrp** command in address family or router configuration mode. To stop such redistribution, use the **no** form of this command.

redistribute dvmrp [**route-map** *map-name*]

no redistribute dvmrp [**route-map** *map-name*]

Syntax Description

route-map *map-name*	(Optional) Name of the route map that contains various BGP attribute settings.

Defaults

DVMRP routes are not redistributed into multiprotocol BGP.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
11.1(20)CC	This command was introduced.
12.0(7)T	Address family configuration mode was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use this command if you have a subset of DVMRP routes in an autonomous system that you want to take the multiprotocol BGP path. Define a route map to further specify which DVMRP routes get redistributed.

Examples

The following router configuration mode example redistributes DVMRP routes to BGP peers that match access list 1:

```
router bgp 109
 redistribute dvmrp route-map dvmrp-into-mbgp
route-map dvmrp-into-mbgp
 match ip address 1
```

The following address family configuration mode example redistributes DVMRP routes to multiprotocol BGP peers that match access list 1:

```
router bgp 109
address-family ipv4 multicast
 redistribute dvmrp route-map dvmrp-into-mbgp

route-map dvmrp-into-mbgp
 match ip address 1
```

router bgp

To configure the Border Gateway Protocol (BGP) routing process, use the **router bgp** command in global configuration mode. To remove a BGP routing process, use the **no** form of this command.

router bgp *autonomous-system-number*

no router bgp *autonomous-system-number*

Syntax Description

autonomous-system-number	Number of an autonomous system that identifies the router to other BGP routers and tags the routing information that is passed along. Number in the range from 1 to 65535. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the "Usage Guidelines" section.

Command Default

No BGP routing process is enabled by default.

Command Modes

Global configuration (config)

Command History

Release	Modification
10.0	This command was introduced.
12.2(25)SG	This command was integrated into Cisco IOS Release 12.2(25)SG.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(31)SB2	This command was integrated into Cisco IOS Release 12.2(31)SB2.
12.2(33)SRB	This command was modified. Support for IPv6 was added.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.2(33)SB	This command was modified. Support for IPv6 was added.
12.0(32)S12	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.

Release	Modification
Cisco IOS XE Release 2.3	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

This command allows you to set up a distributed routing core that automatically guarantees the loop-free exchange of routing information between autonomous systems.

Prior to January 2009, BGP autonomous system numbers that were allocated to companies were 2-octet numbers in the range from 1 to 65535 as described in RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*. Due to increased demand for autonomous system numbers, the Internet Assigned Number Authority (IANA) will start in January 2009 to allocate four-octet autonomous system numbers in the range from 65536 to 4294967295. RFC 5396, *Textual Representation of Autonomous System (AS) Numbers*, documents three methods of representing autonomous system numbers. Cisco has implemented the following two methods:

- Asplain—Decimal value notation where both 2-byte and 4-byte autonomous system numbers are represented by their decimal value. For example, 65526 is a 2-byte autonomous system number and 234567 is a 4-byte autonomous system number.
- Asdot—Autonomous system dot notation where 2-byte autonomous system numbers are represented by their decimal value and 4-byte autonomous system numbers are represented by a dot notation. For example, 65526 is a 2-byte autonomous system number and 1.169031 is a 4-byte autonomous system number (this is dot notation for the 234567 decimal number).

For details about the third method of representing autonomous system numbers, see RFC 5396.

Note In Cisco IOS releases that include 4-byte ASN support, command accounting and command authorization that include a 4-byte ASN number are sent in the asplain notation irrespective of the format that is used on the command-line interface.

Asdot Only Autonomous System Number Formatting

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, Cisco IOS XE Release 2.3, and later releases, the 4-octet (4-byte) autonomous system numbers are entered and displayed only in asdot notation, for example, 1.10 or 45000.64000. When using regular expressions to match 4-byte autonomous system numbers the asdot format includes a period which is a special character in regular expressions. A backslash must be entered before the period for example, 1\.14, to ensure the regular expression match does not fail. Table 6 shows

the format in which 2-byte and 4-byte autonomous system numbers are configured, matched in regular expressions, and displayed in **show** command output in Cisco IOS images where only asdot formatting is available.

Table 6 Asdot Only 4-Byte Autonomous System Number Format

Format	Configuration Format	Show Command Output and Regular Expression Match Format
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535

Asplain as Default Autonomous System Number Formatting

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain as the default display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain and asdot format. In addition, the default format for matching 4-byte autonomous system numbers in regular expressions is asplain, so you must ensure that any regular expressions to match 4-byte autonomous system numbers are written in the asplain format. If you want to change the default **show** command output to display 4-byte autonomous system numbers in the asdot format, use the **bgp asnotation dot** command under router configuration mode. When the asdot format is enabled as the default, any regular expressions to match 4-byte autonomous system numbers must be written using the asdot format, or the regular expression match will fail. Table 7 and Table 8 show that although you can configure 4-byte autonomous system numbers in either asplain or asdot format, only one format is used to display **show** command output and control 4-byte autonomous system number matching for regular expressions, and the default is asplain format. To display 4-byte autonomous system numbers in **show** command output and to control matching for regular expressions in the asdot format, you must configure the **bgp asnotation dot** command. After enabling the **bgp asnotation dot** command, a hard reset must be initiated for all BGP sessions by entering the **clear ip bgp *** command.

Note If you are upgrading to an image that supports 4-byte autonomous system numbers, you can still use 2-byte autonomous system numbers. The **show** command output and regular expression match are not changed and remain in asplain (decimal value) format for 2-byte autonomous system numbers regardless of the format configured for 4-byte autonomous system numbers.

Table 7 Default Asplain 4-Byte Autonomous System Number Format

Format	Configuration Format	Show Command Output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 65536 to 4294967295
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 65536 to 4294967295

Table 8 Asdot 4-Byte Autonomous System Number Format

Format	Configuration Format	Show Command Output and Regular Expression Match Format
asplain	2-byte: 1 to 65535 4-byte: 65536 to 4294967295	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535
asdot	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535	2-byte: 1 to 65535 4-byte: 1.0 to 65535.65535

Reserved and Private Autonomous System Numbers

In Cisco IOS Release 12.0(32)S12, 12.0(32)SY8, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, 12.4(24)T, Cisco IOS XE Release 2.3 and later releases, the Cisco implementation of BGP supports RFC 4893. RFC 4893 was developed to allow BGP to support a gradual transition from 2-byte autonomous system numbers to 4-byte autonomous system numbers. A new reserved (private) autonomous system number, 23456, was created by RFC 4893 and this number cannot be configured as an autonomous system number in the Cisco IOS CLI.

RFC 5398, *Autonomous System (AS) Number Reservation for Documentation Use*, describes new reserved autonomous system numbers for documentation purposes. Use of the reserved numbers allow configuration examples to be accurately documented and avoids conflict with production networks if these configurations are literally copied. The reserved numbers are documented in the IANA autonomous system number registry. Reserved 2-byte autonomous system numbers are in the contiguous block, 64496 to 64511 and reserved 4-byte autonomous system numbers are from 65536 to 65551 inclusive.

Private 2-byte autonomous system numbers are still valid in the range from 64512 to 65534 with 65535 being reserved for special use. Private autonomous system numbers can be used for internal routing domains but must be translated for traffic that is routed out to the Internet. BGP should not be configured to advertise private autonomous system numbers to external networks. Cisco IOS software does not remove private autonomous system numbers from routing updates by default. We recommend that ISPs filter private autonomous system numbers.

Note Autonomous system number assignment for public and private networks is governed by the IANA. For information about autonomous-system numbers, including reserved number assignment, or to apply to register an autonomous system number, see the following URL: http://www.iana.org/.

Examples

The following example configures a BGP process for autonomous system 45000 and configures two external BGP neighbors in different autonomous systems using 2-byte autonomous system numbers:

```
router bgp 45000
 neighbor 192.168.1.2 remote-as 40000
 neighbor 192.168.3.2 remote-as 50000
 neighbor 192.168.3.2 description finance
 !
 address-family ipv4
  neighbor 192.168.1.2 activate
  neighbor 192.168.3.2 activate
  no auto-summary
  no synchronization
  network 172.17.1.0 mask 255.255.255.0
  exit-address-family
```

The following example configures a BGP process for autonomous system 65538 and configures two external BGP neighbors in different autonomous systems using 4-byte autonomous system numbers in asplain notation. This example is supported in Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases.

```
router bgp 65538
 neighbor 192.168.1.2 remote-as 65536
 neighbor 192.168.3.2 remote-as 65550
 neighbor 192.168.3.2 description finance
 !
 address-family ipv4
  neighbor 192.168.1.2 activate
  neighbor 192.168.3.2 activate
  no auto-summary
  no synchronization
  network 172.17.1.0 mask 255.255.255.0
  exit-address-family
```

The following example configures a BGP process for autonomous system 1.2 and configures two external BGP neighbors in different autonomous systems using 4-byte autonomous system numbers in asdot notation. This example is supported in Cisco IOS Release 12.0(32)SY8, 12.0(32)S12, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, 12.4(24)T, and Cisco IOS XE Release 2.3, and later releases.

```
router bgp 1.2
 neighbor 192.168.1.2 remote-as 1.0
 neighbor 192.168.3.2 remote-as 1.14
 neighbor 192.168.3.2 description finance
 !
 address-family ipv4
  neighbor 192.168.1.2 activate
  neighbor 192.168.3.2 activate
  no auto-summary
  no synchronization
  network 172.17.1.0 mask 255.255.255.0
  exit-address-family
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
neighbor remote-as	Adds an entry to the BGP or multiprotocol BGP neighbor table.
network (BGP and multiprotocol BGP)	Specifies the list of networks for the BGP routing process.

route-server-context

To create a route-server context in order to provide flexible policy handling for a BGP route server, use the **route-server-context** command in router configuration mode. To remove the route server context, use the **no** form of this command.

route-server-context *context-name*

no route-server-context *context-name*

Syntax Description

context-name	Name of the route server context.

Command Default

No route server context exists.

Command Modes

Router configuration (config-router)

Command History

Release	Modification
Cisco IOS XE 3.3S	This command was introduced.

Usage Guidelines

Flexible (customized) policy support for a BGP route server is made possible with the use of the **route-server-context** command. The **route-server-context** command creates a context, which represents the virtual table used to store prefixes and paths that require special handling due to individualized policy configurations.

The context is referenced by the BGP neighbors assigned to use that context (in the **neighbor route-server-client** command). Thus, multiple neighbors sharing the same policy can share the same route server context.

In order to configure flexible policy handling, create a route server context, which includes an import map. The import map references a standard route map.

Examples

In the following example, the local router is a BGP route server. Its neighbors at 10.10.10.12 and 10.10.10.13 are its route server clients. A route server context named ONLY_AS27_CONTEXT is created and applied to the neighbor at 10.10.10.13. The context uses an import map that references a route map named only_AS27_routemap. The route map matches routes permitted by access list 27. Access list 27 permits routes that have 27 in the autonomous system path.

```
router bgp 65000
   route-server-context ONLY_AS27_CONTEXT
      address-family ipv4 unicast
         import-map only_AS27_routemap
      exit-address-family
   exit-route-server-context
   !
   neighbor 10.10.10.12 remote-as 12
   neighbor 10.10.10.12 description Peer12
```

```
   neighbor 10.10.10.13 remote-as 13
   neighbor 10.10.10.13 description Peer13
   neighbor 10.10.10.21 remote-as 21
   neighbor 10.10.10.27 remote-as 27
   !
   address-family ipv4
      neighbor 10.10.10.12 activate
      neighbor 10.10.10.12 route-server-client
      neighbor 10.10.10.13 activate
      neighbor 10.10.10.13 route-server-client context ONLY_AS27_CONTEXT
      neighbor 10.10.10.21 activate
      neighbor 10.10.10.27 activate
   exit-address-family
!
ip as-path access-list 27 permit 27
!
route-map only_AS27_routemap permit 10
   match as-path 27
!
```

Related Commands

Command	Description
description (route-server-context)	Specifies a description for a route-server-context.
neighbor route-server-client	Specifies on a BGP route server that a neighbor is a route server client.

scope

To define the scope for a Border Gateway Protocol (BGP) routing session and to enter router scope configuration mode, use the **scope** command in router configuration mode. To remove the scope configuration, use the **no** form of this command.

scope {**global** | **vrf** *vrf-name*}

no scope {**global** | **vrf** *vrf-name*}

Syntax Description

global	Configures BGP to use the global routing table or a specific topology table.
vrf	Configures BGP to use a specific VRF routing table.
vrf-name	Name of an existing VRF.

Command Default

No scope is defined for a BGP routing session.

Command Modes

Router configuration

Command History

Release	Modification
12.2(33)SRB	This command was introduced.

Usage Guidelines

A new configuration hierarchy, named scope, has been introduced into the BGP protocol. To implement Multi-Topology Routing (MTR) support for BGP, the scope hierarchy is required, but the scope hierarchy is not limited to MTR use. The scope hierarchy introduces some new configuration modes such as router scope configuration mode. Router scope configuration mode is entered by configuring the **scope** command in router configuration mode, and a collection of routing tables is created when this command is entered. The scope is configured to isolate routing calculation for a single network (globally) or on a per-VRF basis, and BGP commands configured in routing scope configuration mode are referred to as scoped commands. The scope hierarchy can contain one or more address families.

The BGP command-line interface (CLI) has been modified to provide backwards compatibility for pre-MTR BGP configuration and to provide a hierarchal implementation of MTR. From router scope configuration mode, MTR is configured first by entering the **address-family** command to enter the desired address family and then by entering the **topology** command to define the topology

Note Configuring a scope for a BGP routing process removes CLI support for pre-MTR-based configuration.

Examples

The following example defines a global scope that includes both unicast and multicast topology configurations. Another scope is specifically defined only for the VRF named DATA.

```
Router(config)# router bgp 45000
Router(config-router)# scope global
Router(config-router-scope)# bgp default ipv4-unicast
```

```
Router(config-router-scope)# neighbor 172.16.1.2 remote-as 45000
Router(config-router-scope)# neighbor 192.168.3.2 remote-as 50000
Router(config-router-scope)# address-family ipv4 unicast
Router(config-router-scope-af)# topology VOICE
Router(config-router-scope-af)# bgp tid 100
Router(config-router-scope-af)# neighbor 172.16.1.2 activate
Router(config-router-scope-af)# exit
Router(config-router-scope)# address-family ipv4 multicast
Router(config-router-scope-af)# topology base
Router(config-router-scope-af-topo)# neighbor 192.168.3.2 activate
Router(config-router-scope-af-topo)# exit
Router(config-router-scope-af)# exit
Router(config-router-scope)# exit
Router(config-router)# scope vrf DATA
Router(config-router-scope)# neighbor 192.168.1.2 remote-as 40000
Router(config-router-scope)# address-family ipv4
Router(config-router-scope-af)# neighbor 192.168.1.2 activate
Router(config-router-scope-af)# end
```

Related Commands

Command	Description
bgp tid	Configures BGP to accept routes with a specified topology ID.
topology (BGP)	Configures a process to route IP traffic under the specified topology instance.

set as-path

To modify an autonomous system path for BGP routes, use the **set as-path** command in route-map configuration mode. To not modify the autonomous system path, use the **no** form of this command.

set as-path {**tag** | **prepend** *as-path-string*}

no set as-path {**tag** | **prepend** *as-path-string*}

Syntax Description

tag	Converts the tag of a route into an autonomous system path. Applies only when redistributing routes into BGP.
prepend	Appends the string following the keyword **prepend** to the autonomous system path of the route that is matched by the route map. Applies to inbound and outbound BGP route maps.
as-path-string	Number of an autonomous system to prepend to the AS_PATH attribute. The range of values for this argument is any valid autonomous system number from 1 to 65535. Multiple values can be entered; up to 10 AS numbers can be entered. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command.

Command Default

An autonomous system path is not modified.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
11.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.0(32)S12	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.

Release	Modification
Cisco IOS XE Release 2.3	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

The only global BGP metric available to influence the best path selection is the autonomous system path length. By varying the length of the autonomous system path, a BGP speaker can influence the best path selection by a peer further away.

By allowing you to convert the tag into an autonomous system path, the **set as-path tag** variation of this command modifies the autonomous system length. The **set as-path prepend** variation allows you to "prepend" an arbitrary autonomous system path string to BGP routes. Usually the local autonomous system number is prepended multiple times, increasing the autonomous system path length.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Examples

The following example converts the tag of a redistributed route into an autonomous system path:

```
route-map set-as-path-from-tag
 set as-path tag
!
router bgp 100
 redistribute ospf 109 route-map set-as-path-from-tag
```

The following example prepends 100 100 100 to all the routes that are advertised to 10.108.1.1:

```
route-map set-as-path
 match as-path 1
 set as-path prepend 100 100 100
!
router bgp 100
 neighbor 10.108.1.1 route-map set-as-path out
```

The following example prepends 65538, 65538, and 65538 to all the routes that are advertised to 192.168.1.2. This example requires Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, or a later release.

```
route-map set-as-path
 match as-path 1.1
 set as-path prepend 65538 65538 65538
 exit
router bgp 65538
 neighbor 192.168.1.2 route-map set-as-path out
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
router bgp	Configures the BGP routing process.
set tag (IP)	Sets a tag value of the destination routing protocol.

set comm-list delete

To remove communities from the community attribute of an inbound or outbound update, use the **set comm-list delete** command in route-map configuration mode. To remove a previous **set comm-list delete** command, use the **no** form of this command.

set comm-list {*community-list-number* | *community-list-name*} **delete**

no set comm-list {*community-list-number* | *community-list-name*} **delete**

Syntax Description

community-list-number	A standard or expanded community list number. The range of standard community list numbers is from 1 to 99. The range of expanded community list number is from 100 to 500.
community-list-name	A standard or expanded community list name.

Command Default

No communities are removed.

Command Modes

Route-map configuration

Command History

Release	Modification
12.0	This command was introduced.
12.0(10)S	Named community list support was added.
12.0(16)ST	Named community list support was integrated into Cisco IOS Release 12.0(16)ST.
12.1(9)E	Named community list support was integrated into Cisco IOS Release 12.1(9)E.
12.2(8)T	Named community list support was integrated into Cisco IOS Release 12.2(8)T.
12.0(22)S	The maximum number of expanded community lists was increased from 199 to 500 in Cisco IOS Release 12.0(22)S.
12.2(14)S	The maximum number of expanded community lists was increased from 199 to 500 and named community list support were integrated into Cisco IOS Release 12.2(14)S.
12.2(15)T	The maximum number of expanded community lists was increased from 199 to 500 in Cisco IOS Release 12.2(15)T.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This **set** route-map configuration command removes communities from the community attribute of an inbound or outbound update using a route map to filter and determine the communities to be deleted. Depending upon whether the route map is applied to the inbound or outbound update for a neighbor, each community that passes the route map **permit** clause and matches the given community list will be removed from the community attribute being received from or sent to the Border Gateway Protocol (BGP) neighbor.

Each entry of a standard community list should list only one community when used with the **set comm-list delete** command. For example, in order to be able to delete communities 10:10 and 10:20, you must use the following format to create the entries:

```
ip community-list 500 permit 10:10
ip community-list 500 permit 10:20
```

The following format for a community list entry, while acceptable otherwise, does not work with the **set comm-list delete** command:

```
config ip community-list 500 permit 10:10 10:20
```

When both the **set community** *community-number* and **set comm-list delete** commands are configured in the same sequence of a route map attribute, the deletion operation (**set comm-list delete**) is performed before the set operation (**set community** *community-number*).

Examples

In the following example, the communities 100:10 and 100:20 (if present) will be deleted from updates received from 172.16.233.33. Also, except for 100:50, all communities beginning with 100: will be deleted from updates sent to 172.16.233.33.

```
router bgp 100
 neighbor 172.16.233.33 remote-as 120
 neighbor 172.16.233.33 route-map ROUTEMAPIN in
 neighbor 172.16.233.33 route-map ROUTEMAPOUT out
!
ip community-list 500 permit 100:10
ip community-list 500 permit 100:20
!
ip community-list 120 deny   100:50
ip community-list 120 permit 100:.*
!
route-map ROUTEMAPIN permit 10
 set comm-list 500 delete
!
route-map ROUTEMAPOUT permit 10
 set comm-list 120 delete
```

Related Commands

Command	Description
set community	Sets the BGP communities attribute.

set community

To set the BGP communities attribute, use the **set community** route map configuration command. To delete the entry, use the **no** form of this command.

set community {*community-number* [**additive**] [*well-known-community*] | **none**}

no set community

Syntax Description

community-number	Specifies that community number. Valid values are from 1 to 4294967200, **no-export**, or **no-advertise**.
additive	(Optional) Adds the community to the already existing communities.
well-known-community	(Optional) Well know communities can be specified by using the following keywords: • internet • local-as • no-advertise • no-export
none	(Optional) Removes the community attribute from the prefixes that pass the route map.

Command Default

No BGP communities attributes exist.

Command Modes

Route-map configuration

Command History

Release	Modification
10.3	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

You must have a match clause (even if it points to a "permit everything" list) if you want to set tags.

Use the **route-map** global configuration command, and the **match** and **set** route map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **set** route map configuration commands specify the redistribution *set actions* to be performed when all of the match criteria of a route map are met. When all match criteria are met, all set actions are performed.

Examples

In the following example, routes that pass the autonomous system path access list 1 have the community set to 109. Routes that pass the autonomous system path access list 2 have the community set to no-export (these routes will not be advertised to any external BGP [eBGP] peers).

```
route-map set_community 10 permit
 match as-path 1
 set community 109

route-map set_community 20 permit
 match as-path 2
 set community no-export
```

In the following similar example, routes that pass the autonomous system path access list 1 have the community set to 109. Routes that pass the autonomous system path access list 2 have the community set to local-as (the router will not advertise this route to peers outside the local autonomous system.

```
route-map set_community 10 permit
 match as-path 1
 set community 109

route-map set_community 20 permit
 match as-path 2
 set community local-as
```

Related Commands

Command	Description
ip community-list	Creates a community list for BGP and control access to it.
match community	Matches a BGP community.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set comm-list delete	Removes communities from the community attribute of an inbound or outbound update.
show ip bgp community	Displays routes that belong to specified BGP communities.

set dampening

To set the BGP route dampening factors, use the **set dampening** route map configuration command. To disable this function, use the **no** form of this command.

set dampening *half-life reuse suppress max-suppress-time*

no set dampening

Syntax Description

half-life	Time (in minutes) after which a penalty is decreased. Once the route has been assigned a penalty, the penalty is decreased by half after the half life period (which is 15 minutes by default). The process of reducing the penalty happens every 5 seconds. The range of the half life period is from 1 to 45 minutes. The default is 15 minutes.
reuse	Unsuppresses the route if the penalty for a flapping route decreases enough to fall below this value. The process of unsuppressing routes occurs at 10-second increments. The range of the reuse value is from 1 to 20000; the default is 750.
suppress	Suppresses a route when its penalty exceeds this limit. The range is from 1 to 20000; the default is 2000.
max-suppress-time	Maximum time (in minutes) a route can be suppressed. The range is from 1 to 20000; the default is four times the *half-life* value. If the *half-life* value is allowed to default, the maximum suppress time defaults to 60 minutes.

Defaults

This command is disabled by default.

Command Modes

Route-map configuration

Command History

Release	Modification
11.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use the **route-map** global configuration command, and the **match** and **set** route-map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

When a BGP peer is reset, the route is withdrawn and the flap statistics cleared. In this instance, the withdrawal does not incur a penalty even though route flap dampening is enabled.

Examples

The following example sets the half life to 30 minutes, the reuse value to 1500, the suppress value to 10000; and the maximum suppress time to 120 minutes:

```
route-map tag
 match as path 10
 set dampening 30 1500 10000 120
!
router bgp 100
 neighbor 172.16.233.52 route-map tag in
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
match community	Matches a BGP community.
match interface (IP)	Distributes routes that have their next hop out one of the interfaces specified.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
match ip next-hop	Redistributes any routes that have a next hop router address passed by one of the access lists specified.
match ip route-source	Redistributes routes that have been advertised by routers and access servers at the address specified by the access lists.
match metric (IP)	Redistributes routes with the metric specified.
match route-type (IP)	Redistributes routes of the specified type.
match tag	Redistributes routes in the routing table that match the specified tags.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set automatic-tag	Automatically computes the tag value.
set community	Sets the BGP communities attribute.
set ip next-hop	Specifies the address of the next hop.
set level (IP)	Indicates where to import routes.
set local-preference	Specifies a preference value for the autonomous system path.
set metric (BGP, OSPF, RIP)	Sets the metric value for a routing protocol.
set metric-type	Sets the metric type for the destination routing protocol.
set origin (BGP)	Sets the BGP origin code.
set tag (IP)	Sets the value of the destination routing protocol.
set weight	Specifies the BGP weight for the routing table.
show route-map	Displays all route maps configured or only the one specified.

set extcommunity

To set Border Gateway Protocol (BGP) extended community attributes, use the **set extcommunity** command in route-map configuration mode. To delete the entry, use the **no** form of this command.

set extcommunity {**rt** [*extended-community-value*] [**additive**] | **soo** [*extended-community-value*]}

no set extcommunity

Syntax Description

rt	Specifies the route target (RT) extended community attribute.
soo	Specifies the site of origin (SOO) extended community attribute.
extended-community-value	(Optional) Specifies the value to be set. The value can be one of the following combinations: • *autonomous-system-number*:*network-number* • *ip-address*:*network-number* • *ipv6-address*:*network-number* The colon is used to separate the autonomous system number and network number, the IP address and network number, or the IPv6 address and network number. • In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, 4-byte autonomous system numbers are supported in the range from 65536 to 4294967295 in asplain notation and in the range from 1.0 to 65535.65535 in asdot notation. • In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, 4-byte autonomous system numbers are supported in the range from 1.0 to 65535.65535 in asdot notation only. For more details about autonomous system number formats, see the **router bgp** command.
additive	(Optional) Adds a route target to the existing route target list without replacing any existing route targets.

Command Default

Specifying new route targets with the **rt** keyword replaces existing route targets by default, unless the **additive** keyword is used. The use of the **additive** keyword adds the new route target to the existing route target list but does not replace any existing route targets.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
12.1	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SRB	Support for IPv6 was added.

Release	Modification
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.
12.2(33)SB	Support for IPv6 was added. and this command was integrated into Cisco IOS Release 12.2(33)SB.
12.0(32)S12	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.0(32)SY8	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.4(24)T	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
Cisco IOS XE Release 2.3	This command was modified. Support for 4-byte autonomous system numbers in asdot notation only was added.
12.2(33)SXI1	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

Extended community attributes are used to configure, filter, and identify routes for virtual routing and forwarding instances (VRFs) and Multiprotocol Label Switching (MPLS) Virtual Private Networks (VPNs).

The **set extcommunity** command is used to configure set clauses that use extended community attributes in route maps. All of the standard rules of match and set clauses apply to the configuration of extended community attributes.

The route target (RT) extended community attribute is configured with the **rt** keyword. This attribute is used to identify a set of sites and VRFs that may receive routes that are tagged with the configured route target. Configuring the route target extended attribute with a route allows that route to be placed in the per-site forwarding tables that are used for routing traffic that is received from corresponding sites.

The site of origin (SOO) extended community attribute is configured with the **soo** keyword. This attribute uniquely identifies the site from which the Provider Edge (PE) router learned the route. All routes learned from a particular site must be assigned the same SOO extended community attribute, whether a site is connected to a single PE router or multiple PE routers. Configuring this attribute prevents routing loops from occurring when a site is multihomed. The SOO extended community attribute is configured on the interface and is propagated into BGP through redistribution. The SOO can be applied to routes that are learned from VRFs. The SOO should not be configured for stub sites or sites that are not multihomed.

In Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, the Cisco implementation of 4-byte autonomous system numbers uses asplain—65538 for example—as the default regular expression match and output display format for autonomous system numbers, but you can configure 4-byte autonomous system numbers in both the asplain format and the asdot format as described in RFC 5396. To change the default regular expression

match and output display of 4-byte autonomous system numbers to asdot format, use the **bgp asnotation dot** command followed by the **clear ip bgp *** command to perform a hard reset of all current BGP sessions.

In Cisco IOS Release 12.0(32)S12, 12.4(24)T, and Cisco IOS XE Release 2.3, the Cisco implementation of 4-byte autonomous system numbers uses asdot—1.2 for example—as the only configuration format, regular expression match, and output display, with no asplain support.

Examples

The following example sets the route target to extended community attribute 100:2 for routes that are permitted by the route map:

```
Router(config)# access-list 2 permit 192.168.78.0 255.255.255.0
Router(config)# route-map MAP_NAME permit 10
Router(config-route-map)# match ip-address 2
Router(config-route-map)# set extcommunity rt 100:2
```

The following example sets the route target to extended community attribute 100:3 for routes that are permitted by the route map. The use of the **additive** keyword adds route target 100:3 to the existing route target list but does not replace any existing route targets.

```
Router(config)# access-list 3 permit 192.168.79.0 255.255.255.0
Router(config)# route-map MAP_NAME permit 10
Router(config-route-map)# match ip-address 3
Router(config-route-map)# set extcommunity rt 100:3 additive
```

Note Configuring route targets with the **set extcommunity** command will replace existing route targets, unless the **additive** keyword is used.

The following example sets the site of origin to extended community attribute 100:4 for routes that are permitted by the route map:

```
Router(config)# access-list 4 permit 192.168.80.0 255.255.255.0
Router(config)# route-map MAP_NAME permit 10
Router(config-route-map)# match ip-address 4
Router(config-route-map)# set extcommunity soo 100:4
```

In IPv6, the following example sets the SoO to extended community attribute 100:28 for routes that are permitted by the route map:

```
(config)# router bgp 100
(config-router)# address-family ipv6 vrf red
(config-router-af)# neighbor 8008::72a route-map setsoo in
(config-router-af)# exit
(config-router)# route-map setsoo permit 10
(config-router)# set extcommnunity soo 100:28
```

The following example available in Cisco IOS Release 12.0(32)SY8, 12.0(33)S3, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, Cisco IOS XE Release 2.4, and later releases, shows how to create a VRF with a route-target that uses a 4-byte autonomous system number, 65537 in asplain format, and how to set the route-target to extended community value 65537:100 for routes that are permitted by the route map.

```
Router(config)# ip vrf vpn_red
Router(config-vrf)# rd 64500:100
Router(config-vrf)# route-target both 65537:100
Router(config-vrf)# exit
Router(config)# route-map rt_map permit 10
Router(config-route-map)# set extcommunity rt 65537:100
Router(config-route-map)# end
```

The following example available in Cisco IOS Release 12.0(32)SY8, 12.0(32)S12, 12.2(33)SRE, 12.2(33)XNE, 12.2(33)SXI1, 12.4(24)T, Cisco IOS XE Release 2.3, and later releases, shows how to create a VRF with a route-target that uses a 4-byte autonomous system number, 1.1 in asdot format, and how to set the SoO to extended community attribute 1.1:100 for routes that are permitted by the route map.

```
Router(config)# ip vrf vpn_red
Router(config-vrf)# rd 64500:100
Router(config-vrf)# route-target both 1.1:100
Router(config-vrf)# exit
Router(config)# route-map soo_map permit 10
Router(config-route-map)# set extcommunity soo 1.1:100
Router(config-route-map)# end
```

Related Commands

Command	Description
bgp asnotation dot	Changes the default display and the regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
ip extcommunity-list	Creates an extended community list and controls access to it.
match extcommunity	Matches a BGP VPN extended community list.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
router bgp	Configures the BGP routing process.
route-target	Creates a route target extended community for a VRF.
show ip extcommunity-list	Displays routes that are permitted by the extended community list.
show route-map	Displays all route maps configured or only the one specified.

set extcommunity cost

To create a set clause to apply the cost community attribute to routes that pass through a route map, use the **set extcommunity cost** command in route-map configuration mode. To delete the cost community set clause, use the **no** form of this command.

set extcommunity cost [**igp** | **pre-bestpath**] *community-id cost-value*

no set extcommunity cost [**igp**] *community-id cost-value*

Syntax Description

igp	(Optional) Specifies the IGP point of insertion (POI). The configuration of this keyword forces the cost community to be evaluated after the IGP distance to the next hop has been compared. If this keyword is not specified, IGP is the default POI.
community-id	The ID for the configured extended community. The range is from 0 to 255.
cost-value	The configured cost that is set for matching paths in the route map. The range is from 0 to 4294967295.

Command Default

The default cost value is applied to routes that are not configured with the cost community attribute when cost community filtering is enabled. The default *cost-value* is half of the maximum value (4294967295) or 2147483647.

Command Modes

Route-map configuration

Command History

Release	Modification
12.0(24)S	This command was introduced into Cisco IOS Release 12.0(24)S.
12.3(2)T	This command was integrated.
12.2(18)S	This command was integrated.
12.0(27)S	Support for mixed EIGRP MPLS VPN network topologies that contain back door routes was introduced into Cisco IOS Release 12.0(27)S.
12.3(8)T	Support for mixed EIGRP MPLS VPN network topologies that contain back door routes was introduced into Cisco IOS Release 12.3(8)T.
12.2(25)S	Support for mixed EIGRP MPLS VPN network topologies that contain back door routes was introduced into Cisco IOS Release 12.2(25)S.
12.2(27)SBC	This command was integrated into Cisco IOS Release 12.2(27)SBC.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The cost community attribute is applied to internal routes by configuring the **set extcommunity cost** command in a route map. The cost community set clause is configured with a cost community ID number (0-255) and a cost community number value (0-4294967295). The path with the lowest cost community number is preferred. In the case where two paths have been configured with the same cost community value, the path selection process will then prefer the path with the lower community ID.

The BGP Cost Community feature can be configured only within the same autonomous-system or confederation. The cost community is a non-transitive extended community. The cost community is passed to internal BGP (iBGP) and confederation peers only and is not passed to external BGP (eBGP) peers. The cost community allows you to customize the local preference and best path selection process for specific paths. The cost extended community attribute is propagated to iBGP peers when extended community exchange is enabled with the **neighbor send-community** command.

The following commands can be used to apply the route map with the cost community set clause:

- **aggregate-address**
- **neighbor default-originate route-map** {**in** | **out**}
- **neighbor route-map**
- **network route-map**
- **redistribute route-map**

Multiple cost community set clauses may be configured with the **set extcommunity cost** command in a single route map block or sequence. However, each set clause must be configured with a different ID value for each point of insertion (POI).

Aggregate routes and multipaths are supported by the BGP Cost Community feature. The cost community attribute can be applied to either type of route. The cost community attribute is passed to the aggregate or multipath route from component routes that carry the cost community attribute. Only unique IDs are passed, and only the highest cost of any individual component route will be applied to the aggregate on a per-ID basis. If multiple component routes contain the same ID, the highest configured cost is applied to the route. If one or more component routes does not carry the cost community attribute or if the component routes are configured with different IDs, then the default value (2147483647) will be advertised for the aggregate or multipath route.

Note The BGP cost community attribute must be supported on all routers in an autonomous system or confederation before cost community filtering is configured. The cost community should be applied consistently throughout the local autonomous system or confederation to avoid potential routing loops.

Support for EIGRP MPLS VPN Back Door Links

The “pre-bestpath” point of insertion (POI) has been introduced in the BGP Cost Community feature to support mixed EIGRP VPN network topologies that contain VPN and backdoor links. This POI is applied automatically to EIGRP routes that are redistributed into BGP. The “pre-best path” POI carries the EIGRP route type and metric. This POI influences the best path calculation process by influencing BGP to consider this POI before any other comparison step. No configuration is required. This feature is enabled automatically for EIGRP VPN sites when a supporting is installed to a PE, CE, or back door router.

Examples

The following example configuration shows the configuration of the **set extcommunity cost** command. The following example applies the cost community ID of 1 and cost community value of 100 to routes that are permitted by the route map. This configuration will cause the best path selection process to prefer this route over other equal cost paths that were not permitted by this route map sequence.

```
Router(config)# router bgp 50000
Router(config-router)# neighbor 10.0.0.1 remote-as 50000
Router(config-router)# neighbor 10.0.0.1 update-source Loopback 0
Router(config-router)# address-family ipv4
Router(config-router-af)# neighbor 10.0.0.1 activate
Router(config-router-af)# neighbor 10.0.0.1 route-map COST1 in
Router(config-router-af)# neighbor 10.0.0.1 send-community both
Router(config-router-af)# exit
Router(config)# route-map COST1 permit 10
Router(config-route-map)# match ip-address 1
Router(config-route-map)# set extcommunity cost 1 100
```

Related Commands

Command	Description
aggregate-address	Creates an aggregate entry in a BGP or multicast BGP database.
bgp bestpath cost-community ignore	Configures a router that is running BGP to not evaluate the cost community attribute during the best path selection process.
neighbor default-originate	Allows a BGP speaker (the local router) to send the default route 0.0.0.0 to a neighbor for use as a default route.
neighbor route-map	Applies a route map to incoming or outgoing routes.
network (BGP and multiprotocol BGP)	Specifies the networks to be advertised by the BGP and multiprotocol BGP routing processes.
redistribute (IP)	Redistributes routes from one routing domain into another routing domain.
show ip bgp	Displays entries in the BGP routing table.

set ip next-hop (BGP)

To indicate where to output packets that pass a match clause of a route map for policy routing, use the **set ip next-hop** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set ip next-hop *ip-address* [*... ip-address*] [**peer-address**]

no set ip next-hop *ip-address* [*... ip-address*] [**peer-address**]

Syntax Description

ip-address	IP address of the next hop to which packets are output. It need not be an adjacent router.
peer-address	(Optional) Sets the next hop to be the BGP peering address.

Defaults

This command is disabled by default.

Command Modes

Route-map configuration

Command History

Release	Modification
11.0	This command was introduced.
12.0	The **peer-address** keyword was added.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.
Cisco IOS XE Release 2.1	This command was introduced on Cisco ASR 1000 Series Routers.

Usage Guidelines

An ellipsis (...) in the command syntax indicates that your command input can include multiple values for the *ip-address* argument.

Use the **ip policy route-map** interface configuration command, the **route-map** global configuration command, and the **match** and **set** route-map configuration commands to define the conditions for policy routing packets. The **ip policy route-map** command identifies a route map by name. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which policy routing occurs. The **set** commands specify the *set actions*—the particular routing actions to perform if the criteria enforced by the **match** commands are met.

If the first next hop specified with the **set ip next-hop** command is down, the optionally specified IP addresses are tried in turn.

When the **set ip next-hop command** is used with the **peer-address** keyword in an inbound route map of a BGP peer, the next hop of the received matching routes will be set to be the neighbor peering address, overriding any third-party next hops. So the same route map can be applied to multiple BGP peers to override third-party next hops.

When the **set ip next-hop** command is used with the **peer-address** keyword in an outbound route map of a BGP peer, the next hop of the advertised matching routes will be set to be the peering address of the local router, thus disabling the next hop calculation. The **set ip next-hop** command has finer granularity than the (per-neighbor) **neighbor next-hop-self** command, because you can set the next hop for some routes, but not others. The **neighbor next-hop-self** command sets the next hop for all routes sent to that neighbor.

The set clauses can be used in conjunction with one another. They are evaluated in the following order:

1. **set ip next-hop**
2. **set interface**
3. **set ip default next-hop**
4. **set default interface**

Note To avoid a common configuration error for reflected routes, do not use the **set ip next-hop** command in a route map to be applied to BGP route reflector clients.

Examples

In the following example, three routers are on the same FDDI LAN (with IP addresses 10.1.1.1, 10.1.1.2, and 10.1.1.3). Each is in a different autonomous system. The **set ip next-hop peer-address** command specifies that traffic from the router (10.1.1.3) in remote autonomous system 300 for the router (10.1.1.1) in remote autonomous system 100 that matches the route map is passed through the router bgp 200, rather than sent directly to the router (10.1.1.1) in autonomous system 100 over their mutual connection to the LAN.

```
router bgp 200
neighbor 10.1.1.3 remote-as 300
neighbor 10.1.1.3 route-map set-peer-address out
neighbor 10.1.1.1 remote-as 100
route-map set-peer-address permit 10
set ip next-hop peer-address
```

Related Commands

Command	Description
ip policy route-map	Identifies a route map to use for policy routing on an interface.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
match length	Bases policy routing on the Level 3 length of a packet.
neighbor next-hop-self	Disables next hop processing of BGP updates on the router.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol to another, or enables policy routing.
set default interface	Indicates where to output packets that pass a match clause of a route map for policy routing and that have no explicit route to the destination.

set interface	Indicates where to output packets that pass a match clause of a route map for policy routing.
set ip default next-hop	Indicates where to output packets that pass a match clause of a route map for policy routing and for which the Cisco IOS software has no explicit route to a destination.

set metric (BGP-OSPF-RIP)

To set the metric value for a routing protocol, use the **set metric** command in route-map configuration mode. To return to the default metric value, use the **no** form of this command.

set metric *metric-value*

no set metric *metric-value*

Syntax Description

metric-value	Metric value; an integer from –294967295 to 294967295. This argument applies to all routing protocols except Enhanced Interior Gateway Routing Protocol (EIGRP).

Defaults

The dynamically learned metric value.

Command Modes

Route-map configuration

Command History

Release	Modification
10.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

We recommend that you consult your Cisco technical support representative before changing the default value.

Use the **route-map** global configuration command, and the **match** and **set** route-map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **set** route-map configuration commands specify the redistribution *set actions* to be performed when all the match criteria of a route map are met. When all match criteria are met, all set actions are performed.

Examples

The following example sets the metric value for the routing protocol to 100:

```
route-map set-metric
 set metric 100
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
match community	Matches a BGP community.
match interface (IP)	Distributes routes that have their next hop out one of the interfaces specified.
match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
match ip next-hop	Redistributes any routes that have a next hop router address passed by one of the access lists specified.
match ip route-source	Redistributes routes that have been advertised by routers and access servers at the address specified by the access lists.
match metric (IP)	Redistributes routes with the metric specified.
match route-type (IP)	Redistributes routes of the specified type.
match tag	Redistributes routes in the routing table that match the specified tags.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set automatic-tag	Automatically computes the tag value.
set community	Sets the BGP communities attribute.
set ip next-hop	Specifies the address of the next hop.
set level (IP)	Indicates where to import routes.
set local-preference	Specifies a preference value for the autonomous system path.
set metric (BGP, OSPF, RIP)	Sets the metric value for a routing protocol.
set metric-type	Sets the metric type for the destination routing protocol.
set origin (BGP)	Sets the BGP origin code.
set tag (IP)	Sets the value of the destination routing protocol.

set metric-type internal

To set the Multi Exit Discriminator (MED) value on prefixes advertised to external BGP (eBGP) neighbors to match the Interior Gateway Protocol (IGP) metric of the next hop, use the **set metric-type internal** command in route-map configuration mode. To return to the default, use the **no** form of this command.

set metric-type internal

no set metric-type internal

Syntax Description

This command has no arguments or keywords.

Defaults

This command is disabled by default.

Command Modes

Route-map configuration

Command History

Release	Modification
10.3	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

This command will cause BGP to advertise a MED value that corresponds to the IGP metric associated with the next hop of the route. This command applies to generated, internal BGP (iBGP)-, and eBGP-derived routes.

If this command is used, multiple BGP speakers in a common autonomous system can advertise different MED values for a particular prefix. Also, note that if the IGP metric changes, BGP will readvertise the route every 10 minutes.

Use the **route-map** global configuration command and the **match** and **set** route-map configuration commands to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **set** route-map configuration commands specify the redistribution *set actions* to be performed when all of the match criteria of the route map are met. When all match criteria are met, all set actions are performed.

Note This command is not supported for redistributing routes into Border Gateway Protocol (BGP).

Examples

In the following example, the MED value for all the advertised routes to neighbor 172.16.2.3 is set to the corresponding IGP metric of the next hop:

```
router bgp 109
 network 172.16.0.0
 neighbor 172.16.2.3 remote-as 200
 neighbor 172.16.2.3 route-map setMED out
!
route-map setMED permit 10
 match as-path 1
 set metric-type internal
!
 ip as-path access-list 1 permit .*
```

Related Commands

Command	Description
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.

set origin (BGP)

To set the BGP origin code, use the **set origin** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set origin {**igp** | **egp** *autonomous-system-number* | **incomplete**}

no set origin {**igp** | **egp** *autonomous-system-number* | **incomplete**}

Syntax Description

igp	Remote Interior Gateway Protocol (IGP) system.
egp	Local Exterior Gateway Protocol (EGP) system.
autonomous-system-number	Number of a remote autonomous system number. The range of values for this argument is any valid autonomous system number from 1 to 65535.
incomplete	Unknown heritage.

Command Default

The origin of the route is based on the path information of the route in the main IP routing table.

Command Modes

Route-map configuration (config-route-map)

Command History

Release	Modification
10.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(14)SX	This command was integrated into Cisco IOS Release 12.2(14)SX.
12.4(2)T	This command was modified. The **egp** keyword and *autonomous-system-number* argument were removed.
12.0(33)S3	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
Cisco IOS XE Release 2.4	This command was modified. Support for asplain notation was added and the default format for 4-byte autonomous system numbers is now asplain.
12.2(33)SRE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.
12.2(33)XNE	This command was modified. Support for 4-byte autonomous system numbers in asplain and asdot notation was added.

Usage Guidelines

You must have a match clause (even if it points to a “permit everything” list) if you want to set the origin of a route. Use this command to set a specific origin when a route is redistributed into BGP. When routes are redistributed, the origin is usually recorded as incomplete, identified with a ? in the BGP table.

Use the **route-map** global configuration command, and the **match** and **set** route-map configuration commands, to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands

specify the *match criteria*—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the *set actions*—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **set** route-map configuration commands specify the redistribution *set actions* to be performed when all of the match criteria of a route map are met. When all match criteria are met, all set actions are performed.

Examples

The following example sets the origin of routes that pass the route map to IGP:

```
route-map set_origin
 match as-path 10
 set origin igp
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
router bgp	Configures the BGP routing process.
set as-path	Modifies an autonomous system path for BGP routes.

set traffic-index

To indicate how to classify packets that pass a match clause of a route map for Border Gateway Protocol (BGP) policy accounting, use the **set traffic-index** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set traffic-index *bucket-number*

no set traffic-index *bucket-number*

Syntax Description

bucket-number	Number that represents a bucket into which packet and byte statistics are collected for a specific traffic classification. The range is from 1 to 64.

Command Default

Routing traffic is not classified.

Command Modes

Route-map configuration

Command History

Release	Modification
12.0(9)S	This command was introduced.
12.0(17)ST	This command was integrated into Cisco IOS Release 12.0(17)ST.
12.2(13)T	This command was integrated into Cisco IOS Release 12.2(13)T.
12.0(22)S	Support for 64 buckets was added for the Cisco 12000 series Internet router.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.
12.3(4)T	This command was integrated into Cisco IOS Release 12.3(4)T and support for 64 buckets was added for all platforms.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

Use the **set traffic-index** route-map configuration command, the **route-map** global configuration command, and a **match** route-map configuration command to define the conditions for BGP policy accounting. The **match** commands specify the *match criteria*—the conditions under which policy routing occurs. The **set traffic-index** command specifies the *set actions*—the particular routing actions to perform if the criteria specified by the **match** commands are met.

Examples

In the following example, an index for BGP policy accounting is set according to autonomous system path criteria:

```
route-map buckets permit 10
 match as-path 1
 set traffic-index 1
```

Related Commands

Command	Description
bgp-policy	Enables BGP policy accounting or policy propagation on an interface.
route-map	Defines the conditions for redistributing routes from one routing protocol to another, or enables policy routing.

set weight

To specify the BGP weight for the routing table, use the **set weight** command in route-map configuration mode. To delete an entry, use the **no** form of this command.

set weight *number*

no set weight *number*

Syntax Description

number	Weight value. It can be an integer ranging from 0 to 65535.

Defaults

The weight is not changed by the specified route map.

Command Modes

Route-map configuration

Command History

Release	Modification
10.0	This command was introduced.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2SX	This command is supported in the Cisco IOS Release 12.2SX train. Support in a specific 12.2SX release of this train depends on your feature set, platform, and platform hardware.

Usage Guidelines

The implemented weight is based on the first matched autonomous system path. Weights indicated when an autonomous system path is matched override the weights assigned by global **neighbor** commands. In other words, the weights assigned with the **set weight** route-map configuration command override the weights assigned using the **neighbor weight** command.

Examples

The following example sets the BGP weight for the routes matching the autonomous system path access list to 200:

```
route-map set-weight
 match as-path 10
 set weight 200
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
match community	Matches a BGP community.
match interface (IP)	Distributes routes that have their next hop out one of the interfaces specified.

match ip address	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on packets.
match ip next-hop	Redistributes any routes that have a next hop router address passed by one of the access lists specified.
match ip route-source	Redistributes routes that have been advertised by routers and access servers at the address specified by the access lists.
match metric (IP)	Redistributes routes with the metric specified.
match route-type (IP)	Redistributes routes of the specified type.
match tag	Redistributes routes in the routing table that match the specified tags.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
set automatic-tag	Automatically computes the tag value.
set community	Sets the BGP communities attribute.
set ip next-hop	Specifies the address of the next hop.
set level (IP)	Indicates where to import routes.
set local-preference	Specifies a preference value for the autonomous system path.
set metric (BGP, OSPF, RIP)	Sets the metric value for a routing protocol.
set metric-type	Sets the metric type for the destination routing protocol.
set origin (BGP)	Sets the BGP origin code.
set tag (IP)	Sets the value of the destination routing protocol.
set weight	Specifies the BGP weight for the routing table.

show bgp all community

To display routes for all address families belonging to a particular Border Gateway Protocol (BGP) community, use the **show bgp all community** command in user EXEC or privileged EXEC configuration mode.

show bgp all community [*community-number...*[*community-number*]] [**local-as**] [**no-advertise**] [**no-export**] [**exact-match**]

Syntax Description

community-number	(Optional) Displays the routes pertaining to the community numbers specified. • You can specify multiple community numbers. The range is from 1 to 4294967295 or AA:NN (autonomous system:community number, which is a 2-byte number).
local-as	(Optional) Displays only routes that are not sent outside of the local autonomous system (well-known community).
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).
exact-match	(Optional) Displays only routes that match exactly with the BGP community list specified. **Note** The availability of keywords in the command depends on the command mode. The **exact-match** keyword is not available in user EXEC mode.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.3(2)T	This command was introduced.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(25)SG	This command was integrated into Cisco IOS Release 12.2(25)SG.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(33)SXH	This command was integrated into Cisco IOS Release 12.2(33)SXH.

Usage Guidelines

You can enter the **local-as**, **no-advertise** and **no-export** keywords in any order. You can set the communities using the **set community** command.

When using the **bgp all community** command, be sure to enter the numerical communities before the well-known communities.

For example, the following string is not valid:

```
Router# show bgp all community local-as 111:12345
```

Use the following string instead:

```
Router# show bgp all community 111:12345 local-as
```

Examples

The following is sample output from the **show bgp all community** command, specifying communities of 1, 2345, and 6789012:

```
Router# show bgp all community 1 2345 6789012 no-advertise local-as no-export exact-match

For address family: IPv4 Unicast

BGP table version is 5, local router ID is 30.0.0.5
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
*> 10.0.3.0/24      10.0.0.4                               0 4 3 ?
*> 10.1.0.0/16      10.0.0.4                 0             0 4 ?
*> 10.12.34.0/24    10.0.0.6                 0             0 6 ?
```

Table 9 describes the significant fields shown in the display.

Table 9 show bgp all community Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	The router ID of the router on which the BGP communities are set to display. A 32-bit number written as 4 octets separated by periods (dotted-decimal format).
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP session.
Origin codes	Indicates the origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from the Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from the Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	The network address and network mask of a network entity. The type of address depends on the address family.

Table 9 show bgp all community Field Descriptions (continued)

Field	Description
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. The type of address depends on the address family.
Metric	The value of the inter autonomous system metric. This field is not used frequently.
LocPrf	Local preference value as set with the **set local-preference** command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

Related Commands

Command	Description
set community	Sets BGP communities.
set local-preference	Specifies a preference value for the autonomous system path.

show bgp all neighbors

To display information about Border Gateway Protocol (BGP) connections to neighbors of all address families, use the **show bgp all neighbors** command in user EXEC or privileged EXEC mode.

show bgp all neighbors [*ip-address* | *ipv6-address*] [**advertised-routes** | **dampened-routes** | **flap-statistics** | **paths** [*reg-exp*] | **policy** [**detail**] | **received prefix-filter** | **received-routes** | **routes**]

Syntax Description

ip-address	(Optional) IP address of a neighbor. If this argument is omitted, information about all neighbors is displayed.
ipv6-address	(Optional) Address of the IPv6 BGP-speaking neighbor. This argument must be in the form documented in RFC 2373 where the address is specified in hexadecimal using 16-bit values between colons.
advertised-routes	(Optional) Displays all routes that have been advertised to neighbors.
dampened-routes	(Optional) Displays the dampened routes received from the specified neighbor (for external BGP peers only).
flap-statistics	(Optional) Displays the flap statistics of the routes learned from the specified neighbor (for external BGP peers only).
paths *reg-exp*	(Optional) Displays autonomous system paths learned from the specified neighbor. An optional regular expression can be used to filter the output.
policy	(Optional) Displays the policies applied to neighbor per address family.
detail	(Optional) Displays detailed policy information such as route maps, prefix lists, community lists, Access Control Lists (ACLs), and autonomous system path filter lists.
received prefix-filter	(Optional) Displays the prefix-list (outbound route filter [ORF]) sent from the specified neighbor.
received-routes	(Optional) Displays all received routes (both accepted and rejected) from the specified neighbor.
routes	(Optional) Displays all routes that are received and accepted. The output displayed when this keyword is entered is a subset of the output displayed by the **received-routes** keyword.

Command Default

The output of this command displays information for all neighbors.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.3(26)	This command was introduced.
12.2(18)S	This command was integrated into Cisco IOS Release 12.2(18)S and was made available in privileged EXEC mode.

Release	Modification
12.2(19)S	This command was made available in user EXEC mode.
12.2(28)SB	This command was integrated into Cisco IOS Release 12.2(28)SB.
12.2(33)SRA	This command was integrated into Cisco IOS Release 12.2(33)SRA.
12.2(18)SXF	This command was integrated into Cisco IOS Release 12.2(18)SXF.
12.4(11)T	This command was integrated into Cisco IOS Release 12.4(11)T. The **policy** keyword was added.
12.2(33)SRB	The **policy** keyword was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE Release 2.1	This command was integrated into Cisco IOS XE Release 2.1.

Usage Guidelines

Use the **show bgp all neighbors** command to display BGP and TCP connection information for neighbor sessions specific to address families such as IPv4, IPv6, Network Service Access Point (NSAP), Virtual Private Network (VPN) v4, and VPNv6.

Examples

The following example shows output of the **show bgp all neighbors** command:

```
Router# show bgp all neighbors

For address family: IPv4 Unicast
BGP neighbor is 172.16.232.53,  remote AS 100, external link
 Member of peer-group internal for session parameters
  BGP version 4, remote router ID 172.16.232.53
  BGP state = Established, up for 13:40:17
  Last read 00:00:09, hold time is 180, keepalive interval is 60 seconds
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
                         Sent       Rcvd
    Opens:                  3          3
    Notifications:          0          0
    Updates:                0          0
    Keepalives:           113        112
    Route Refresh:          0          0
    Total:                116        11
  Default minimum time between advertisement runs is 5 seconds

  Connections established 22; dropped 21
  Last reset 13:47:05, due to BGP Notification sent, hold time expired
External BGP neighbor may be up to 2 hops away.
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 3FFE:700:20:1::12, Local port: 55345
Foreign host: 3FFE:700:20:1::11, Foreign port: 179

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

Event Timers (current time is 0x1A0D543C):
Timer          Starts    Wakeups            Next
Retrans          1218          5             0x0
TimeWait            0          0             0x0
AckHold          3327       3051             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
```

```
DeadWait            0          0             0x0

iss: 1805423033  snduna: 1805489354  sndnxt: 1805489354     sndwnd:  15531
irs:  821333727  rcvnxt:  821591465  rcvwnd:      15547  delrcvwnd:    837

SRTT: 300 ms, RTTO: 303 ms, RTV: 3 ms, KRTT: 0 ms
minRTT: 8 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1420 bytes):
Rcvd: 4252 (out of order: 0), with data: 3328, total data bytes: 257737
Sent: 4445 (retransmit: 5), with data: 4445, total data bytes: 244128

For address family: IPv6 Unicast

For address family: IPv4 MDT

For address family: VPNv4 Unicast

For address family: VPNv6 Unicast

For address family: IPv4 Multicast

For address family: IPv6 Multicast

For address family: NSAP Unicast
```

Table 10 describes the significant fields shown in the display.

Table 10 show bgp all neighbors Field Descriptions

Field	Description
For address family:	Address family to which the following fields refer.
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system number of the neighbor.
external link	External Border Gateway Protocol (eBGP) peer.
BGP version	BGP version being used to communicate with the remote router.
remote router ID	IP address of the neighbor.
BGP state	State of this BGP connection.
up for	Time, in hh:mm:ss, that the underlying TCP connection has been in existence.
Last read	Time, in hh:mm:ss, since BGP last received a message from this neighbor.
hold time	Time, in seconds, that BGP will maintain the session with this neighbor without receiving messages.
keepalive interval	Time interval, in seconds, at which keepalive messages are transmitted to this neighbor.
Message statistics	Statistics organized by message type.
InQ depth is	Number of messages in the input queue.
OutQ depth is	Number of messages in the output queue.
Sent	Total number of transmitted messages.

Table 10 show bgp all neighbors Field Descriptions (continued)

Field	Description
Rcvd	Total number of received messages.
Opens	Number of open messages sent and received.
Notifications	Number of notification (error) messages sent and received.
Updates	Number of update messages sent and received.
Keepalives	Number of keepalive messages sent and received.
Route Refresh	Number of route refresh request messages sent and received.
Total	Total number of messages sent and received.
Default minimum time between...	Time, in seconds, between advertisement transmissions.
Connections established	Number of times a TCP and BGP connection has been successfully established.
dropped	Number of times that a valid session has failed or been taken down.
Last reset	Time, in hh:mm:ss, since this peering session was last reset. The reason for the reset is displayed on this line.
External BGP neighbor may be...	Indicates that the BGP Time-to-live (TTL) security check is enabled. The maximum number of hops that can separate the local and remote peer is displayed on this line.
Connection state	Connection status of the BGP peer.
Local host, Local port	IP address of the local BGP speaker and the port number.
Foreign host, Foreign port	Neighbor address and BGP destination port number.
Enqueued packets for retransmit:	Packets queued for retransmission by TCP.
Event Timers	TCP event timers. Counters are provided for starts and wakeups (expired timers).
Retrans	Number of times a packet has been retransmitted.
TimeWait	Time waiting for the retransmission timers to expire.
AckHold	Acknowledgment hold timer.
SendWnd	Transmission (send) window.
KeepAlive	Number of keepalive packets.
GiveUp	Number times a packet is dropped due to no acknowledgment.
PmtuAger	Path MTU discovery timer.
DeadWait	Expiration timer for dead segments.
iss:	Initial packet transmission sequence number.
snduna:	Last transmission sequence number that has not been acknowledged.
sndnxt:	Next packet sequence number to be transmitted.
sndwnd:	TCP window size of the remote host.
irs:	Initial packet receive sequence number.

Table 10 show bgp all neighbors Field Descriptions (continued)

Field	Description
rcvnxt:	Last receive sequence number that has been locally acknowledged.
rcvwnd:	TCP window size of the local host.
delrcvwnd:	Delayed receive window—data the local host has read from the connection, but has not yet subtracted from the receive window 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.
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:	Length of time the local host will delay an acknowledgment to carry (piggyback) additional data.
IP Precedence value:	IP precedence of the BGP packets.
Datagrams	Number of update packets received from a neighbor.
Rcvd:	Number of received packets.
with data	Number of update packets sent with data.
total data bytes	Total amount of data received, in bytes.
Sent	Number of update packets sent.
with data	Number of update packets received with data.
total data bytes	Total amount of data sent, in bytes.

Related Commands

Command	Description
router bgp	Configures the BGP routing process.

show bgp nsap

To display entries in the Border Gateway Protocol (BGP) routing table for the network service access point (NSAP) address family, use the **show bgp nsap** command in EXEC mode.

show bgp nsap [*nsap-prefix*]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast [*nsap-prefix*]

Syntax Description

unicast	Specifies NSAP unicast address prefixes.
nsap-prefix	(Optional) NSAP prefix number, entered to display a particular network in the BGP routing table for the NSAP address family. This argument may be any length up to 20 octets.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap** command provides output similar to the **show ip bgp** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap** command:

```
Router# show bgp nsap

BGP table version is 6, local router ID is 10.1.57.11
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
*> 49.0101          49.0101.1111.1111.1111.1111.00
                                                       0 65101 i
* i49.0202.2222     49.0202.3333.3333.3333.3333.00
                                              100      0 ?
*>                  49.0202.2222.2222.2222.2222.00
                                                   32768 ?
* i49.0202.3333     49.0202.3333.3333.3333.3333.00
                                              100      0 ?
*>                  49.0202.2222.2222.2222.2222.00
                                                   32768 ?
```

```
*> 49.0303          49.0303.4444.4444.4444.4444.00
                                                        0 65303 i
*  49.0404          49.0303.4444.4444.4444.4444.00
                                                        0 65303 65404 i
*>i                 49.0404.9999.9999.9999.9999.00
                                              100       0 65404 i
```

Table 11 describes the significant fields shown in the display.

Table 11 show bgp nsap Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show bgp nsap** command, showing information for NSAP prefix 49.6005.1234.4567:

```
Router# show bgp nsap 49.6005.1234.4567

BGP routing table entry for 49.6005.1234.4567, version 2
Paths: (1 available, best #1)
  Not advertised to any peer
  Local
    49.6005.1234.4567.5678.1111.2222.3333.00 from 0.0.0.0 (10.1.1.1)
      Origin IGP, localpref 100, weight 32768, valid, sourced, local, best
```

Note If a prefix has not been advertised to any peer, the display shows “Not advertised to any peer.”

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast
```

show bgp nsap community

To display routes that belong to specified network service access point (NSAP) Border Gateway Protocol (BGP) communities, use the **show bgp nsap community** command in EXEC mode.

show bgp nsap community [*community-number*] [**exact-match** | **local-as** | **no-advertise** | **no-export**]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast community [*community-number*] [**exact-match** | **local-as** | **no-advertise** | **no-export**]

Syntax Description

community-number	(Optional) Valid value is a community number in the range from 1 to 4294967295 or AA:NN (autonomous system-community number/2-byte number).
exact-match	(Optional) Displays only routes that have an exact match.
local-as	(Optional) Displays only routes that are not sent outside of the local autonomous system (well-known community).
no-advertise	(Optional) Displays only routes that are not advertised to any peer (well-known community).
no-export	(Optional) Displays only routes that are not exported outside of the local autonomous system (well-known community).
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap community** command provides output similar to the **show ip bgp community** command, except that it is specific to the NSAP address family.

Communities are set with the **route-map** and **set community** commands. Communities are sent using the **neighbor send-community** and **neighbor route-map out** commands. You must enter the numerical communities before the well-known communities. For example, the following string does not work:

```
Router> show bgp nsap community local-as 111:12345
```

Use the following string instead:

```
Router> show bgp nsap community 111:12345 local-as
```

Examples

The following is sample output from the **show bgp nsap community** command:

```
Router# show bgp nsap community no-export

BGP table version is 5, local router ID is 10.1.57.14
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
*> 49.0101.11       49.0101.2222.2222.2222.2222.00
                                                          0 101 i
```

Table 12 describes the significant fields shown in the display.

Table 12 show bgp nsap community Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast community no-export
```

Related Commands

Command	Description
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another.
set community	Sets the BGP communities attribute.
show bgp nsap community-list	Displays BGP community list information for the NSAP address family.

show bgp nsap community-list

To display routes that are permitted by the Border Gateway Protocol (BGP) community list for network service access point (NSAP) prefixes, use the **show bgp nsap community-list** command in EXEC mode.

show bgp nsap community-list *community-list-number* [**exact-match**]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast community-list *community-list-number* [**exact-match**]

Syntax Description

community-list-number	Community list number in the range from 1 to 199.
exact-match	(Optional) Displays only routes that have an exact match.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap community-list** command provides output similar to the **show ip bgp community-list** command, except that it is specific to the NSAP address family.

Examples

The following is sample output of the **show bgp nsap community-list** command:

```
Router# show bgp nsap community-list 1

BGP table version is 6, local router ID is 10.0.22.33
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
*> 49.0a0a.bb       49.0a0a.bbbb.bbbb.bbbb.bbbb.00
                                                          0 606
```

Table 13 describes the significant fields shown in the display.

Table 13 show bgp nsap community-list Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast community-list 1
```

show bgp nsap dampened-paths

Effective with Cisco IOS Release 12.2(33)SRB, the **show bgp nsap dampened-paths** command is replaced by the **show bgp nsap dampening** command. See the **show bgp nsap dampening** command for more information.

To display network service access point (NSAP) address family Border Gateway Protocol (BGP) dampened routes in the BGP routing table, use the **show bgp nsap dampened-paths** command in EXEC mode.

show bgp nsap dampened-paths

Syntax Description

This command has no arguments or keywords.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	This command was replaced by the **show bgp nsap dampening** command in Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

In Cisco IOS Release 12.2(33)SRB and later releases, the **show bgp nsap dampened-paths** command is replaced by the **show bgp nsap dampening** command. A keyword, **dampened-paths**, can be used with the new **show bgp nsap dampened-paths** command to display NSAP address family BGP dampened routes.

Examples

The following is sample output from the **show bgp nsap dampened-paths** command in privileged EXEC mode:

```
Router# show bgp nsap dampened-paths

BGP table version is 20, local router ID is 10.1.57.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
*d 49.0404          10.2.4.2         00:25:50 65202 65404 i
```

Table 14 describes the significant fields shown in the display.

Table 14 show bgp nsap dampened-paths Field Descriptions

Field	Description
BGP table version	Internal version number for the table. This number is incremented whenever the table changes.
local router ID	IP address of the router where route dampening is enabled.
*d	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system path of the route that is being dampened.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear bgp nsap dampening	Clears BGP NSAP prefix route dampening information and unsuppresses the suppressed routes.

show bgp nsap dampening

To display network service access point (NSAP) address family Border Gateway Protocol (BGP) dampened routes in the BGP routing table, use the **show bgp nsap dampening** command in user EXEC or privileged EXEC mode.

show bgp nsap unicast dampening {**dampened-paths** | **flap-statistics** [**regexp** *regexp* | **quote-regexp** *regexp* | **filter-list** *access-list-number* | *nsap-prefix*] | **parameters**}

Syntax Description

unicast	Specifies NSAP unicast address prefixes.
dampened-paths	Displays paths suppressed due to dampening.
flap-statistics	Displays flap statistics of routes.
regexp *regexp*	(Optional) Displays flap statistics for all the paths that match the regular expression.
quote-regexp *regexp*	(Optional) Displays flap statistics for all the paths that match the regular expression as a quoted string of characters.
filter-list *access-list-number*	(Optional) Displays flap statistics for all the paths that pass the access list.
nsap-prefix	(Optional) Displays flap statistics for a single entry at this NSAP network number.
parameters	Displays details of configured dampening parameters.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(33)SRB	This command was introduced.

Examples

The following is sample output from the **show bgp nsap dampened-paths** command in privileged EXEC mode:

```
Router# show bgp nsap unicast dampening dampened-paths

BGP table version is 20, local router ID is 10.1.57.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
*d 49.0404          10.2.4.2         00:25:50 65202 65404 i
```

Table 15 describes the significant fields shown in the display.

Table 15 show bgp nsap unicast dampening dampened-paths Field Descriptions

Field	Description
BGP table version	Internal version number for the table. This number is incremented whenever the table changes.
local router ID	IP address of the router where route dampening is enabled.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system path of the route that is being dampened.

The following is sample output from the **show bgp nsap unicast dampening flap-statistics** command:

```
Router# show bgp nsap unicast dampening flap-statistics

BGP table version is 20, local router ID is 10.1.57.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From            Flaps Duration Reuse    Path
*d 49.0404          10.2.4.2        3     00:09:45 00:23:40 65202 65404
```

Table 16 describes the significant fields shown in the display.

Table 16 show bgp nsap unicast dampening flap-statistics Field Descriptions

Field	Description
BGP table version	Internal version number for the table. This number is incremented whenever the table changes.
local router ID	IP address of the router where route dampening is enabled.

Table 16 show bgp nsap unicast dampening flap-statistics Field Descriptions

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the router noticed the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	Autonomous system path of the route that is being dampened.

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear bgp nsap dampening	Clears BGP NSAP prefix route dampening information and unsuppresses the suppressed routes.

show bgp nsap filter-list

To display routes in the Border Gateway Protocol (BGP) routing table for the network service access point (NSAP) address family that conform to a specified filter list, use the **show bgp nsap filter-list** command in privileged EXEC mode.

show bgp nsap filter-list *access-list-number*

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast filter-list *access-list-number*

Syntax Description

access-list-number	Number of an autonomous system path access list. It can be a number from 1 to 199.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Examples

The following is sample output from the **show bgp nsap filter-list** command:

```
Router# show bgp nsap filter-list 1

BGP table version is 3, local router ID is 10.0.11.33
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
*> 49.0b0b          49.0b0b.bbbb.bbbb.bbbb.bbbb.00
                                                          0 707 i
```

Table 17 describes the significant fields shown in the display.

Table 17 show bgp nsap filter-list Field Descriptions

Field	Description
BGP table version	Internal version number for the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Table 17 show bgp nsap filter-list Field Descriptions (continued)

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Set through the use of autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast filter-list 1
```

show bgp nsap flap-statistics

To display Border Gateway Protocol (BGP) flap statistics for network service access point (NSAP) prefixes, use the **show bgp nsap flap-statistics** command in EXEC mode.

show bgp nsap flap-statistics [**regexp** *regexp* | **quote-regexp** *regexp* | **filter-list** *access-list-number* | *nsap-prefix*]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast flap-statistics [**regexp** *regexp* | **quote-regexp** *regexp* | **filter-list** *access-list-number* | *nsap-prefix*]

Syntax Description

regexp *regexp*	(Optional) Displays flap statistics for all the paths that match the regular expression.
quote-regexp *regexp*	(Optional) Displays flap statistics for all the paths that match the regular expression as a quoted string of characters.
filter-list *access-list-number*	(Optional) Displays flap statistics for all the paths that pass the access list.
nsap-prefix	(Optional) Displays flap statistics for a single entry at this NSAP network number.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap flap-statistics** command provides output similar to the **show ip bgp flap-statistics** command, except that it is specific to the NSAP address family.

If no arguments or keywords are specified, the router displays flap statistics for all NSAP prefix routes.

Examples

The following is sample output from the **show bgp nsap flap-statistics** command without arguments or keywords:

```
Router# show bgp nsap flap-statistics

BGP table version is 20, local router ID is 10.1.57.13
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          From            Flaps Duration Reuse    Path
*d 49.0404          10.2.4.2        3     00:09:45 00:23:40 65202 65404
```

Table 18 describes the significant fields shown in the display.

Table 18 show bgp nsap flap-statistics Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	Route to the network indicated is dampened.
From	IP address of the peer that advertised this path.
Flaps	Number of times the route has flapped.
Duration	Time (in hours:minutes:seconds) since the router noticed the first flap.
Reuse	Time (in hours:minutes:seconds) after which the path will be made available.
Path	AS-path of the route that is being dampened.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast flap-statistics
```

Related Commands

Command	Description
bgp dampening	Enables BGP route dampening or changes various BGP route dampening factors.
clear bgp nsap flap-statistics	Clears BGP flap statistics for NSAP prefix routes.

show bgp nsap inconsistent-as

To display Border Gateway Protocol (BGP) network service access point (NSAP) prefix routes with inconsistent originating autonomous systems, use the **show bgp nsap inconsistent-as** command in EXEC mode.

show bgp nsap inconsistent-as

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast inconsistent-as

Syntax Description

unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap inconsistent-as** command provides output similar to the **show ip bgp inconsistent-as** command, except that it is specific to the NSAP address family.

Use the **show bgp nsap inconsistent-as** command to discover any BGP routing table entries that contain inconsistent autonomous system path information. Inconsistent autonomous path information is useful for troubleshooting networks because it highlights a configuration error in the network.

Examples

The following is sample output from the **show bgp nsap inconsistent-as** command. In this example, the network prefix of 49.0a0a has two entries in the BGP routing table showing different originating paths. The originating path information should be the same in both entries.

```
Router# show bgp nsap inconsistent-as

BGP table version is 3, local router ID is 10.1.57.17
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
*  49.0a0a          49.0a0a.cccc.cccc.cccc.00
                                                          0 30 i
*>                  49.0a0a.aaaa.aaaa.aaaa.00
                                                          0 10 i
```

Table 19 describes the significant fields shown in the display.

Table 19 show bgp nsap inconsistent-as Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast inconsistent-as
```

show bgp nsap neighbors

To display information about Border Gateway Protocol (BGP) network service access point (NSAP) prefix connections to neighbors, use the **show bgp nsap neighbors** command in EXEC mode.

show bgp nsap neighbors [*ip-address* [**routes** | **flap-statistics** | **advertised-routes** | **paths** *regexp* | **dampened-routes**]]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast neighbors [*ip-address* [**routes** | **flap-statistics** | **advertised-routes** | **paths** *regexp* | **dampened-routes**]]

Syntax Description

ip-address	(Optional) IP address of the BGP-speaking neighbor. If you omit this argument, all neighbors are displayed.
routes	(Optional) Displays all routes received and accepted.
flap-statistics	(Optional) Displays flap statistics for the routes learned from the neighbor.
advertised-routes	(Optional) Displays all the routes the networking device advertised to the neighbor.
paths *regexp*	(Optional) Regular expression used to match the paths received.
dampened-routes	(Optional) Displays the dampened routes to the neighbor at the NSAP prefix address specified.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap neighbors** command provides output similar to the **show ip bgp neighbors** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap neighbors** command:

```
Router# show bgp nsap neighbors 10.0.2.3

BGP neighbor is 10.0.2.3,  remote AS 64500, external link
  BGP version 4, remote router ID 172.17.1.2
  BGP state = Established, up for 00:12:50
  Last read 00:00:50, hold time is 180, keepalive interval is 60 seconds
```

```
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Address family NSAP Unicast: advertised and received
  Received 17 messages, 0 notifications, 0 in queue
  Sent 17 messages, 0 notifications, 0 in queue
  Route refresh request: received 0, sent 0
  Default minimum time between advertisement runs is 30 seconds

 For address family: NSAP Unicast
  BGP table version 5, neighbor version 5
  Index 2, Offset 0, Mask 0x4
  2 accepted prefixes consume 114 bytes
  Prefix advertised 2, suppressed 0, withdrawn 0
  Number of NLRIs in the update sent: max 1, min 0

  Connections established 1; dropped 0
  Last reset never
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Local host: 10.0.2.2, Local port: 11000
Foreign host: 10.0.2.3, Foreign port: 179

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

Event Timers (current time is 0x115940):
Timer          Starts    Wakeups            Next
Retrans            22          1             0x0
TimeWait            0          0             0x0
AckHold            19          7             0x0
SendWnd             0          0             0x0
KeepAlive           0          0             0x0
GiveUp              0          0             0x0
PmtuAger            0          0             0x0
DeadWait            0          0             0x0

iss: 2052706884  snduna: 2052707371  sndnxt: 2052707371     sndwnd:  15898
irs: 1625021348  rcvnxt: 1625021835  rcvwnd:      15898  delrcvwnd:    486

SRTT: 279 ms, RTTO: 446 ms, RTV: 167 ms, KRTT: 0 ms
minRTT: 0 ms, maxRTT: 300 ms, ACK hold: 200 ms
Flags: higher precedence, nagle

Datagrams (max data segment is 1460 bytes):
Rcvd: 30 (out of order: 0), with data: 19, total data bytes: 486
Sent: 29 (retransmit: 1, fastretransmit: 0), with data: 20, total data bytes: 46
```

Table 20 describes the significant fields shown in the display.

Table 20 show bgp nsap neighbors Field Descriptions

Field	Description
BGP neighbor	IP address of the BGP neighbor and its autonomous system number.
remote AS	Autonomous system of the neighbor.
link	If the neighbor is in the same autonomous system as the router, then the link between them is internal; otherwise, it is considered external.
BGP version	BGP version being used to communicate with the remote router; the router ID (an IP address) of the neighbor is also specified.
remote router ID	A 32-bit number written as 4 octets separated by periods (dotted decimal format).

Table 20 show bgp nsap neighbors Field Descriptions (continued)

Field	Description
BGP state	Internal state of this BGP connection.
up for	Amount of time (in hours:minutes:seconds) that the underlying TCP connection has been in existence.
Last read	Time (in hours:minutes:seconds) that BGP last read a message from this neighbor.
hold time	Maximum amount of time, in seconds, that can elapse between messages from the peer.
keepalive interval	Time period, in seconds, between sending keepalive packets, which help ensure that the TCP connection is up.
Neighbor capabilities	BGP capabilities advertised and received from this neighbor.
Route refresh	Indicates that the neighbor supports dynamic soft reset using the route refresh capability.
Address family NSAP Unicast	NSAP unicast-specific properties of this neighbor.
Received	Number of total BGP messages received from this peer, including keepalives.
notifications	Number of error messages received from the peer.
Sent	Total number of BGP messages that have been sent to this peer, including keepalives.
notifications	Number of error messages the router has sent to this peer.
Route refresh request	Number of route refresh requests sent and received from this neighbor.
advertisement runs	Value of minimum advertisement interval.
For address family	Address family to which the following fields refer.
BGP table version	Indicates that the neighbor has been updated with this version of the primary BGP routing table.
neighbor version	Number used by the software to track the prefixes that have been sent and those that must be sent to this neighbor.
Community attribute (not shown in sample output)	Appears if the **neighbor send-community** command is configured for this neighbor.
Inbound path policy (not shown in sample output)	Indicates that an inbound filter list or route map is configured.
Outbound path policy (not shown in sample output)	Indicates that an outbound filter list, route map, or unsuppress map is configured.
bgp-in (not shown in sample output)	Name of the inbound update prefix filter list for the NSAP unicast address family.
aggregate (not shown in sample output)	Name of the outbound update prefix filter list for the NSAP unicast address family.
uni-out (not shown in sample output)	Name of the outbound route map for the NSAP unicast address family.

Table 20 show bgp nsap neighbors Field Descriptions (continued)

Field	Description
accepted prefixes	Number of prefixes accepted.
Prefix advertised	Number of prefixes advertised.
suppressed	Number of prefixes suppressed.
withdrawn	Number of prefixes withdrawn.
history paths (not shown in sample output)	Number of path entries held to remember history.
Connections established	Number of times the router has established a TCP connection and the two peers have agreed to speak BGP with each other.
dropped	Number of times that a good connection has failed or been taken down.
Last reset	Elapsed time since this peering session was last reset.
Connection state	State of the BGP peer.
unread input bytes	Number of bytes of packets still to be processed.
Local host, Local port	Peering address of local router, plus port.
Foreign host, Foreign port	Peering address of the neighbor.
Event Timers	Table that displays the number of starts and wakeups for each timer.
iss	Initial send sequence number.
snduna	Last send sequence number the local host sent but for which it has not received an acknowledgment.
sndnxt	Sequence number the local host will send next.
sndwnd	TCP window size of the remote host.
irs	Initial receive sequence number.
rcvnxt	Last receive sequence number the local host has acknowledged.
rcvwnd	TCP window size of the local host.
delrcvwnd	Delayed receive window—data the local host has read from the connection but has not yet subtracted from the receive window 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.
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 (in milliseconds) the local host will delay an acknowledgment in order to “piggyback” data on it.
Flags	IP precedence of the BGP packets.

Table 20 show bgp nsap neighbors Field Descriptions (continued)

Field	Description
Datagrams: Rcvd	Number of update packets received from neighbor.
with data	Number of update packets received with data.
total data bytes	Total bytes of data.
Sent	Number of update packets sent.
with data	Number of update packets with data sent.
total data bytes	Total number of data bytes.

The following is sample output from the **show bgp nsap neighbors** command with the **advertised-routes** keyword:

```
Router# show bgp nsap neighbors 10.0.2.3 advertised-routes

BGP table version is 5, local router ID is 172.17.1.2
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
*> 49.0101          49.0101.1111.1111.1111.1111.00
                                                           0 101 i
*> 49.0202          49.0202.2222.2222.2222.2222.00
                                                           32768 i
```

The following is sample output from the **show bgp nsap neighbors** command with the **routes** keyword:

```
Router# show bgp nsap neighbors 10.0.2.3 routes

BGP table version is 5, local router ID is 172.17.1.2
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
*> 49.0303          49.0303.3333.3333.3333.3333.00
                                                           0 303 i
*> 49.0404          49.0303.3333.3333.3333.3333.00
                                                           0 303 404 i

Total number of prefixes 2
```

Table 21 describes the significant fields shown in the display.

Table 21 show bgp nsap neighbors Field Descriptions with advertised-routes and routes keywords

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Table 21 show bgp nsap neighbors Field Descriptions with advertised-routes and routes keywords (continued)

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show bgp nsap neighbors** command with the **paths** keyword:

```
Router# show bgp nsap neighbors 10.0.3.3 paths ^101

Address    Refcount Metric Path
0x62281590        1      0 101 i
```

Note The caret (^) symbol in the example is a regular expression that is entered by simultaneously pressing the Shift and 6 keys on your keyboard. A caret (^) symbol at the beginning of a regular expression matches the start of a line.

Table 22 describes the significant fields shown in the display.

Table 22 show bgp nsap neighbors paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Refcount	Number of routes using that path.
Metric	The Multiple Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The AS-path for that route, followed by the origin code for that route.

The following sample output from the **show bgp nsap neighbors** command shows the NSAP prefix dampened routes for the neighbor at 10.0.2.2:

```
Router# show bgp nsap neighbors 10.0.2.2 dampened-routes

BGP table version is 10, local router ID is 172.17.1.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From             Reuse    Path
*d 49.0101          10.0.2.2         00:25:50 202 101 i
```

The following sample output from the **show bgp nsap neighbors** command shows the NSAP prefix flap statistics for the neighbor at 10.0.2.2:

```
Router# show bgp nsap neighbors 10.0.2.2 flap-statistics

BGP table version is 10, local router ID is 10.1.57.14
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          From            Flaps Duration Reuse    Path
*d 49.0101          10.0.2.2        3     00:07:00 00:24:50 202 101
```

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast neighbors 10.0.2.3
```

Related Commands

Command	Description
neighbor activate	Enables the exchange of information with a neighboring router.

show bgp nsap paths

To display all the Border Gateway Protocol (BGP) network service access point (NSAP) prefix paths in the database, use the **show bgp nsap paths** command in EXEC mode.

show bgp nsap paths [*AS-path-regexp*]

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast paths [*AS-path-regexp*]

Syntax Description

AS-path-regexp	(Optional) Regular expression that is used to match the received paths in the database.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap paths** command provides output similar to the **show ip bgp paths** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap paths** command without a specified regular expression:

```
Router# show bgp nsap paths

Address    Hash Refcount Metric Path
0x622803FC    0        1      0 i
0x62280364 1197        1      0 202 101 i
0x62280448 1739        1      0 202 i
0x622803B0 1941        1      0 404 i
```

Table 23 describes the significant fields shown in the display.

Table 23 show bgp nsap paths Field Descriptions

Field	Description
Address	Internal address where the path is stored.
Hash	Hash bucket where the path is stored.

Table 23 show bgp nsap paths Field Descriptions (continued)

Field	Description
Refcount	Number of routes using that path.
Metric	The Multiple Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
Path	The AS-path for that route, followed by the origin code for that route.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast paths
```

show bgp nsap quote-regexp

To display Border Gateway Protocol (BGP) network service access point (NSAP) prefix routes matching the AS-path regular expression as a quoted string of characters, use the **show bgp nsap quote-regexp** command in privileged EXEC mode.

show bgp nsap quote-regexp *as-path-regexp*

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast quote-regexp *as-path-regexp*

Syntax Description

as-path-regexp	Regular expression to match the BGP autonomous system paths. The regular expression is contained within quotes.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap quote-regexp** command provides output similar to the **show ip bgp quote-regexp** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap quote-regexp** command that shows paths equal to 202:

```
Router# show bgp nsap quote-regexp "202"

BGP table version is 10, local router ID is 10.1.57.14
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
*d 49.0101          49.0202.2222.2222.2222.2222.00
                                                           0 202 101 i
*> 49.0202          49.0202.2222.2222.2222.2222.00
                                                           0 202 i
```

Table 24 describes the significant fields shown in the display.

Table 24 show bgp nsap quote-regexp Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast quote-regexp "202"
```

Related Commands

Command	Description
show bgp nsap regexp	Displays NSAP prefix routes matching the AS-path regular expression.

show bgp nsap regexp

To display Border Gateway Protocol (BGP) network service access point (NSAP) prefix routes matching the AS-path regular expression, use the **show bgp nsap regexp** command in privileged EXEC mode.

show bgp nsap regexp *AS-path-regexp*

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast regexp *AS-path-regexp*

Syntax Description

AS-path-regexp	Regular expression to match the BGP autonomous system paths.
unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap regexp** command provides output similar to the **show ip bgp regexp** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap regexp** command that shows paths beginning with 202 or containing 101:

```
Router# show bgp nsap regexp ^202 101

BGP table version is 10, local router ID is 10.1.57.14
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
*d 49.0101          49.0202.2222.2222.2222.2222.00
                                                          0 202 101 i
```

Note The caret (^) symbol in the example is a regular expression that is entered by simultaneously pressing the Shift and 6 keys on your keyboard. A caret (^) symbol at the beginning of a regular expression matches the start of a line.

Table 25 describes the significant fields shown in the display.

Table 25 show bgp nsap regexp Field Descriptions

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry is history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from an Interior Gateway Protocol (IGP) and was advertised with a **network** router configuration command. e—Entry originated from an Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a route that is redistributed into BGP from an IGP.
Network	NSAP prefix address of a network entity.
Next Hop	CLNS network entity title (NET) consisting of area address and system ID of the next system that is used when forwarding a packet to the destination network. This entry may cause a line break with the values of the following fields being displayed on the next line under their respective titles.
Metric	If shown, the value of the interautonomous system metric. This field is frequently not used.
LocPrf	Local preference value as set with the **set local-preference** route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast regexp ^202 101
```

Related Commands

Command	Description
show bgp nsap quote-regexp	Displays BGP NSAP prefix routes matching the AS-path regular expression.

show bgp nsap summary

To display the status of all Border Gateway Protocol (BGP) network service access point (NSAP) prefix connections, use the **show bgp nsap summary** command in EXEC mode.

show bgp nsap summary

Syntax in Cisco IOS Release 12.2(33)SRB

show bgp nsap unicast summary

Syntax Description

unicast	Specifies NSAP unicast address prefixes.

Command Modes

User EXEC (>)
Privileged EXEC (#)

Command History

Release	Modification
12.2(8)T	This command was introduced.
12.2(33)SRB	The **unicast** keyword was added and this command was integrated into Cisco IOS Release 12.2(33)SRB.
Cisco IOS XE 2.6	This command was integrated into Cisco IOS XE Release 2.6.

Usage Guidelines

The **show bgp nsap summary** command provides output similar to the **show ip bgp summary** command, except that it is specific to the NSAP address family.

Examples

The following is sample output from the **show bgp nsap summary** command:

```
Router# show bgp nsap summary

BGP router identifier 10.2.4.2, local AS number 65202
BGP table version is 26, main routing table version 26
5 network entries and 8 paths using 1141 bytes of memory
6 BGP path attribute entries using 360 bytes of memory
4 BGP AS-PATH entries using 96 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
Dampening enabled. 0 history paths, 0 dampened paths
BGP activity 16/261 prefixes, 34/26 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.2.1        4 65101    1162    1162       26    0    0 18:17:07        1
10.2.3.3        4 65202    1183    1188       26    0    0 18:23:28        3
10.2.4.4        4 65303    1163    1187       26    0    0 18:23:14        2
```

Table 26 describes the significant fields shown in the display.

Table 26 show bgp nsap summary Field Descriptions

Field	Description
BGP router identifier	IP address of the networking device.
local AS number	Number of the local autonomous system.
BGP table version	Internal version number of the BGP database.
main routing table version	Last version of the BGP database that was injected into the main routing table.
network entries	Number of network entries and paths in the main routing table including the associated memory usage.
BGP path attribute entries	Number of BGP path attribute entries in the main routing table including the associated memory usage.
BGP route-map cache entries	Number of BGP route map cache entries in the main routing table including the associated memory usage.
BGP filter-list cache entries	Number of BGP filter list cache entries in the main routing table including the associated memory usage.
Dampening	Indicates whether route dampening is enabled, the number of history paths, and number of dampened paths.
BGP activity	Displays the number of BGP prefixes and paths, followed by the BGP scan interval in seconds.
Neighbor	IP address of a neighbor.
V	BGP version number communicated to that neighbor.
AS	Autonomous system.
MsgRcvd	BGP messages received from that neighbor.
MsgSent	BGP messages sent to that neighbor.
TblVer	Last version of the BGP database that was sent to that neighbor.
InQ	Number of messages from that neighbor waiting to be processed.
OutQ	Number of messages waiting to be sent to that neighbor.
Up/Down	The length of time that the BGP session has been in state Established, or the current state if it is not Established.
State/PfxRcd	Current state of the BGP session/the number of prefixes the router has received from a neighbor or peer group. When the maximum number (as set by the **neighbor maximum-prefix** command) is reached, the string "PfxRcd" appears in the entry, the neighbor is shut down, and the connection is Idle. An (Admin) entry with Idle status indicates that the connection has been shut down using the **neighbor shutdown** command.

In this example for Cisco IOS Release 12.2(33)SRB there is a new keyword, **unicast**, that is required. The output for the following command is the same as in the first example.

```
Router# show bgp nsap unicast summary
```

Related Commands

Command	Description
clear bgp nsap	Resets an NSAP BGP TCP connection.
neighbor maximum-prefix	Controls how many prefixes can be received from a neighbor.
neighbor shutdown	Disables a neighbor or peer group.