



# Configuring Advanced BGP

---

This chapter contains the following sections:

- [About Advanced BGP, on page 2](#)
- [Prerequisites for Advanced BGP, on page 14](#)
- [Guidelines and Limitations for Advanced BGP, on page 14](#)
- [Default Settings, on page 17](#)
- [Configuring Advanced BGP, on page 17](#)
- [Configuring BGP Additional Paths, on page 33](#)
- [Configuring eBGP, on page 37](#)
- [Configuring AS Confederations, on page 39](#)
- [Configuring Route Reflector, on page 40](#)
- [Configuring Next-Hops on Reflected Routes Using an Outbound Route-Map, on page 41](#)
- [Configuring Route Dampening, on page 44](#)
- [Configuring Load Sharing and ECMP, on page 44](#)
- [Configuring Maximum Prefixes, on page 45](#)
- [Configuring DSCP, on page 45](#)
- [Configuring Dynamic Capability, on page 46](#)
- [Configuring Aggregate Addresses, on page 46](#)
- [Suppressing BGP Routes, on page 47](#)
- [Configuring BGP Conditional Advertisement, on page 48](#)
- [Configuring Route Redistribution, on page 50](#)
- [Advertising the Default Route, on page 51](#)
- [Configuring BGP Attribute Filtering and Error Handling, on page 52](#)
- [Tuning BGP, on page 55](#)
- [Configuring Policy-Based Administrative Distance, on page 60](#)
- [Configuring Multiprotocol BGP, on page 61](#)
- [Configuring BMP, on page 62](#)
- [BGP Local Route Leaking, on page 64](#)
- [BGP Graceful Shutdown, on page 72](#)
- [Configuring a Graceful Restart, on page 83](#)
- [Configuring Virtualization, on page 85](#)
- [Verifying the Advanced BGP Configuration, on page 86](#)
- [Monitoring BGP Statistics, on page 88](#)
- [Configuration Examples, on page 89](#)

- [Related Topics, on page 90](#)
- [Additional References, on page 90](#)

## About Advanced BGP

BGP is an interdomain routing protocol that provides loop-free routing between organizations or autonomous systems. Cisco NX-OS supports BGP version 4. BGP version 4 includes multiprotocol extensions that allow BGP to carry routing information for IP multicast routes and multiple Layer 3 protocol address families. BGP uses TCP as a reliable transport protocol to create TCP sessions with other BGP-enabled devices called BGP peers. When connecting to an external organization, the router creates external BGP (eBGP) peering sessions. BGP peers within the same organization exchange routing information through internal BGP (iBGP) peering sessions.

## Peer Templates

BGP peer templates allow you to create blocks of common configuration that you can reuse across similar BGP peers. Each block allows you to define a set of attributes that a peer then inherits. You can choose to override some of the inherited attributes as well, making it a very flexible scheme for simplifying the repetitive nature of BGP configurations.

Cisco NX-OS implements three types of peer templates:

- The peer-session template defines BGP peer session attributes, such as the transport details, remote autonomous system number of the peer, and session timers. A peer-session template can also inherit attributes from another peer-session template (with locally defined attributes that override the attributes from an inherited peer-session).
- A peer-policy template defines the address-family dependent policy aspects for a peer including the inbound and outbound policy, filter-lists, and prefix-lists. A peer-policy template can inherit from a set of peer-policy templates. Cisco NX-OS evaluates these peer-policy templates in the order specified by the preference value in the inherit configuration. The lowest number is preferred over higher numbers.
- The peer template can inherit the peer-session and peer-policy templates to allow for simplified peer definitions. It is not mandatory to use a peer template but it can simplify the BGP configuration by providing reusable blocks of configuration.

## Authentication

You can configure authentication for a BGP neighbor session. This authentication method adds an MD5 authentication digest to each TCP segment sent to the neighbor to protect BGP against unauthorized messages and TCP security attacks.



---

**Note** The MD5 password must be identical between BGP peers.

---

## Route Policies and Resetting BGP Sessions

You can associate a route policy to a BGP peer. Route policies use route maps to control or modify the routes that BGP recognizes. You can configure a route policy for inbound or outbound route updates. The route policies can match on different criteria, such as a prefix or AS\_path attribute, and selectively accept or deny the routes. Route policies can also modify the path attributes.

When you change a route policy applied to a BGP peer, you must reset the BGP sessions for that peer. Cisco NX-OS supports the following three mechanisms to reset BGP peering sessions:

- **Hard reset**—A hard reset tears down the specified peering sessions, including the TCP connection, and deletes routes coming from the specified peer. This option interrupts packet flow through the BGP network. Hard reset is disabled by default.
- **Soft reconfiguration inbound**—A soft reconfiguration inbound triggers routing updates for the specified peer without resetting the session. You can use this option if you change an inbound route policy. Soft reconfiguration inbound saves a copy of all routes received from the peer before processing the routes through the inbound route policy. If you change the inbound route policy, Cisco NX-OS passes these stored routes through the modified inbound route policy to update the route table without tearing down existing peering sessions. Soft reconfiguration inbound can use significant memory resources to store the unfiltered BGP routes. Soft reconfiguration inbound is disabled by default.
- **Route Refresh**—A route refresh updates the inbound routing tables dynamically by sending route refresh requests to supporting peers when you change an inbound route policy. The remote BGP peer responds with a new copy of its routes that the local BGP speaker processes with the modified route policy. Cisco NX-OS automatically sends an outbound route refresh of prefixes to the peer.
- BGP peers advertise the route refresh capability as part of the BGP capability negotiation when establishing the BGP peer session. Route refresh is the preferred option and enabled by default.



---

**Note** BGP also uses route maps for route redistribution, route aggregation, route dampening, and other features. See [Configuring Route Policy Manager](#), for more information on route maps.

---

## eBGP

External BGP (eBGP) allows you to connect BGP peers from different autonomous systems to exchange routing updates. Connecting to external networks enables traffic from your network to be forwarded to other networks and across the Internet.


Typically eBGP peerings need to be over directly connected interfaces so that convergence will be faster when the interface goes down.

## iBGP

Internal BGP (iBGP) allows you to connect BGP peers within the same autonomous system. You can use iBGP for multihomed BGP networks (networks that have more than one connection to the same external autonomous system).

The figure shows an iBGP network within a larger BGP network.

Figure 1: iBGP Network



iBGP networks are fully meshed. Each iBGP peer has a direct connection to all other iBGP peers to prevent network loops.

For single-hop iBGP peers with update-source configured under neighbor configuration mode, the peer supports fast external fall-over.

You should use loopback interfaces for establishing iBGP peering sessions because loopback interfaces are less susceptible to interface flapping. An interface flap occurs when the interface is administratively brought up or down because of a failure or maintenance issue. See the [Configuring eBGP, on page 37](#) section for information on multihop, fast external fallovers, and limiting the size of the AS\_path attribute.




**Note** You should configure a separate interior gateway protocol in the iBGP network.

## AS Confederations

A fully meshed iBGP network becomes complex as the number of iBGP peers grows. You can reduce the iBGP mesh by dividing the autonomous system into multiple subautonomous systems and grouping them into a single confederation. A confederation is a group of iBGP peers that use the same autonomous system number to communicate to external networks. Each subautonomous system is fully meshed within itself and has a few connections to other subautonomous systems in the same confederation.

The figure shows the BGP network, split into two subautonomous systems and one confederation.

Figure 2: AS Confederation



In this example, AS10 is split into two subautonomous systems, AS1 and AS2. Each subautonomous system is fully meshed, but there is only one link between the subautonomous systems. By using AS confederations, you can reduce the number of links compared to the fully meshed autonomous system.

## Route Reflector


You can alternately reduce the iBGP mesh by using a route reflector configuration where route reflectors pass learned routes to neighbors so that all iBGP peers do not need to be fully meshed.

When you configure an iBGP peer to be a route reflector, it becomes responsible for passing iBGP learned routes to a set of iBGP neighbors.

The figure shows a simple iBGP configuration with four meshed iBGP speakers (routers A, B, C, and D). Without route reflectors, when router A receives a route from an external neighbor, it advertises the route to all three iBGP neighbors.

In the figure, router B is the route reflector. When the route reflector receives routes advertised from router A, it advertises (reflects) the routes to routers C and D. Router A no longer has to advertise to both routers C and D.

**Figure 3: Route Reflector**



The route reflector and its client peers form a cluster. You do not have to configure all iBGP peers to act as client peers of the route reflector. You must configure any nonclient peer as fully meshed to guarantee that complete BGP updates reach all peers.

## Capabilities Negotiation

A BGP speaker can learn about BGP extensions that are supported by a peer by using the capabilities negotiation feature. Capabilities negotiation allows BGP to use only the set of features supported by both BGP peers on a link.

If a BGP peer does not support capabilities negotiation, Cisco NX-OS attempts a new session to the peer without capabilities negotiation if you have configured the address family as IPv4. Any other multiprotocol configuration (such as IPv6) requires capabilities negotiation.

## Route Dampening

Route dampening is a BGP feature that minimizes the propagation of flapping routes across an internetwork. A route flaps when it alternates between the available and unavailable states in rapid succession.

For example, consider a network with three BGP autonomous systems: AS1, AS2, and AS3. Suppose that a route in AS1 flaps (it becomes unavailable). Without route dampening, AS1 sends a withdraw message to AS2. AS2 propagates the withdrawal message to AS3. When the flapping route reappears, AS1 sends an advertisement message to AS2, which sends the advertisement to AS3. If the route repeatedly becomes unavailable, and then available, AS1 sends many withdrawal and advertisement messages that propagate through the other autonomous systems.

Route dampening can minimize flapping. Suppose that the route flaps. AS2 (in which route dampening is enabled) assigns the route a penalty of 1000. AS2 continues to advertise the status of the route to neighbors. Each time that the route flaps, AS2 adds to the penalty value. When the route flaps so often that the penalty exceeds a configurable suppression limit, AS2 stops advertising the route, regardless of how many times that it flaps. The route is now dampened.

The penalty placed on the route decays until the reuse limit is reached. At that time, AS2 advertises the route again. When the reuse limit is at 50 percent, AS2 removes the dampening information for the route.




---

**Note** The router does not apply a penalty to a resetting BGP peer when route dampening is enabled, even though the peer reset withdraws the route.

---

## Load Sharing and Multipath

BGP can install multiple equal-cost eBGP or iBGP paths into the routing table to reach the same destination prefix. Traffic to the destination prefix is then shared across all the installed paths.

The BGP best-path algorithm considers the paths as equal-cost paths if the following attributes are identical:

- Weight
- Local preference
- AS\_path
- Origin code
- Multi-exit discriminator (MED)
- IGP cost to the BGP next hop

BGP selects only one of these multiple paths as the best path and advertises the path to the BGP peers. For more information, see the [BGP Additional Paths](#) section.




---

**Note** Paths that are received from different AS confederations are considered as equal-cost paths if the external AS\_path values and the other attributes are identical.

---




---

**Note** When you configure a route reflector for iBGP multipath, and the route reflector advertises the selected best path to its peers, the next hop for the path is not modified.

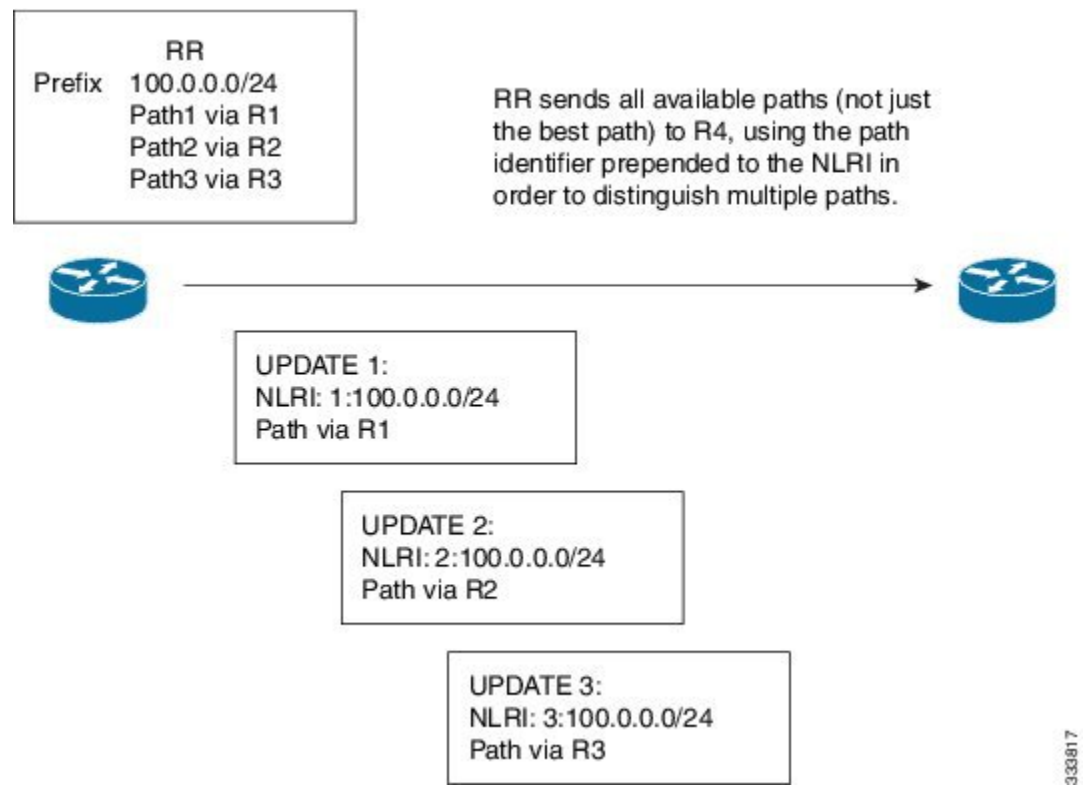
---

## BGP Additional Paths

Only one BGP best path is advertised, and the BGP speaker accepts only one path for a given prefix from a given peer. If a BGP speaker receives multiple paths for the same prefix within the same session, it uses the most recent advertisement.

BGP supports the additional paths feature, which allows the BGP speaker to propagate and accept multiple paths for the same prefix without the new paths replacing any previous ones. This feature allows BGP speaker peers to negotiate whether they support advertising and receiving multiple paths per prefix and advertising such paths. A special 4-byte path ID is added to the network layer reachability information (NLRI) to differentiate multiple paths for the same prefix sent across a peer session. The following figure illustrates the BGP additional paths capability.

**Figure 4: BGP Route Advertisement with the Additional Paths Capability**



For information on configuring BGP additional paths, see the [Configuring BGP Additional Paths](#), on page 33 section.

## Route Aggregation

You can configure aggregate addresses. Route aggregation simplifies route tables by replacing a number of more specific addresses with an address that represents all the specific addresses. For example, you can replace these three more specific addresses, 10.1.1.0/24, 10.1.2.0/24, and 10.1.3.0/24 with one aggregate address, 10.1.0.0/16.

Aggregate prefixes are present in the BGP route table so that fewer routes are advertised.



**Note** Cisco NX-OS does not support automatic route aggregation.

Route aggregation can lead to forwarding loops. To avoid this problem, when BGP generates an advertisement for an aggregate address, it automatically installs a summary discard route for that aggregate address in the



local routing table. BGP sets the administrative distance of the summary discard to 220 and sets the route type to discard. BGP does not use discard routes for next-hop resolution.

A summary entry is created in the BGP table when you issue the **aggregate-address** command, but the summary entry is not eligible for advertisement until a subset of the aggregate is found in the table.

## BGP Conditional Advertisement

BGP conditional advertisement allows you to configure BGP to advertise or withdraw a route based on whether or not a prefix exists in the BGP table. This feature is useful, for example, in multihomed networks, in which you want BGP to advertise some prefixes to one of the providers only if information from the other provider is not present.

Consider an example network with three BGP autonomous systems: AS1, AS2, and AS3, where AS1 and AS3 connect to the Internet and to AS2. Without conditional advertisement, AS2 propagates all routes to both AS1 and AS3. With conditional advertisement, you can configure AS2 to advertise certain routes to AS3 only if routes from AS1 do not exist (if for example, the link to AS1 fails).

BGP conditional advertisement adds an exist or not-exist test to each route that matches the configured route map. See the [Configuring BGP Conditional Advertisement](#) section for more information.

## BGP Next-Hop Address Tracking

BGP monitors the next-hop address of installed routes to verify next-hop reachability and to select, install, and validate the BGP best path. BGP next-hop address tracking speeds up this next-hop reachability test by triggering the verification process when routes change in the Routing Information Base (RIB) that may affect BGP next-hop reachability.

BGP receives notifications from the RIB when the next-hop information changes (event-driven notifications). BGP is notified when any of the following events occurs:

- The next hop becomes unreachable.
- The next hop becomes reachable.
- The fully recursed Interior Gateway Protocol (IGP) metric to the next hop changes.
- The first hop IP address or first hop interface changes.
- The next hop becomes connected.
- The next hop becomes unconnected.
- The next hop becomes a local address.
- The next hop becomes a nonlocal address.



---

**Note** Reachability and recursed metric events trigger a best-path recalculation.

---

Event notifications from the RIB are classified as critical and noncritical. Notifications for critical and noncritical events are sent in separate batches. However, a noncritical event is sent with the critical events if the noncritical event is pending and there is a request to read the critical events.



- Critical events are related to next-hop reachability, such as the loss of next hops resulting in a switchover to a different path. A change in the IGP metric for a next hop resulting in a switchover to a different path can also be considered a critical event.
- Non-critical events are related to next hops being added without affecting the best path or changing the IGP metric to a single next hop.

See the [Configuring BGP Next-Hop Address Tracking](#) section for more information.

## Route Redistribution

You can configure BGP to redistribute static routes or routes from other protocols. You must configure a route map with the redistribution to control which routes are passed into BGP. A route map allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. See [Configuring Route Policy Manager](#), for more information.

You can use route maps to override the default behavior in both scenarios, but be careful when doing so as incorrect use of route maps can result in network loops. The following examples show how to use route maps to change the default behavior.

You can change the default behavior for scenario 1 by modifying the route map as follows:

```
route-map foo permit 10
  match route-type internal
router ospf 1
  redistribute bgp 100 route-map foo
```

Similarly, you can change the default behavior for scenario 2 by modifying the route map as follows:

```
route-map foo deny 10
  match route-type internal
router ospf 1
  vrf bar
  redistribute bgp 100 route-map foo
```

## Labeled and Unlabeled Unicast Routes

In release 7.0(3)I7(6), SAFI-1 (unlabeled unicast) and SAFI-4 (labeled unicast routing) are now supported for IPv4 BGP on a single session. For more information, see the *Cisco Nexus 9000 Series NX-OS Label Switching Configuration Guide, Release 7.x*.

## BFD

This feature supports bidirectional forwarding detection (BFD) for IPv4 and IPv6. BFD is a detection protocol designed to provide fast forwarding-path failure detection times. BFD provides subsecond failure detection between two adjacent devices and can be less CPU-intensive than protocol hello messages because some of the BFD load can be distributed onto the data plane on supported modules.

BFD for BGP is supported on eBGP peers and iBGP single-hop peers. Configure the **update-source** option in neighbor configuration mode for iBGP single-hop peers using BFD.

Beginning with Cisco NX-OS Release 9.3(3), BFD for BGP is also supported for BGP IPv4 and IPv6 prefix peers. This support enables BGP to use multihop BFD, which improves BGP convergence times. Both single-hop and multihop BGP are supported for prefix peers.

Beginning with Cisco NX-OS Release 9.3(3), BFD supports BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families. However, BFD multihop is not supported with unnumbered BGP.

See the [Cisco Nexus 9000 Series NX-OS Interfaces Configuration Guide](#) for more information.

## Tuning BGP

You can modify the default behavior of BGP through BGP timers and by adjusting the best-path algorithm.

### BGP Timers

BGP uses different types of timers for neighbor session and global protocol events. Each established session has a minimum of two timers for sending periodic keepalive messages and for timing out sessions when peer keepalives do not arrive within the expected time. In addition, there are other timers for handling specific features. Typically, you configure these timers in seconds. The timers include a random adjustment so that the same timers on different BGP peers trigger at different times.

### Tuning the Best-Path Algorithm

You can modify the default behavior of the best-path algorithm through optional configuration parameters, including changing how the algorithm handles the multi-exit discriminator (MED) attribute and the router ID.

## Multiprotocol BGP

BGP on Cisco NX-OS supports multiple address families. Multiprotocol BGP (MP-BGP) carries different sets of routes depending on the address family. For example, BGP can carry one set of routes for IPv4 unicast routing, one set of routes for IPv4 multicast routing, and one set of routes for IPv6 multicast routing. You can use MP-BGP for reverse-path forwarding (RPF) checks in IP multicast networks.



---

**Note** Because Multicast BGP does not propagate multicast state information, you need a multicast protocol, such as Protocol Independent Multicast (PIM).

---

Use the router address-family and neighbor address-family configuration modes to support multiprotocol BGP configurations. MP-BGP maintains separate RIBs for each configured address family, such as a unicast RIB and a multicast RIB for BGP.

A multiprotocol BGP network is backward compatible but BGP peers that do not support multiprotocol extensions cannot forward routing information, such as address family identifier information, that the multiprotocol extensions carry.

### RFC 5549

BGP supports RFC 5549, which allows an IPv4 prefix to be carried over an IPv6 next hop. Because BGP is running on every hop, all routers can forward IPv4 and IPv6 traffic. Therefore, there is no need to support IPv6 tunnels between any routers. BGP installs IPv4 over an IPv6 route to the Unicast Route Information Base (URIB).

Beginning with Cisco NX-OS Release 9.2(2), Cisco Nexus 9500 platform switches with -R line cards support RFC 5549.

Currently, NX-OS does not support IPv6 recursive next-hops (RNH) for an IPv4 route.

## RFC 6368

### Introduction

This section describes how the Internal Border Gateway Protocol (iBGP) between Provider Edge (PE) and Customer Edge (CE) feature is implemented in Cisco NX-OS.

In current deployments, when BGP is used as the Provider/Customer Edge routing protocol, these peering sessions are configured as an external peering between the VPN provider autonomous system (AS) and the customer network autonomous system.

RFC 6368 adds support for these peers to be configured as iBGP peers instead.

Beginning with Cisco NX-OS Release 10.1(2), RFC 6368 support is enabled for EVPN-VxLANv4 and EVPN-VxLANv6.

Cisco NX-OS DC fabric deployment supports EVPN-VxLAN, EVPN-SR/MPLS, L3VPN-MPLS, and L3VPN-SRv6 as part of VPN configuration within DC.

### Framework

Beginning with Cisco NX-OS Release 10.1(2), deploying iBGP PE-CE feature:

- You can have one single Autonomous System Number (ASN) on the multiple sites of the VRF, without the deployment of External Border Gateway Protocol (eBGP) with as-override.
- You can give internal route reflection towards the CE routers, acting as if the Provider core is one transparent Route Reflector (RR).

With this feature, the VRF sites can have the same ASN as the provider core. However, in case the ASN of the VRF sites are different than the ASN of the provider core, it can be made to appear the same with the use of the feature local Autonomous System (AS).

### Implement iBGP PE-CE

Here are the two major parts to make this feature work:

- A new attribute `ATTR_SET` added to the BGP protocol to carry the VPN BGP attributes across the provider core in a transparent manner.
- Make the PE router a RR for the iBGP sessions towards the CE routers in the VRF.

The new `ATTR_SET` attribute allows the provider to carry all the BGP attributes of the customer transparently and does not interfere with the provider attributes and BGP policies. Such attributes are the cluster list, local preference, and so on.

#### BGP Customer Route Attribute

`ATTR_SET` is the new BGP attribute used to carry the VPN BGP attributes of the provider customer. It is an optional transitive attribute. In this attribute, Local Preference, Med, Origin, AS Path, Originator ID, Cluster list attributes will be carried across the provider network. The `ATTR_SET` attribute has the format:

```
+-----+
```

```
| Attr Flags (O|T) Code = 128 |
+-----+
| Attr. Length (1 or 2 octets) |
+-----+
| Origin AS (4 octets)      |
+-----+
|Path Attributes (variable) |
+-----+
```

- Attribute Flags are regular BGP attribute flags.
- Attribute length indicates whether the length is one or two octets.
- Origin AS field is to prevent a leak of one route that originated in one AS to be leaked to another AS without proper manipulation of the `AS_PATH`.
- The variable-length path attributes field carries VPN BGP attributes that must be carried across the provider core.

For more information on the implementation of iBGP PE-CE, see [IOS Implementation of the iBGP PE-CE Feature](#).

This example shows BGP neighbor configuration on PE device for iBGP Customer Edge device:

```
router bgp 200
vrf nxbgp3-leaf2-2
address-family ipv4 unicast
redistribute static route-map ALLOW-ALL
address-family ipv6 unicast
redistribute static route-map ALLOW-ALL
neighbor 101.101.101.101 remote-as 200
description ibgp sample config
internal-vpn-client (1)
address-family ipv4 unicast
route-reflector-client (2)
next-hop-self (3)
```

## BGP Monitoring Protocol

The BGP Monitoring Protocol (BMP) monitors BGP updates and peer statistics and is supported for all Cisco Nexus 9000 Series switches.

Using this protocol, the BGP speaker connects to external BMP servers and sends them information regarding BGP events. A maximum of two BMP servers can be configured in a BGP speaker, and each BGP peer can be configured for monitoring by all or a subset of the BMP servers. The BGP speaker does not accept any information from the BMP server.

## Graceful Restart and High Availability

Cisco NX-OS supports nonstop forwarding and graceful restart for BGP.

You can use nonstop forwarding (NSF) for BGP to forward data packets along known routes in the Forward Information Base (FIB) while the BGP routing protocol information is being restored following a failover. With NSF, BGP peers do not experience routing flaps. During a failover, the data traffic is forwarded through intelligent modules while the standby supervisor becomes active.

If a Cisco NX-OS router experiences a cold reboot, the network does not forward traffic to the router and removes the router from the network topology. In this scenario, BGP experiences a nongraceful restart and removes all routes. When Cisco NX-OS applies the startup configuration, BGP reestablishes peering sessions and relearns the routes.

A Cisco NX-OS router that has dual supervisors can experience a stateful supervisor switchover. During the switchover, BGP uses nonstop forwarding to forward traffic based on the information in the FIB, and the system is not removed from the network topology. A router whose neighbor is restarting is referred to as a "helper." After the switchover, a graceful restart operation begins. When it is in progress, both routers reestablish their neighbor relationship and exchange their BGP routes. The helper continues to forward prefixes pointing to the restarting peer, and the restarting router continues to forward traffic to peers even though those neighbor relationships are restarting. When the restarting router has all route updates from all BGP peers that are graceful restart capable, the graceful restart is complete, and BGP informs the neighbors that it is operational again.

When a router detects that a graceful restart operation is in progress, both routers exchange their topology tables. When the router has route updates from all BGP peers, it removes all the stale routes and runs the best-path algorithm on the updated routes.

After the switchover, Cisco NX-OS applies the running configuration, and BGP informs the neighbors that it is operational again.

For single-hop iBGP peers with update-source configured under neighbor configuration mode, the peer supports fast external fall-over.

Beginning with Cisco NX-OS Release 9.3(3), BGP prefix peers support graceful restarts.

With the additional BGP paths feature, if the number of paths advertised for a given prefix is the same before and after restart, the choice of path ID guarantees the final state and removal of stale paths. If fewer paths are advertised for a given prefix after a restart, stale paths can occur on the graceful restart helper peer.

## Low Memory Handling

BGP reacts to low memory for the following conditions:

- Minor alert—BGP does not establish any new eBGP peers. BGP continues to establish new iBGP peers and confederate peers. Established peers remain, but reset peers are not re-established.
- Severe alert—BGP shuts down select established eBGP peers every two minutes until the memory alert becomes minor. For each eBGP peer, BGP calculates the ratio of total number of paths received to the number of paths selected as best paths. The peers with the highest ratio are selected to be shut down to reduce memory usage. You must clear a shutdown eBGP peer before you can bring the eBGP peer back up to avoid oscillation.



---

**Note** You can exempt important eBGP peers from this selection process.

---

- Critical alert—BGP gracefully shuts down all the established peers. You must clear a shutdown BGP peer before you can bring the BGP peer back up.

See the [Tuning BGP](#) section for more information on how to exempt a BGP peer from a shutdown due to a low memory condition.

## Virtualization Support

You can configure one BGP instance. BGP supports virtual routing and forwarding (VRF) instances.

## Prerequisites for Advanced BGP

Advanced BGP has the following prerequisites:

- You must enable BGP (see the [Enabling BGP](#) section).
- You should have a valid router ID configured on the system.
- You must have an AS number, either assigned by a Regional Internet Registry (RIR) or locally administered.
- You must have reachability (such as an interior gateway protocol [IGP], a static route, or a direct connection) to the peer that you are trying to make a neighbor relationship with.
- You must explicitly configure an address family under a neighbor for the BGP session establishment.

## Guidelines and Limitations for Advanced BGP

Advanced BGP has the following configuration guidelines and limitations:

- There are three scenarios in which the command behavior has changed beginning with Cisco NX-OS Release 9.3(5):

```
• Router bgp 1
  Template peer abc
    Ttl-security hops 30
  Neighbor 1.2.3.4
    Inherit peer abc
```

If you later enter the **ebgp-multihop 20** command, the configuration is blocked due to the presence of **ttl-security hops 30** command. Beginning with the Cisco NX-OS Release 9.3(5), the configuration is no longer blocked. However, the **ttl-security hops** command has priority and would be the enabled functionality.

```
• Router bgp 1
  Template peer abc
    Ebgp-multihops 20
  Neighbor 1.2.3.4
    Inherit peer abc
```

If you later enter the **ttl-security hops 30** command, the configuration is blocked due to the presence of **ebgp-multihop 20** command. Beginning with Cisco NX-OS Release 9.3(5), the configuration is no longer blocked. However again, the **ttl-security hops** command has priority and would be the enabled functionality.

```
• Router bgp 1
  Template peer abc
    Remote-as 1
  Neighbor 1.2.3.4
    Inherit peer abc
```

If you later enter the **ttl-security hops 30** or **ebgp-multihop 20** commands, they are blocked. Beginning with Cisco NX-OS Release 9.3(5), the configuration is not blocked. However, their functionalities are turned off as the **remote-as** command has priority which makes the peer an iBGP peer.

- Prefix peering operates only in passive TCP mode. It accepts incoming connections from remote peers if the peer address falls within the prefix.
- Beginning with Cisco NX-OS 9.3(5), a packet with a TTL value of 1 to a vPC peer is hardware that is forwarded.
- Configuring the **advertise-maps** command multiple times is not supported.
- Names in the prefix-list are case-insensitive. We recommend using unique names. Do not use the same name by modifying uppercase and lowercase characters. For example, CTCPrimaryNetworks and CtcPrimaryNetworks are not two different entries.
- The dynamic AS number prefix peer configuration overrides the individual AS number configuration that is inherited from a BGP template.
- If you configure a dynamic AS number for prefix peers in an AS confederation, BGP establishes sessions with only the AS numbers in the local confederation.
- BGP sessions that are created through a dynamic AS number prefix peer ignore any configured eBGP multihop time-to-live (TTL) value or a disabled check for directly connected peers.
- Configure a router ID for BGP to avoid automatic router ID changes and session flaps.
- Use the maximum-prefix configuration option per peer to restrict the number of routes that are received and system resources used.
- Configure the update source to establish a session with eBGP multihop sessions.
- Specify a BGP route map if you configure a redistribution.
- Configure the BGP router ID within a VRF.
- If you decrease the keepalive and hold timer values, the network might experience session flaps.
- When you redistribute BGP to IGP, iBGP is redistributed as well. To override this behavior, you must insert an extra deny statement into the route map.
- To enable BFD for iBGP single-hop peers, you must configure the **update-source** option on the physical interface.
- Beginning with Cisco NX-OS Release 9.3(3), BFD for BGP is supported for BGP IPv4 and IPv6 prefix peers.
- The following guidelines and limitations apply to the **remove-private-as** command:
  - It applies only to eBGP peers.
  - It can be configured only in neighbor configuration mode and not in neighbor-address-family mode.
  - If the AS-path includes both private and public AS numbers, the private AS numbers are not removed.
  - If the AS-path contains the AS number of the eBGP neighbor, the private AS numbers are not removed.



- Private AS numbers are removed only if all AS numbers in that AS-path belong to a private AS number range. Private AS numbers are not removed if a peer's AS number or a non-private AS number is found in the AS-path segment.
- If you use the **aggregate-address** command to configure aggregate addresses and the **suppress-fib-pending** command to suppress BGP routes, lossless traffic for aggregates cannot be ensured on BGP or system triggers.
- When you enable FIB suppression on the switch and route programming fails in the hardware, BGP advertises routes that are not programmed locally in the hardware.
- If you disable a command in the neighbor, template peer, template peer-session, or template peer-policy configuration mode (and the **inherit peer** or **inherit peer-session** command is present), you must use the **default** keyword to return the command to its default state. For example, to disable the **update-source loopback 0** command from the running configuration, you must enter the **default update-source loopback 0** command.
- When next-hop-self is configured for route-reflector clients, the route reflector advertises routes to its clients with itself as the next hop.
- The following guidelines and limitations apply to weighted ECMP:
  - Weighted ECMP is supported only for the IPv4 address family.
  - BGP uses the Link Bandwidth EXTCOMM defined in the draft-ietf-idr-link-bandwidth-06.txt to implement the weighted ECMP feature.
  - BGP accepts the Link Bandwidth EXTCOMM from both iBGP and eBGP peers.
- The following guidelines and limitations apply to BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families:
  - This feature does not support having the same link-local address configured across multiple interfaces.
  - This feature is not supported on logical interfaces (loopback). Only Ethernet interfaces, port-channel interfaces, subinterfaces, and breakout interfaces are supported.
  - Beginning with Cisco NX-OS Release 9.3(6), VLAN interfaces are supported.
  - This feature is supported only for IPv6-enabled interfaces with link-local addresses.
  - This feature is not supported when the configured prefix peer and interface have the same remote peer.
  - The following commands are not supported in neighbor interface configuration mode:
    - **disable-connected-check**
    - **maximum-peers**
    - **update-source**
    - **ebgp-multihop**
  - BFD multihop and the following commands are not supported for BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families:
    - **bfd-multihop**

- **bfd multihop interval**
- **bfd multihop authentication**
- BGP requires faster convergence time for route advertisements. To speed up detection of the Route Advertisement (RA) link-level protocol, enter the following commands on each IPv6-enabled interface that is using BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families:

```
interface Ethernet port/slot
ipv6 nd ra-interval 4 min 3
ipv6 nd ra-lifetime 10
```

- When configuring the BGP neighbor with link-local, you need to customize the TCAM "in-sup" from 512 to 768.
- The command [**maximum-paths eibgp**] is supported only in MPLS environments.

## Default Settings

The table lists the default settings for advanced BGP parameters.

Parameters	Default
BGP feature	Disabled
BGP additional paths	Disabled
Keep alive interval	60 seconds
Hold timer	180 seconds
Dynamic capability	Enabled

## Configuring Advanced BGP

### Enabling IP Forward on an Interface

To use RFC 5549, you must configure at least one IPv4 address. If you do not want to configure an IPv4 address, you must enable the IP forward feature to use RFC 5549.

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<b>interface</b> <i>type slot/port</i>  <b>Example:</b> switch(config)# interface ethernet 1/2 switch(config-if)#	Enters interface configuration mode.
<b>Step 3</b>	<b>ip forward</b>  <b>Example:</b> switch(config-if)# ip forward	Allows IPv4 traffic on the interface even when there is no IP address configuration on that interface.
<b>Step 4</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> switch(config-if)# copy running-config startup-config	Saves this configuration change.

## Configuring BGP Session Templates

You can use BGP session templates to simplify the BGP configuration for multiple BGP peers with similar configuration needs. BGP templates allow you to reuse common configuration blocks. You configure BGP templates first and then apply these templates to BGP peers.

With BGP session templates, you can configure session attributes such as inheritance, passwords, timers, and security.

A peer-session template can inherit from one other peer-session template. You can configure the second template to inherit from a third template. The first template also inherits this third template. This indirect inheritance can continue for up to seven peer-session templates.

Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).



**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the default form of the command to reset that attribute to the default state.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.

	Command or Action	Purpose
Step 2	<b>router bgp</b> <i>autonomous-system-number</i> <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	<b>template peer-session</b> <i>template-name</i> <b>Example:</b> switch(config-router)# template peer-session BaseSession switch(config-router-stmp)#	Enters peer-session template configuration mode.
Step 4	(Optional) <b>password</b> <i>number password</i> <b>Example:</b> switch(config-router-stmp)# password 0 test	Adds the clear text password test to the neighbor. The password is stored and displayed in type 3 encrypted form (3DES).
Step 5	(Optional) <b>timers</b> <i>keepalive hold</i> <b>Example:</b> switch(config-router-stmp)# timers 30 90	Adds the BGP keepalive and holdtimer values to the peer-session template. The default keepalive interval is 60. The default hold time is 180.
Step 6	<b>exit</b> <b>Example:</b> switch(config-router-stmp)# exit switch(config-router)#	Exits peer-session template configuration mode.
Step 7	<b>neighbor</b> <i>ip-address remote-as as-number</i> <b>Example:</b> switch(config-router)# neighbor 192.168.1.2 remote-as 65535 switch(config-router-neighbor)#	Places the router in the neighbor configuration mode for BGP routing and configures the neighbor IP address.
Step 8	<b>inherit peer-session</b> <i>template-name</i> <b>Example:</b> switch(config-router-neighbor)# inherit peer-session BaseSession switch(config-router-neighbor)#	Applies a peer-session template to the peer.
Step 9	(Optional) <b>description</b> <i>text</i> <b>Example:</b> switch(config-router-neighbor)# description Peer Router A switch(config-router-neighbor)#	Adds a description for the neighbor.
Step 10	(Optional) <b>show bgp peer-session</b> <i>template-name</i> <b>Example:</b>	Displays the peer-policy template.

	Command or Action	Purpose
	<pre>switch(config-router-neighbor)# show bgp peer-session BaseSession</pre>	
<b>Step 11</b>	<p>(Optional) <b>copy running-config startup-config</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# copy running-config startup-config</pre>	<p>Saves this configuration change.</p> <p>Use the <b>show bgp neighbor</b> command to see the template applied.</p>

### Example

This example shows how to configure a BGP peer-session template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# template peer-session BaseSession
switch(config-router-stmp)# timers 30 90
switch(config-router-stmp)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 65536
switch(config-router-neighbor)# inherit peer-session BaseSession
switch(config-router-neighbor)# description Peer Router A
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring BGP Peer-Policy Templates

You can configure a peer-policy template to define attributes for a particular address family. You assign a preference to each peer-policy template and these templates are inherited in the order specified, for up to five peer-policy templates in a neighbor address family.

Cisco NX-OS evaluates multiple peer policies for an address family using the preference value. The lowest preference value is evaluated first. Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template.

Peer-policy templates can configure address family-specific attributes such as AS-path filter lists, prefix lists, route reflection, and soft reconfiguration.



**Note** Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 9000 Series NX-OS Unicast Routing Command Reference*, for details on all commands available in the template.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).



**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the default form of the command to reset that attribute to the default state.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> switch# configure terminal	Enters configuration mode.
<b>Step 2</b>	<b>router bgp <i>autonomous-system-number</i></b> <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>template peer-session <i>template-name</i></b> <b>Example:</b> switch(config-router)# template peer-policy BasePolicy switch(config-router-ptmp)#	Creates a peer-policy template.
<b>Step 4</b>	(Optional) <b>advertise-active-only</b> <b>Example:</b> switch(config-router-ptmp)# advertise-active-only	Advertises only active routes to the peer.
<b>Step 5</b>	(Optional) <b>maximum-prefix <i>number</i></b> <b>Example:</b> switch(config-router-ptmp)# maximum-prefix 20	Sets the maximum number of prefixes allowed from this peer.
<b>Step 6</b>	<b>exit</b> <b>Example:</b> switch(config-router-ptmp)# exit switch(config-router)#	Exits peer-policy template configuration mode.
<b>Step 7</b>	<b>neighbor <i>ip-address</i> remote-as <i>as-number</i></b> <b>Example:</b> switch(config-router)# neighbor 192.168.1.2 remote-as 65535 switch(config-router-neighbor)#	Places the router in the neighbor configuration mode for BGP routing and configures the neighbor IP address.
<b>Step 8</b>	<b>address-family {ipv4   ipv6} {multicast   unicast}</b> <b>Example:</b>	Enters global address family configuration mode for the address family specified.

	Command or Action	Purpose
	<pre>switch(config-router-neighbor)# address-family ipv4 unicast switch(config-router-neighbor-af)#</pre>	
<b>Step 9</b>	<b>inherit peer-policy <i>template-name preference</i></b> <b>Example:</b> <pre>switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1</pre>	Applies a peer-policy template to the peer address family configuration and assigns the preference value for this peer policy.
<b>Step 10</b>	<b>(Optional) show bgp peer-policy <i>template-name</i></b> <b>Example:</b> <pre>switch(config-router-neighbor-af)# show bgp peer-policy BasePolicy</pre>	Displays the peer-policy template.
<b>Step 11</b>	<b>(Optional) copy running-config startup-config</b> <b>Example:</b> <pre>switch(config-router-neighbor-af)# copy running-config startup-config</pre>	Saves this configuration change.  Use the <b>show bgp neighbor</b> command to see the template applied.

### Example

This example shows how to configure a BGP peer-policy template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# template peer-session BasePolicy
switch(config-router-ptmp)# maximum-prefix 20
switch(config-router-ptmp)# exit
switch(config-router)# neighbor 192.168.1.1 remote-as 65536
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring BGP Peer Templates

You can configure BGP peer templates to combine session and policy attributes in one reusable configuration block. Peer templates can also inherit peer-session or peer-policy templates. Any attributes configured for the neighbor take priority over any attributes inherited by that neighbor from a BGP template. You configure only one peer template for a neighbor, but that peer template can inherit peer-session and peer-policy templates.

Peer templates support session and address family attributes, such as eBGP multihop time-to-live, maximum prefix, next-hop self, and timers.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).





**Note** When editing a template, you can use the **no** form of a command at either the peer or template level to explicitly override a setting in a template. You must use the default form of the command to reset that attribute to the default state.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> switch# configure terminal	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp</b> <i>autonomous-system-number</i> <b>Example:</b> switch(config)# router bgp 65535	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>template peer</b> <i>template-name</i> <b>Example:</b> switch(config-router)# template peer BasePeer	Enters peer template configuration mode.
<b>Step 4</b>	(Optional) <b>inherit peer-session</b> <i>template-name</i> <b>Example:</b> switch(config-router-neighbor)# inherit peer-session BaseSession	Adds a peer-session template to the peer template.
<b>Step 5</b>	(Optional) <b>address-family</b> { <i>ipv4 ipv6</i> } { <i>multicast unicast</i> } <b>Example:</b> switch(config-router-neighbor)# address-family ipv4 unicast switch(config-router-neighbor-af)	Configures the global address family configuration mode for the specified address family.
<b>Step 6</b>	(Optional) <b>inherit peer-policy</b> <i>template-name</i> <b>Example:</b> switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1	Applies a peer-policy template to the neighbor address family configuration.
<b>Step 7</b>	<b>exit</b> <b>Example:</b> switch(config-router-neighbor-af)# exit	Exits BGP neighbor address family configuration mode.
<b>Step 8</b>	(Optional) <b>timers</b> <i>keepalive hold</i> <b>Example:</b> switch(config-router-neighbor)# timers 45 100	Adds the BGP timer values to the peer. These values override the timer values in the peer-session template, BaseSession.

	Command or Action	Purpose
<b>Step 9</b>	<b>exit</b> <b>Example:</b> switch(config-router-neighbor)# exit	Exits BGP neighbor configuration mode.
<b>Step 10</b>	<b>neighbor ip-address remote-as as-number</b> <b>Example:</b> switch(config-router)# neighbor 192.168.1.2 remote-as 65535 switch(config-router-neighbor)#	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.
<b>Step 11</b>	<b>inherit peer template-name</b> <b>Example:</b> switch(config-router-neighbor)# inherit peer BasePeer	Inherits the peer template.
<b>Step 12</b>	(Optional) <b>timers keepalive hold</b> <b>Example:</b> switch(config-router-neighbor)# timers 60 120	Adds the BGP timer values to this neighbor. These values override the timer values in the peer template and the peer-session template.
<b>Step 13</b>	(Optional) <b>show bgp peer-template template-name</b> <b>Example:</b> switch(config-router-neighbor)# show bgp peer-template BasePeer	Displays the peer template.
<b>Step 14</b>	(Optional) <b>copy running-config startup-config</b> <b>Example:</b> switch(config-router-neighbor)# copy running-config startup-config	Saves this configuration change. Use the <b>show bgp neighbor</b> command to see the template applied.

### Example

This example shows how to configure a BGP peer template and apply it to a BGP peer:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# template peer BasePeer
switch(config-router-neighbor)# inherit peer-session BaseSession
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 65536
switch(config-router-neighbor)# inherit peer BasePeer
switch(config-router-neighbor)# copy running-config startup-config
```

## Configuring Prefix Peering

BGP supports the definition of a set of peers using a prefix for both IPv4 and IPv6. This feature allows you to not have to add each neighbor to the configuration.

When defining a prefix peering, you must specify the remote AS number with the prefix. BGP accepts any peer that connects from that prefix and autonomous system if the prefix peering does not exceed the configured maximum peers allowed.

When a BGP peer that is part of a prefix peering disconnects, Cisco NX-OS holds its peer structures for a defined prefix peer timeout value. An established peer can reset and reconnect without danger of being blocked because other peers have consumed all slots for that prefix peering.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>timers prefix-peer-timeout</b> <i>value</i>  <b>Example:</b> <pre>switch(config-router-neighbor) # timers prefix-peer-timeout 120</pre>	Configures the BGP prefix peering timeout value in router configuration mode. The range is from 0 to 1200 seconds. The default value is 30.  <b>Note</b> For prefix peers, set the prefix peer timeout to be greater than the configured graceful restart timer. If the prefix peer timeout is greater than the graceful restart timer, a peer's route is retained during its restart. If the prefix peer timeout is less than the graceful restart timer, the peer's route is purged by the prefix peer timeout, which may occur before the restart is complete.
<b>Step 2</b>	<b>maximum-peers</b> <i>value</i>  <b>Example:</b> <pre>switch(config-router-neighbor) # maximum-peers 120</pre>	Configures the maximum number of peers for this prefix peering in neighbor configuration mode. The range is from 1 to 1000.

### Example

This example shows how to configure a prefix peering that accepts up to 10 peers:

```
switch(config) # router bgp 65536
switch(config-router) # timers prefix-peer-timeout 120
switch(config-router) # neighbor 10.100.200.0/24 remote-as 65536
switch(config-router-neighbor) # maximum-peers 10
switch(config-router-neighbor) # address-family ipv4 unicast
switch(config-router-neighbor-af) #
```

Use the **show bgp ipv4 unicast neighbors** command to show the details of the configuration for that prefix peering with a list of the currently accepted instances and the counts of active, maximum concurrent, and total accepted peers.

## Configuring BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families

You can configure BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families for automatic BGP neighbor discovery using unnumbered interfaces. Doing so allows you to set up BGP sessions using an interface name as a BGP peer (rather than interface-scoped addresses). This feature relies on ICMPv6 neighbor discovery (ND) route advertisement (RA) for automatic neighbor discovery and on RFC 5549 for sending IPv4 routes with IPv6 next hop.

### Before you begin

You must enable BGP.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal	Enters configuration mode.
<b>Step 2</b>	<b>router bgp <i>autonomous-system-number</i></b>  <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker. The AS number can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.
<b>Step 3</b>	<b>neighbor <i>interface-name</i> remote-as {<i>as-number</i>   <b>route-map</b> <i>map-name</i>}</b>  <b>Example:</b> switch(config-router)# neighbor Ethernet1/1 remote-as route-map Testmap switch(config-router-neighbor)#	Places the router in the neighbor configuration mode for BGP routing and configures the interface for BGP peering.  <b>Note</b> You can specify only Ethernet interfaces, port-channel interfaces, subinterfaces, and breakout interfaces.  Beginning with Cisco NX-OS Release 9.3(6), you can specify a route map, which can contain AS lists and ranges. See <a href="#">Dynamic AS Numbers for Prefix Peers and Interface Peers</a> for more information about using dynamic AS numbers.
<b>Step 4</b>	<b>inherit peer <i>template-name</i></b>  <b>Example:</b> switch(config-router-neighbor)# inherit peer PEER	Inherits the peer template.

	Command or Action	Purpose
<b>Step 5</b>	<b>address-family {ipv4   ipv6} unicast</b> <b>Example:</b> <pre>switch(config-router-neighbor) # address-family ipv4 unicast switch(config-router-neighbor-af) #</pre>	Enters global address family configuration mode for the address family specified.
<b>Step 6</b>	(Optional) <b>show bgp {ipv4   ipv6} unicast neighbors interface</b> <b>Example:</b> <pre>switch(config-router-neighbor-af) # show bgp ipv4 unicast neighbors e1/25</pre> <b>Example:</b> <pre>switch(config-router-neighbor-af) # show bgp ipv6 unicast neighbors 3FFE:700:20:1::11</pre>	Displays information about BGP peers.
<b>Step 7</b>	(Optional) <b>show ip bgp neighbors interface-name</b> <b>Example:</b> <pre>switch(config-router-neighbor-af) # show ip bgp neighbors Ethernet1/1</pre>	Displays the interface used as a BGP peer.
<b>Step 8</b>	(Optional) <b>show ipv6 routers [interface interface]</b> <b>Example:</b> <pre>switch(config-router-neighbor-af) # show ipv6 routers interface Ethernet1/1</pre>	Displays the link-local address of remote IPv6 routers, which is learned through IPv6 ICMP router advertisement.
<b>Step 9</b>	(Optional) <b>copy running-config startup-config</b> <b>Example:</b> <pre>switch(config-router-neighbor-af) # copy running-config startup-config</pre>	Saves this configuration change.

### Example

This example shows how to configure BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families using a route map:

iBGP Interface Peering Configuration for Leaf 1:

```
switch# configure terminal
switch(config)# route-map Testmap permit 10
switch(config-route-map)# match as-number 100-200, 300, 400
switch(config-route-map)# exit
switch(config)# router bgp 65000
switch(config-router)# neighbor Ethernet1/1 remote-as route-map Testmap
switch(config-router-neighbor)# inherit peer PEER
switch(config-router-neighbor)# address-family ipv4 unicast
```

```
switch(config-router-neighbor)# address-family ipv6 unicast
switch(config-router-neighbor-af)# copy running-config startup-config
```

This example shows sample output for BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families:

```
switch(config-router-neighbor)# show bgp ipv4 unicast neighbors e1/15.1
BGP neighbor is fe80::2, remote AS 100, ibgp link, Peer index 4
Peer is an instance of interface peering Ethernet1/15.1
BGP version 4, remote router ID 5.5.5.5
Neighbor previous state = OpenConfirm
BGP state = Established, up for 2d16h
Neighbor vrf: default
Peer is directly attached, interface Ethernet1/15.1
Last read 00:00:54, hold time = 180, keepalive interval is 60 seconds
Last written 00:00:08, keepalive timer expiry due 00:00:51
Received 3869 messages, 0 notifications, 0 bytes in queue
Sent 3871 messages, 0 notifications, 0(0) bytes in queue
Enhanced error processing: On
0 discarded attributes
Connections established 2, dropped 1
Last reset by peer 2d16h, due to session closed
Last error length received: 0
Reset error value received 0
Reset error received major: 104 minor: 0
Notification data received:
Last reset by us never, due to No error
Last error length sent: 0
Reset error value sent: 0
Reset error sent major: 0 minor: 0
--More--
```

Interface Configuration:

IPv6 needs to be enabled on the corresponding interface using one of the following commands:

- **ipv6 address** *ipv6-address*
- **ipv6 address use-link-local-only**
- **ipv6 link-local** *link-local-address*

```
switch# configure terminal
switch(config)# interface Ethernet1/1
switch(config-if)# ipv6 address use-link-local-only
```




---

**Note** If an IPv4 address is not configured on the interface, the **ip forward** command must be configured on the interface to enable IPv4 forwarding.

---




---

**Note** IPv6 ND timers can be tuned to speed up neighbor discovery and for BGP faster route convergence.

---

```
switch(config-if)# ipv6 nd ra-interval 4 min 3
switch(config-if)# ipv6 nd ra-lifetime 10
```

---



**Note** Beginning with Cisco NX-OS Release 9.3(6), for customer deployments with parallel links, the following command must be added in interface mode:

```
switch(config-if)# ipv6 link-local use-bia
```

The command makes IPv6 LLA unique across different interfaces.

## Configuring BGP Authentication

You can configure BGP to authenticate route updates from peers using MD5 digests.

To configure BGP to use MD5 digests, use the following command in neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>password</b> {0   3   7} <i>string</i>  <b>Example:</b> switch(config-router-neighbor)# password BGPpassword	Configures an MD5 password for BGP neighbor sessions.

## Resetting a BGP Session

If you modify a route policy for BGP, you must reset the associated BGP peer sessions. If the BGP peers do not support route refresh, you can configure a soft reconfiguration for inbound policy changes. Cisco NX-OS automatically attempts a soft reset for the session.

To configure soft reconfiguration inbound, use the following command in neighbor address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>soft-reconfiguration inbound</b>  <b>Example:</b> switch(config-router-neighbor-af)# soft-reconfiguration inbound	Enables soft reconfiguration to store the inbound BGP route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<b>Step 2</b>	(Optional) <b>clear bgp</b> {ipv4   ipv6 } {unicast   multicast <i>ip-address</i> soft {in   out}  <b>Example:</b> switch# clear bgp ip unicast 192.0.2.1 soft in	Resets the BGP session without tearing down the TCP session.



	Command or Action	Purpose
<b>Step 3</b>	<b>clear bgp {ipv4   ipv6} {unicast   multicast} ip-address soft (in   out)</b>  <b>Example:</b> <pre>switch# clear bgp ip unicast 192.0.2.1 soft in</pre>	Resets the BGP session without tearing down the TCP session.

## Modifying the Next-Hop Address

You can modify the next-hop address used in a route advertisement in the following ways:

- Disable next-hop calculation and use the local BGP speaker address as the next-hop address.
- Set the next-hop address as a third-party address. Use this feature in situations where the original next-hop address is on the same subnet as the peer that the route is being sent to. Using this feature saves an extra hop during forwarding.

To modify the next-hop address, use the following commands in address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>next-hop-self</b>  <b>Example:</b> <pre>switch(config-router-neighbor-af) # next-hop-self</pre>	Uses the local BGP speaker address as the next-hop address in route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<b>Step 2</b>	<b>next-hop-third-party</b>  <b>Example:</b> <pre>switch(config-router-neighbor-af) # next-hop-third-party</pre>	Sets the next-hop address as a third-party address. Use this command for single-hop eBGP peers that do not have <b>next-hop-self</b> configured.

## Configuring BGP Next-Hop Address Tracking

BGP next-hop address tracking is enabled by default and cannot be disabled.

You can modify the delay interval between RIB checks to increase the performance of BGP next-hop tracking.

To modify the BGP next-hop address tracking, use the following commands in address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>nexthop trigger-delay {critical   non-critical} milliseconds</b>  <b>Example:</b>	Specifies the next-hop address tracking delay timer for critical next-hop reachability routes and for noncritical routes. The range is from 1

	Command or Action	Purpose
	<code>switch(config-router-af)# nexthop trigger-delay critical 5000</code>	to 4294967295 milliseconds. The critical timer default is 3000. The noncritical timer default is 10000.

## Configuring Next-Hop Filtering

BGP next-hop filtering allows you to specify that when a next-hop address is checked with the RIB, the underlying route for that next-hop address is passed through the route map. If the route map rejects the route, the next-hop address is treated as unreachable.

BGP marks all next hops that are rejected by the route policy as invalid and does not calculate the best path for the routes that use the invalid next-hop address.

To configure BGP next-hop filtering, use the following command in address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>nexthop route-map <i>name</i></b>  <b>Example:</b> <code>switch(config-router-af)# nexthop route-map nextHopLimits</code>	Specifies a route map to match the BGP next-hop route to. The name can be any case-sensitive, alphanumeric string up to 63 characters.

## Configuring Next-Hop Resolution via Default Route

BGP next-hop resolution allows you to specify if the IP default route is used for BGP next-hop resolution.

To configure BGP next-hop resolution, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>[no] nexthop suppress-default-resolution</b>  <b>Example:</b> <code>switch(config-router)# nexthop suppress-default-resolution</code>	Prevents resolution of BGP next hop through the IP default route.  When this command is enabled: <ul style="list-style-type: none"> <li>The output of the <b>show bgp process detail</b> command includes the following line: Use default route for nexthop resolution: No</li> <li>The output of the <b>show routing clients bgp</b> command includes the following line: Owned rnh will never resolve to 0.0.0.0/0</li> </ul>

## Controlling Reflected Routes Through Next-Hop-Self

NX-OS enables controlling the iBGP routes being sent to a specific peer through the **next-hop-self** [all] arguments. By using these arguments, you can selectively change the next-hop of routes even if the route is reflected.

Command	Purpose
<b>next-hop-self</b> [all] <b>Example:</b> <pre>switch(config-router-af)# next-hop-self all</pre>	Uses the local BGP speaker address as the next-hop address in route updates.  The all keyword is optional. If you specify all, all routes are sent to the peer with next-hop-self. If you do not specify all, the next hops of reflected routes are not changed.

## Shrinking Next-Hop Groups When A Session Goes Down

You can configure BGP to shrink ECMP groups in an accelerated way when a session goes down.

This feature applies to the following BGP path failure events:

- Any single or multiple Layer 3 link failures
- Line card failures
- BFD failure detections for BGP neighbors
- Administrative shutdown of BGP neighbors (using the shutdown command)

The accelerated handling of the first two events (Layer 3 link failures and line card failures) is enabled by default and does not require a configuration command to be enabled.

To configure the accelerated handling of the last two events, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>neighbor-down fib-accelerate</b> <b>Example:</b> <pre>switch(config-router)# neighbor-down fib-accelerate</pre>	Withdraws the corresponding next hop from all next-hop groups (ECMP groups and single next-hop routes) whenever a BGP session goes down.  <b>Note</b> This command applies to both IPv4 and IPv6 routes.

## Disabling Capabilities Negotiation

You can disable capabilities negotiations to interoperate with older BGP peers that do not support capabilities negotiation.

To disable capabilities negotiation, use the following command in neighbor configuration mode:

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>dont-capability-negotiate</b>  <b>Example:</b> <pre>switch(config-router-neighbor) # dont-capability-negotiate</pre>	Disables capabilities negotiation. You must manually reset the BGP sessions after configuring this command.

## Disabling Policy Batching

In BGP deployments where prefixes have unique attributes, BGP tries to identify routes with similar attributes to bundle in the same BGP update message. To avoid the overhead of this additional BGP processing, you can disable batching.

Cisco recommends that you disable policy batching for BGP deployments that have a large number of routes with unique next hops.

To disable policy batching, use the following command in router configuration mode:

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>disable-policy-batching</b>  <b>Example:</b> <pre>switch(config-router) # disable-policy-batching</pre>	Disables the batching evaluation of prefix advertisements to all peers.

## Configuring BGP Additional Paths

BGP supports sending and receiving multiple paths per prefix and advertising such paths.

### Advertising the Capability of Sending and Receiving Additional Paths

You can configure BGP to advertise the capability of sending and receiving additional paths to and from the BGP peers. To do so, use the following commands in neighbor address-family configuration mode:

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>[no] capability additional-paths send</b> <b>[disable]</b>  <b>Example:</b>	Advertises the capability to send additional paths to the BGP peer. The <b>disable</b> option disables the advertising capability of sending additional paths.

	Command or Action	Purpose
	<code>switch(config-router-neighbor-af) # capability additional-paths send</code>	The <b>no</b> form of this command disables the capability of sending additional paths.
<b>Step 2</b>	<b>[no] capability additional-paths receive</b> <b>[disable]</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # capability additional-paths receive</code>	Advertises the capability to receive additional paths from the BGP peer. The <b>disable</b> option disables the advertising capability of receiving additional paths.  The <b>no</b> form of this command disables the capability of receiving additional paths.
<b>Step 3</b>	<b>show bgp neighbor</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # show bgp neighbor</code>	Displays whether the local peer has advertised the additional paths send or receive capability to the remote peer.

### Example

This example shows how to configure BGP to advertise the capability to send and receive additional paths to and from the BGP peer:

```
switch# configure terminal
switch(config)# router bgp 100
switch(config-router)# neighbor 10.131.31.2 remote-as 100
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# capability additional-paths send
switch(config-router-neighbor-af)# capability additional-paths receive
```

## Configuring the Sending and Receiving of Additional Paths

You can configure the capability of sending and receiving additional paths to and from the BGP peers. To do so, use the following commands in address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>[no] additional-paths send</b>  <b>Example:</b> <code>switch(config-router-af) # additional-paths send</code>	Enables the send capability of additional paths for all of the neighbors under this address family for which the capability has not been disabled.  The <b>no</b> form of this command disables the send capability.
<b>Step 2</b>	<b>[no] additional-paths receive</b>  <b>Example:</b> <code>switch(config-router-af) # additional-paths receive</code>	Enables the receive capability of additional paths for all of the neighbors under this address family for which the capability has not been disabled.  The <b>no</b> form of this command disables the receive capability.

	Command or Action	Purpose
<b>Step 3</b>	<b>show bgp neighbor</b>  <b>Example:</b> <pre>switch(config-router-af)# show bgp neighbor</pre>	Displays whether the local peer as advertised the additional paths send or receive capability to the remote peer.

### Example

This example shows how to enable the additional paths send and receive capability for all neighbors under the specified address family for which this capability has not been disabled:

```
switch# configure terminal
switch(config)# router bgp 100
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# additional-paths send
switch(config-router-af)# additional-paths receive
```

## Configuring Advertised Paths

You can specify the paths that are advertised for BGP. To do so, use the following commands in route-map configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>[no] set ip next-hop unchanged</b>  <b>Example:</b> <pre>switch(config-route-map)# set ip next-hop unchanged</pre>	Specifies and unchanged next-hop IP address.
<b>Step 2</b>	<b>[no] set path-selection { all   backup   best2   multipaths }   advertise</b>  <b>Example:</b> <pre>switch(config-route-map)# set path-selection all advertise</pre>	Specifies that all paths be advertised for a given prefix. You can use one of the following options: <ul style="list-style-type: none"> <li>• <b>all</b>—Advertises all available valid paths.</li> <li>• <b>backup</b>—Advertises paths marked as backup paths. This option requires that backup paths be enabled using the <code>additional-path install backup</code> command.</li> <li>• <b>best2</b>—Advertises the second best path, which is the best path of the remaining available paths, except the already calculated best path.</li> <li>• <b>multipaths</b>—Advertises all multipaths. This option requires that multipaths be enabled using the <code>maximum-paths</code> command.</li> </ul>

	Command or Action	Purpose
		<p><b>Note</b> If there are no multipaths, the backup and best2 options are the same. If there are multipaths, best2 is the first path on the list of multipaths while backup is the best path of all available paths, except the calculated best path and multipaths.</p> <p>The <b>no</b> form of this command specifies that only the best path be advertised.</p>
<b>Step 3</b>	<p><b>show bgp {ipv4   ipv6} unicast</b> [<i>ip-address</i>   <i>ipv6-prefix</i>] [<b>vrf</b> <i>vrf-name</i>]</p> <p><b>Example:</b></p> <pre>switch(config-route-map)# show bgp ipv4 unicast</pre>	Displays the path ID for the additional paths of a prefix and advertisement information for these paths.

### Example

This example show how to specify that all paths be advertised for the prefix list p1:

```
switch# configure terminal
switch(config)# route-map PATH_SELECTION_RMAP
switch(config-route-map)# match ip address prefix-list p1
switch(config-route-map)# set path-selection all advertise
```

## Configuring Additional Path Selection

You can configure the capability fo selecting additional paths for a prefix. To do so, use the following commands in address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>[no] additional-paths selection route-map</b> <i>map-name</i></p> <p><b>Example:</b></p> <pre>switch(config-router-af)# additional paths selection route-map map1</pre>	<p>Configures the capability of selecting additional paths for a prefix.</p> <p>The <b>no</b> form of this command disables the additional paths selection capability.</p>
<b>Step 2</b>	<p><b>show bgp {ipv4   ipv6} unicast</b> [<i>ip-address</i>   <i>ipv6-prefix</i>] [<b>vrf</b> <i>vrf-name</i>]</p> <p><b>Example:</b></p> <pre>switch(config-route-af)# show bgp ipv4 unicast</pre>	Displays the path ID for the additional paths of a prefix and advertisement information for these paths.

### Example

This example shows how to configure additional paths selection under the specified address family:

```

switch# configure terminal
switch(config)# router bgp 100
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# additional-paths selection route-map PATH_SELECTION_RMAP

```

## Configuring eBGP

### Disabling eBGP Single-Hop Checking

You can configure eBGP to disable checking whether a single-hop eBGP peer is directly connected to the local router. Use this option for configuring a single-hop loopback eBGP session between directly connected switches.

To disable checking whether or not a single-hop eBGP peer is directly connected, use the following command in neighbor configuration mode:

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>disable-connected-check</b>  <b>Example:</b> <pre>switch(config-router-neighbor)# disable-connected-check</pre>	Disables checking whether or not a single-hop eBGP peer is directly connected. You must manually reset the BGP sessions after using this command.

### Configuring eBGP Multihop

You can configure the eBGP time-to-live (TTL) value to support eBGP multihop. In some situations, an eBGP peer is not directly connected to another eBGP peer and requires multiple hops to reach the remote eBGP peer. You can configure the eBGP TTL value for a neighbor session to allow these multihop sessions.



**Note** This configuration is not supported for BGP interface peering.

To configure eBGP multihop, use the following command in neighbor configuration mode:

#### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>ebgp-multihop <i>ttl-value</i></b>  <b>Example:</b>	Configures the eBGP TTL value for eBGP multihop. The range is from 2 to 255. You must



	Command or Action	Purpose
	<code>switch(config-router-neighbor)# ebgp-multihop 5</code>	manually reset the BGP sessions after using this command.

## Disabling a Fast External Fallover

By default, the Cisco NX-OS device supports fast external fallover for neighbors in all VRFs and address families (IPv4 or IPv6). Typically, when a BGP router loses connectivity to a directly connected eBGP peer, BGP triggers a fast external fallover by resetting the eBGP session to the peer. You can disable this fast external fallover to limit the instability caused by link flaps.

To disable fast external fallover, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>no fast-external-fallover</b>  <b>Example:</b> <code>switch(config-router)# no fast-external-fallover</code>	Disables a fast external fallover for eBGP peers. This command is enabled by default.

## Limiting the AS-path Attribute

You can configure eBGP to discard routes that have a high number of AS numbers in the AS-path attribute.

To discard routes that have a high number of AS numbers in the AS-path attribute, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>maxas-limit <i>number</i></b>  <b>Example:</b> <code>switch(config-router)# maxas-limit 50</code>	Discards eBGP routes that have a number of AS-path segments that exceed the specified limit. The range is from 1 to 2000.

## Configuring Local AS Support

The local-AS feature allows a router to appear to be a member of a second autonomous system (AS), in addition to its real AS. Local AS allows two ISPs to merge without modifying peering arrangements. Routers in the merged ISP become members of the new autonomous system but continue to use their old AS numbers for their customers.

This feature can only be used for true eBGP peers. You cannot use this feature for two peers that are members of different confederation subautonomous systems.

Furthermore, the remote peer's ASN configured with the remote-as command cannot be identical to the local device's ASN configured with the local-as command.

To configure eBGP local AS support, use the following command in neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>local-as</b> <i>number</i> [ <b>no-prepend</b> [ <b>replace-as</b> [ <b>dual-as</b> ]]]  <b>Example:</b> switch(config-router-neighbor)# local-as 1.1	Configures eBGP to prepend the local AS <i>number</i> to the AS_PATH attribute. The AS <i>number</i> can be a 16-bit integer or a 32-bit integer in the form of a higher 16-bit decimal number and a lower 16-bit decimal number in xx.xx format.

### Example

This example shows how to configure local AS support on a VRF:

```
switch# configure terminal
switch(config)# router bgp 1
switch(config-router)# vrf test
switch(config-router-vrf)# local-as 1
switch(config-router-vrf)# show running-config bgp
```

## Configuring AS Confederations

To configure an AS confederation, you must specify a confederation identifier. To the outside world, the group of autonomous systems within the AS confederation look like a single autonomous system with the confederation identifier as the autonomous system number.

To configure a BGP confederation identifier, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>confederation identifier</b> <i>as-number</i>  <b>Example:</b> switch(config-router)# confederation identifier 4000	In router configuration mode, this command configures a BGP confederation identifier.  The command triggers an automatic notification and session reset for the BGP neighbor sessions.
<b>Step 2</b>	<b>bgp confederation peers</b> <i>as-number</i> [ <i>as-number2...</i> ]  <b>Example:</b> switch(config-router)# bgp confederation peers 5 33 44	In router configuration mode, this command configures the autonomous systems that belong to the AS confederation.  The command specifies a list of autonomous systems that belong to the confederation and it triggers an automatic notification and session reset for the BGP neighbor sessions.

# Configuring Route Reflector

You can configure iBGP peers as route reflector clients to the local BGP speaker, which acts as the route reflector. Together, a route reflector and its clients form a cluster. A cluster of clients usually has a single route reflector. In such instances, the cluster is identified by the router ID of the route reflector. To increase redundancy and avoid a single point of failure in the network, you can configure a cluster with more than one route reflector. You must configure all route reflectors in the cluster with the same 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster.

## Before you begin

You must enable BGP.

## Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>cluster-id <i>cluster-id</i></b>  <b>Example:</b> switch(config-router)# cluster-id 192.0.2.1	Configures the local router as one of the route reflectors that serve the cluster. You specify a cluster ID to identify the cluster. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<b>Step 4</b>	<b>address-family {<i>ipv4</i>   <i>ipv6</i>} {unicast   multicast}</b>  <b>Example:</b> switch(config-router)# address-family ipv4 unicast switch(config-router-af)#	Enters router address family configuration mode for the specified address family.
<b>Step 5</b>	(Optional) <b>client-to-client reflection</b>  <b>Example:</b> switch(config-router-af)# client-to-client reflection	Configures client-to-client route reflection. This feature is enabled by default. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<b>Step 6</b>	<b>exit</b>  <b>Example:</b> switch(config-router-af)# exit switch(config-router)#	Exits router address configuration mode.

	Command or Action	Purpose
<b>Step 7</b>	<b>neighbor ip-address remote-as as-number</b>  <b>Example:</b>  switch(config-router)# neighbor 192.0.2.10 remote-as 65535 switch(config-router-neighbor)#	Configures the IP address and AS number for a remote BGP peer.
<b>Step 8</b>	<b>address-family {ipv4   ipv6} {unicast   multicast}</b>  <b>Example:</b>  switch(config-router-neighbor)# address-family ipv4 unicast switch(config-router-neighbor-af)#	Enters neighbor address family configuration mode for the unicast IPv4 address family.
<b>Step 9</b>	<b>route-reflector-client</b>  <b>Example:</b>  switch(config-router-neighbor-af)# route-reflector-client	Configures the device as a BGP route reflector and configures the neighbor as its client. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
<b>Step 10</b>	(Optional) <b>show bgp {ipv4   ipv6} {unicast   multicast} neighbors</b>  <b>Example:</b>  switch(config-router-neighbor-af)# show bgp ipv4 unicast neighbors	Displays the BGP peers.
<b>Step 11</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b>  switch(config-router-neighbor-af)# copy running-config startup-config	Saves this configuration change.

### Example

This example shows how to configure the router as a route reflector and add one neighbor as a client:

```
switch(config)# router bgp 65536
switch(config-router)# neighbor 192.0.2.10 remote-as 65536
switch(config-router-neighbor)# address-family ip unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring Next-Hops on Reflected Routes Using an Outbound Route-Map

You can change the next-hop on reflected routes on a BGP route reflector using an outbound route-map. You can configure the outbound route-map to specify the peer's local address as the next-hop address.



**Note** The **next-hop-self** command does not enable this functionality for routes being reflected to clients by a route reflector. This functionality can only be enabled using an outbound route-map.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).

Ensure that you are in the correct VDC (or use the **switchto vdc** command).

You must enter the **set next-hop** command to configure an address family-specific next-hop address. For example, for the IPv6 address family, you must enter the **set ipv6 next-hop peer-address** command.

- When setting IPv4 next-hops using route-maps—If **set ip next-hop peer-address** matches the route-map, the next-hop is set to the peer's local address. If no next-hop is set in the route-map, the next-hop is set to the one stored in the path.
- When setting IPv6 next-hops using route-maps—If **set ipv6 next-hop peer-address** matches the route-map, the next-hop is set as follows:
  - For IPv6 peers, the next-hop is set to the peer's local IPv6 address.
  - For IPv4 peers, if **update-source** is configured, the next-hop is set to the source interface's IPv6 address, if any. If no IPv6 address is configured, no next-hop is set
  - For IPv4 peers, if **update-source** is not configured, the next-hop is set to the outgoing interface's IPv6 address, if any. If no IPv6 address is configured, no next-hop is set.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> <pre>switch(config)# router bgp 200 switch(config-router)#</pre>	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>neighbor <i>ip-address</i> remote-as <i>as-number</i></b>  <b>Example:</b> <pre>switch(config-router)# neighbor 192.0.2.12 remote-as 200 switch(config-router-neighbor)#</pre>	Configures the IP address and AS number for a remote BGP peer.
<b>Step 4</b>	(Optional) <b>update-source <i>interface number</i></b>  <b>Example:</b>	Specifies and updates the source of the BGP session.

	Command or Action	Purpose
	<code>switch(config-router-neighbor) # update-source loopback 300</code>	
<b>Step 5</b>	<b>address-family {ipv4   ipv6} {unicast   multicast}</b>  <b>Example:</b> <code>switch(config-router-neighbor) # address-family ipv4 unicast switch(config-router-neighbor-af) #</code>	Enters router address family configuration mode for the specified address family.
<b>Step 6</b>	<b>route-reflector-client</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # route-reflector-client</code>	Configures the device as a BGP route reflector and configures the neighbor as its client. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
<b>Step 7</b>	<b>route-map map-name out</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # route-map setrrnh out</code>	Applies the configured BGP policy to outgoing routes.
<b>Step 8</b>	(Optional) <b>show bgp {ipv4   ipv6} {unicast   multicast} [ip-address   ipv6-prefix] route-map map-name [vrf vrf-name]</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # show bgp ipv4 unicast route-map setrrnh</code>	Displays the BGP routes that match the route map.
<b>Step 9</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> <code>switch(config-router-neighbor-af) # copy running-config startup-config</code>	Saves this configuration change.

### Example

This example shows how to configure the next-hop on reflected routes on a BGP route reflector using an outbound route-map:

```
switch(config)# interface loopback 300
switch(config-if)# ip address 192.0.2.11/32
switch(config-if)# ipv6 address 2001::a0c:1a65/64
switch(config-if)# ip router ospf 1 area 0.0.0.0
switch(config-if)# exit
switch(config)# route-map setrrnh permit 10
switch(config-route-map)# set ip next-hop peer-address
switch(config-route-map)# exit
switch(config)# route-map setrrnhv6 permit 10
switch(config-route-map)# set ipv6 next-hop peer-address
switch(config-route-map)# exit
switch(config)# router bgp 200
switch(config-router)# neighbor 192.0.2.12 remote-as 200
```

```

switch(config-router-neighbor)# update-source loopback 300
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# route-map setrrnh out
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# address-family ipv6 unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# route-map setrrnhv6 out

```

## Configuring Route Dampening

You can configure route dampening to minimize route flaps propagating through your iBGP network.

To configure route dampening, use the following command in address-family or VRF address family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>dampening</b> [ <i>{half-life reuse-limit suppress-limit max-suppress-time   route-map map-name}</i> ]  <b>Example:</b> <pre>switch(config-router-af)# dampening route-map bgpDamp</pre>	Disables capabilities negotiation. The parameter values are as follows: <ul style="list-style-type: none"> <li>• <i>half-life</i>—The range is from 1 to 45.</li> <li>• <i>reuse-limit</i>—The range is from 1 to 20000.</li> <li>• <i>suppress-limit</i>—The range is from 1 to 20000.</li> <li>• <i>max-suppress-time</i>—The range is from 1 to 255.</li> </ul>

## Configuring Load Sharing and ECMP

You can configure the maximum number of paths that BGP adds to the route table for equal-cost multipath (ECMP) load balancing.

To configure the maximum number of paths, use the following command in router address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>maximum-paths</b> [ <i>ibgp</i> ] <i>maxpaths</i>  <b>Example:</b> <pre>switch(config-router-af)# maximum-paths 8</pre>	Configures the maximum number of equal-cost paths for load sharing. The default is 1.

## Configuring Maximum Prefixes

You can configure the maximum number of prefixes that BGP can receive from a BGP peer. If the number of prefixes exceeds this value, you can optionally configure BGP to generate a warning message or tear down the BGP session to the peer.

To configure the maximum allowed prefixes for a BGP peer, use the following command in neighbor address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>maximum-prefix</b> <i>maximum</i> [ <i>threshold</i> ] [ <i>restart time</i>   <b>warning-only</b> ]  <b>Example:</b> <pre>switch(config-router-neighbor-af) # maximum-prefix 12</pre>	Configures the maximum number of prefixes from a peer. The parameter ranges are as follows: <ul style="list-style-type: none"> <li>• <i>maximum</i>—The range is from 1 to 300000.</li> <li>• <i>threshold</i>—The range is from 1 to 100 percent. The default is 75 percent.</li> <li>• <i>time</i>—The range is from 1 to 65535 minutes.</li> </ul> This command triggers an automatic notification and session reset for the BGP neighbor sessions if the prefix is exceeded.

## Configuring DSCP

You can configure a differentiated services code point (DSCP) for a neighbor. You can specify a DSCP value for locally originated packets for IPv4 or IPv6.

To configure the DSCP value, use the following command in neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>dscp</b> <i>dscp_value</i>  <b>Example:</b> <pre>switch(config-router-neighbor) # dscp 63</pre> Below is an example of the corresponding <b>show</b> command:  <pre>show ipv6 bgp neighbors BGP neighbor is 10.1.1.1, remote AS 0, unknown link, Peer index 4   BGP version 4, remote router ID 0.0.0.0</pre>	Sets the differentiated services code point (DSCP) value for the neighbor. The DSCP value can be a number from 0 to 63, or it can be one of the following keywords: <b>ef</b> , <b>af11</b> , <b>af12</b> , <b>af13</b> , <b>af21</b> , <b>af22</b> , <b>af23</b> , <b>af31</b> , <b>af32</b> , <b>af33</b> , <b>af41</b> , <b>af42</b> , <b>af43</b> , <b>cs1</b> , <b>cs2</b> , <b>cs3</b> , <b>cs4</b> , <b>cs5</b> , <b>cs6</b> , or <b>cs7</b> .  The default value is cs6.



	Command or Action	Purpose
	<pre>BGP state = Idle, down for 00:13:34, retry in 0.000000 DSCP (DiffServ CodePoint): 0 Last read never, hold time = 180, keepalive interval is 60 seconds</pre>	

## Configuring Dynamic Capability

You can configure dynamic capability for a BGP peer.

To configure dynamic capability, use the following command in neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>dynamic-capability</b>  <b>Example:</b> <pre>switch(config-router-neighbor)# dynamic-capability</pre>	Enables dynamic capability. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

## Configuring Aggregate Addresses

You can configure aggregate address entries in the BGP route table.

To configure an aggregate address, use the following command in router address-family configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>aggregate-address</b> <i>ip-prefix/length</i> [ <b>as-set</b> ] [ <b>summary-only</b> ] [ <b>advertise-map</b> <i>map-name</i> ] [ <b>attribute-map</b> <i>map-name</i> ] [ <b>suppress-map</b> <i>map-name</i> ]  <b>Example:</b> <pre>switch(config-router-af)# aggregate-address 192.0.2.0/8 as-set</pre>	Creates an aggregate address. The path advertised for this route is an autonomous system set that consists of all elements contained in all paths that are being summarized: <ul style="list-style-type: none"> <li>• The <b>as-set</b> keyword generates autonomous system set path information and community information from contributing paths.</li> <li>• The <b>summary-only</b> keyword filters all more specific routes from updates.</li> <li>• The <b>advertise-map</b> keyword and argument specify the route map used to select attribute information from selected routes.</li> </ul>

	Command or Action	Purpose
		<ul style="list-style-type: none"> <li>• The <b>attribute-map</b> keyword and argument specify the route map used to select attribute information from the aggregate.</li> <li>• The <b>suppress-map</b> keyword and argument conditionally filter more specific routes. If you specify the <b>suppress-map</b> option while performing a BGP route aggregation, you can set the community attribute for a BGP route update. This option enables you to set community attributes on the more-specific routes.</li> <li>• The <b>suppress-map</b> keyword and argument conditionally filter more specific routes. If you specify the <b>suppress-map</b> option while performing a BGP route aggregation, you can either suppress certain more-specific routes from being advertised to its peers, or decide to advertise the more-specific routes with some community attributes set on them, depending upon the <b>suppress-map</b> route-map configuration. A route-map configured with only match clauses will suppress the more-specific routes that satisfy the match criteria. However, if a route-map is configured with match and set clauses, then the routes satisfying the match criteria will be advertised with the appropriate attributes as modified by the route-map. The second option enables you to set community attributes on the more-specific routes.</li> </ul>

## Suppressing BGP Routes

You can configure Cisco NX-OS to advertise newly learned BGP routes only after these routes are confirmed by the Forwarding Information Base (FIB) and programmed in the hardware. After the routes are programmed, subsequent changes to these routes do not require this hardware-programming check.

To suppress BGP routes, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>suppress-fib-pending</b> <b>Example:</b>	Suppresses newly learned BGP routes (IPv4 or IPv6) from being advertised to downstream

	Command or Action	Purpose
	<code>switch(config-router)# suppress-fib-pending</code>	BGP neighbors until the routes have been programmed in the hardware.

## Configuring BGP Conditional Advertisement

You can configure BGP conditional advertisement to limit the routes that BGP propagates. You define the following two route maps:

- **Advertise map**—Specifies the conditions that the route must match before BGP considers the conditional advertisement. This route map can contain any appropriate match statements.
- **Exist map or nonexist map**—Defines the prefix that must exist in the BGP table before BGP propagates a route that matches the advertise map. The nonexist map defines the prefix that must not exist in the BGP table before BGP propagates a route that matches the advertise map. BGP processes only the permit statements in the prefix list match statements in these route maps.

If the route does not pass the condition, BGP withdraws the route if it exists in the BGP table.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> <code>switch# configure terminal switch(config)#</code>	Enters configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> <code>switch(config)# router bgp 65535 switch(config-router)#</code>	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>neighbor <i>ip-address</i> remote-as <i>as-number</i></b>  <b>Example:</b> <code>switch(config-router)# neighbor 192.168.1.2 remote-as 65534 switch(config-router-neighbor)#</code>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.
<b>Step 4</b>	<b>address-family {<i>ipv4</i>   <i>ipv6</i>} {unicast   multicast}</b>  <b>Example:</b> <code>switch(config-router-neighbor)# address-family ipv4 multicast switch(config-router-neighbor-af)#</code>	Enters address family configuration mode.

	Command or Action	Purpose
<b>Step 5</b>	<p><b>advertise-map</b> <i>adv-map</i> {<b>exist-map</b> <i>exist-rmap</i> <b>non-exist-map</b> <i>nonexist-rmap</i>}</p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af)# advertise-map advertise exist-map exist</pre>	<p>Configures BGP to conditionally advertise routes based on the two configured route maps:</p> <ul style="list-style-type: none"> <li>• <i>adv-map</i>—Specifies a route map with <b>match</b> statements that the route must pass before BGP passes the route to the next route map. The <i>adv-map</i> is a case-sensitive, alphanumeric string up to 63 characters.</li> <li>• <i>exist-rmap</i>—Specifies a route map with match statements for a prefix list. A prefix in the BGP table must match a prefix in the prefix list before BGP advertises the route. The <i>exist-rmap</i> is a case-sensitive, alphanumeric string up to 63 characters.</li> <li>• <i>nonexist-rmap</i>—Specifies a route map with match statements for a prefix list. A prefix in the BGP table must not match a prefix in the prefix list before BGP advertises the route. The <i>nonexist-rmap</i> is a case-sensitive, alphanumeric string up to 63 characters.</li> </ul>
<b>Step 6</b>	<p>(Optional) <b>show bgp</b> {<b>ipv4</b>   <b>ipv6</b>} {<b>unicast</b>   <b>multicast</b>} <b>neighbors</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af)# show ip bgp neighbor</pre>	<p>Displays information about BGP and the configured conditional advertisement route maps.</p>
<b>Step 7</b>	<p>(Optional) <b>copy running-config startup-config</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af)# copy running-config startup-config</pre>	<p>Saves this configuration change.</p>

### Example

This example shows how to configure BGP conditional advertisement:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# neighbor 192.0.2.2 remote-as 65537
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# advertise-map advertise exist-map exist
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# exit
switch(config-router)# exit
switch(config)# route-map advertise
switch(config-route-map)# match as-path pathList
switch(config-route-map)# exit
```

```
switch(config)# route-map exit
switch(config-route-map)# match ip address prefix-list plist
switch(config-route-map)# exit
switch(config)# ip prefix-list plist permit 209.165.201.0/27
```

## Configuring Route Redistribution

You can configure BGP to accept routing information from another routing protocol and redistribute that information through the BGP network. Optionally, you can assign a default route for redistributed routes.

### Before you begin

You must enable BGP.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>address-family {ipv4   ipv6 } {unicast   multicast}</b>  <b>Example:</b> switch(config-router)# address-family vpv4 unicast switch(config-router-af)#	Enters address family configuration mode.
<b>Step 4</b>	<b>address-family {ipv4   ipv6} {unicast   multicast}</b>  <b>Example:</b> switch(config-router)# address-family ipv4 unicast switch(config-router-af)#	Enters address-family configuration mode.
<b>Step 5</b>	<b>redistribute {direct   {eigrp   isis   ospf   ospfv3   rip} instance-tag   static} route-map <i>map-name</i></b>  <b>Example:</b> switch(config-router-af)# redistribute eigrp 201 route-map Eigrpmap	Redistributes routes from other protocols into BGP.

	Command or Action	Purpose
<b>Step 6</b>	<b>redistribute</b> {direct   {eigrp   isis   ospf   ospfv3   rip} instance-tag   static} route-map map-name  <b>Example:</b> <pre>switch(config-router-af)# redistribute eigrp 201 route-map Eigrpmap</pre>	Redistributes routes from other protocols into BGP.
<b>Step 7</b>	(Optional) <b>default-metric</b> value  <b>Example:</b> <pre>switch(config-router-af)# default-metric 33</pre>	Generates a default route into BGP.
<b>Step 8</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> <pre>switch(config-router-af)# copy running-config startup-config</pre>	Saves this configuration change.

### Example

This example shows how to redistribute EIGRP into BGP:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# redistribute eigrp 201 route-map Eigrpmap
switch(config-router-af)# copy running-config startup-config
```

## Advertising the Default Route

You can configure BGP to advertise the default route (network 0.0.0.0).

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
<b>Step 2</b>	<b>route-map allow permit</b>  <b>Example:</b>	Enters router map configuration mode and defines the conditions for redistributing routes.

	Command or Action	Purpose
	<code>switch(config)# route-map allow permit</code> <code>switch(config-route-map)#</code>	
<b>Step 3</b>	<b>exit</b>  <b>Example:</b> <code>switch(config-route-map)# exit</code> <code>switch(config)#</code>	Exits router map configuration mode.
<b>Step 4</b>	<b>ip route <i>ip-address network-mask</i> null</b> <b><i>null-interface-number</i></b>  <b>Example:</b> <code>switch(config)# ip route 192.0.2.1</code> <code>255.255.255.0 null 0</code>	Configures the IP address.
<b>Step 5</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> <code>switch(config)# router bgp 65535</code> <code>switch(config-router)#</code>	Enters BGP mode and assigns the AS number to the local BGP speaker.
<b>Step 6</b>	<b>address-family {<i>ipv4</i>   <i>ipv6</i>} unicast</b>  <b>Example:</b> <code>switch(config-router)# address-family</code> <code>ipv4 unicast</code> <code>switch(config-router-af)#</code>	Enters address-family configuration mode.
<b>Step 7</b>	<b>default-information originate</b>  <b>Example:</b> <code>switch(config-router-af)#</code> <code>default-information originate</code>	Advertises the default route.
<b>Step 8</b>	<b>redistribute static route-map allow</b>  <b>Example:</b> <code>switch(config-router-af)# redistribute</code> <code>static route-map allow</code>	Redistributes the default route.
<b>Step 9</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> <code>switch(config-router-af)# copy</code> <code>running-config startup-config</code>	Saves this configuration change.

## Configuring BGP Attribute Filtering and Error Handling

Beginning with Cisco NX-OS Release 9.3(3), you can configure BGP attribute filtering and error handling to provide an increased level of security. The following features are available and implemented in the following order:

- **Path attribute treat-as-withdraw:** Allows you to treat-as-withdraw a BGP update from a specific neighbor if the update contains a specified attribute type. The prefixes contained in the update are removed from the routing table.
- **Path attribute discard:** Allows you to remove specific path attributes in a BGP update from a specific neighbor.
- **Enhanced attribute error handling:** Prevents peer sessions from flapping due to a malformed update.

Attribute types 1, 2, 3, 4, 5, 8, 14, 15, and 16 cannot be configured for path attribute treat-as-withdraw and path attribute discard. Attribute type 9 (Originator) and type 10 (Cluster-id) can be configured for eBGP neighbors only.

## Treating as Withdraw Path Attributes from a BGP Update Message

To "treat-as-withdraw" BGP updates that contain specific path attributes, use the following command in router neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>[no] path-attribute treat-as-withdraw</b> [<i>value</i>   <i>range start end</i>] <b>in</b></p> <p><b>Example:</b></p> <pre>switch#(config-router)# neighbor 10.20.30.40 switch(config-router-neighbor)# path-attribute treat-as-withdraw 100 in</pre> <p><b>Example:</b></p> <pre>switch#(config-router)# neighbor 10.20.30.40 switch(config-router-neighbor)# path-attribute treat-as-withdraw range 21 255 in</pre>	<p>Treats as withdraw any incoming BGP update messages that contain the specified path attribute or range of path attributes and triggers an inbound route refresh to ensure that the routing table is up to date. Any prefixes in a BGP update that are treat-as-withdraw are removed from the BGP routing table.</p> <p>This command is also supported for BGP template peers and BGP template peer sessions.</p>

## Discarding Path Attributes from a BGP Update Message

To discard BGP updates that contain specific path attributes, use the following command in router neighbor configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>[no] path-attribute discard</b> [<i>value</i>   <i>range start end</i>] <b>in</b></p> <p><b>Example:</b></p> <pre>switch#(config-router)# neighbor 10.20.30.40</pre>	<p>Drops specified path attributes in BGP update messages for the specified neighbor and triggers an inbound route refresh to ensure that the routing table is up to date. You can configure a specific attribute or an entire range of unwanted attributes.</p>



	Command or Action	Purpose
	<pre>switch(config-router-neighbor)# path-attribute discard 100 in</pre> <p><b>Example:</b></p> <pre>switch#(config-router)# neighbor 10.20.30.40 switch(config-router-neighbor)# path-attribute discard range 100 255 in</pre>	<p>This command is also supported for BGP template peers and BGP template peer sessions.</p> <p><b>Note</b> When the same path attribute is configured for both discard and treat-as-withdraw, treat-as-withdraw has a higher priority.</p>

## Enabling or Disabling Enhanced Attribute Error Handling

BGP enhanced attribute error handling is enabled by default but can be disabled. This feature, which complies with RFC 7606, prevents peer sessions from flapping due to a malformed update. The default behavior applies to both eBGP and iBGP peers.

To disable or reenable enhanced error handling, use the following command in router configuration mode:

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>[no] enhanced-error</b></p> <p><b>Example:</b></p> <pre>switch(config)# router bgp 1000 switch(config-router)# enhanced-error</pre>	Enables or disables BGP enhanced attribute error handling.

## Displaying Discarded or Unknown Path Attributes

To display information about discarded or unknown path attributes, perform one of the following tasks:

Command	Purpose
<b>show bgp {ipv4   ipv6} unicast path-attribute discard]</b>	Displays all prefixes for which an attribute has been discarded.
<b>show bgp {ipv4   ipv6} unicast path-attribute unknown]</b>	Displays all prefixes that have an unknown attribute.
<b>show bgp {ipv4   ipv6} unicast ip-address</b>	Displays the unknown attributes and discarded attributes associated with a prefix.

The following example shows the prefixes for which an attribute has been discarded:

```
switch# show bgp ipv4 unicast path-attribute discard
Network          Next Hop
1.1.1.1/32       20.1.1.1
1.1.1.2/32       20.1.1.1
1.1.1.3/32       20.1.1.1
```

The following example shows the prefixes that have an unknown attribute:

```
switch# show bgp ipv4 unicast path-attribute unknown
Network          Next Hop
2.2.2.2/32       20.1.1.1
2.2.2.3/32       20.1.1.1
```

The following example shows the unknown attributes and discarded attributes associated with a prefix:

```
switch# show bgp ipv4 unicast 2.2.2.2
BGP routing table entry for 2.2.2.2/32, version 6241
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  1000
    20.1.1.1 from 20.1.1.1 (20.1.1.1)
      Origin IGP, localpref 100, valid, external, best
      unknown transitive attribute: flag 0xE0 type 0x62 length 0x64
        value 0000 0000 0100 0000 0200 0000 0300 0000
              0400 0000 0500 0000 0600 0000 0700 0000
              0800 0000 0900 0000 0A00 0000 0B00 0000
              0C00 0000 0D00 0000 0E00 0000 0F00 0000
              1000 0000 1100 0000 1200 0000 1300 0000
              1400 0000 1500 0000 1600 0000 1700 0000
              1800 0000
      rx pathid: 0, tx pathid: 0x0
      Updated on Jul 20 2019 07:50:43 PST
```

## Tuning BGP

You can tune BGP characteristics through a series of optional parameters.

To tune BGP, use the following optional commands in router configuration mode:

Command	Purpose
<p><b>bestpath</b> [<b>always-compare-med</b>   <b>as-path multipath-relax</b>   <b>compare-routerid</b>   <b>cost-community ignore</b>   <b>igp-metric ignore</b>   <b>med {confed   missing-as-worst   non-deterministic}</b>]</p> <p><b>Example:</b></p> <pre>switch(config-router)# bestpath always-compare-med</pre>	<p>Modifies the best-path algorithm. The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <b>always-compare-med</b> —Compares MED on paths from different autonomous systems.</li> <li>• <b>as-path multipath-relax</b> —Allows load sharing across the providers with different (but equal-length) AS paths. Without this option, the AS paths must be identical for load sharing.</li> <li>• <b>compare-routerid</b> —Compares the router IDs for identical eBGP paths.</li> <li>• <b>cost-community ignore</b> —Ignores the cost community for BGP best-path calculations.</li> <li>• <b>igp-metric ignore</b> —Ignores the Interior Gateway Protocol (IGP) metric for next hop during best-path selection. This option is supported beginning with Cisco NX-OS Release 9.2(2).</li> <li>• <b>med confed</b> —Forces bestpath to do a MED comparison only between paths originated within a confederation.</li> <li>• <b>med missing-as-worst</b> —Treats a missing MED as the highest MED.</li> <li>• <b>med non-deterministic</b> —Does not always pick the best MED path from among the paths from the same autonomous system.</li> </ul>
<p><b>enforce-first-as</b></p> <p><b>Example:</b></p> <pre>switch(config-router)# enforce-first-as</pre>	<p>Enforces the neighbor autonomous system to be the first AS number listed in the AS_path attribute for eBGP.</p>
<p><b>log-neighbor-changes</b></p> <p><b>Example:</b></p> <pre>switch(config-router)# log-neighbor-changes</pre>	<p>Generates a system message when any neighbor changes state.</p> <p><b>Note</b> To suppress neighbor status change messages for a specific neighbor, you can use the <b>log-neighbor-changes disable</b> command in router address-family configuration mode.</p>
<p><b>router-id</b> <i>id</i></p> <p><b>Example:</b></p> <pre>switch(config-router)# router-id 10.165.20.1</pre>	<p>Manually configures the router ID for this BGP speaker.</p>

Command	Purpose
<p><b>timers</b> [<b>bestpath-delay</b> <i>delay</i>   <i>bgpkeepalive holdtime</i>   <b>prefix-peer-timeout</b> <i>timeout</i>]</p> <p><b>Example:</b></p> <pre>switch(config-router)# timers bgp 90 270</pre>	<p>Sets BGP timer values. The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <i>delay</i> —Initial best-path timeout value after a restart. The range is from 0 to 3600 seconds. The default value is 300.</li> <li>• <i>keepalive</i> —BGP session keepalive time. The range is from 0 to 3600 seconds. The default value is 60.</li> <li>• <i>holdtime</i> —BGP session hold time. The range is from 0 to 3600 seconds. The default value is 180.</li> <li>• <i>timeout</i> —Prefix peer timeout value. The range is from 0 to 1200 seconds. The default value is 30.</li> </ul> <p>You must manually reset the BGP sessions after configuring this command.</p>

To tune BGP, use the following optional commands in router address-family configuration mode:

Command	Purpose
<p><b>distance</b> <i>ebgp-distance ibgp-distance local-distance</i></p> <p><b>Example:</b></p> <pre>switch(config-router-af)# distance 20 100 200</pre>	<p>Sets the administrative distance for BGP. The range is from 1 to 255. The defaults are as follows:</p> <ul style="list-style-type: none"> <li>• <i>ebgp-distance</i> —20.</li> <li>• <i>ibgp-distance</i> —200.</li> <li>• <i>local-distance</i> —220. Local-distance is the administrative distance used for aggregate discard routes when they are installed in the RIB.</li> </ul> <p>After you enter the value for the external administrative distance, you must enter the value for the administrative distance for the internal routes or/and the value for the administrative distance for the local routes depending on your requirement; so that the internal/local routes are also considered in the route administration.</p>
<p><b>log-neighbor-changes</b> [<b>disable</b>]</p> <p><b>Example:</b></p> <pre>switch(config-router-af)# log-neighbor-changes disable</pre>	<p>Generates a system message when this specific neighbor changes state.</p> <p>The <b>disable</b> option suppresses neighbor status changes messages for this specific neighbor.</p>

To tune BGP, use the following optional commands in neighbor configuration mode:

Command	Purpose
<p><b>description</b> <i>string</i></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# description main site</pre>	<p>Sets a descriptive string for this BGP peer. The string can be up to 80 alphanumeric characters.</p>
<p><b>low-memory exempt</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# low-memory exempt</pre>	<p>Exempts this BGP neighbor from a possible shutdown due to a low memory condition.</p>
<p><b>transport connection-mode passive</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# transport connection-mode passive</pre>	<p>Allows a passive connection setup only. This BGP speaker does not initiate a TCP connection to a BGP peer. You must manually reset the BGP sessions after configuring this command.</p>
<p><b>[no   default] remove-private-as [all   replace-as]</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# remove-private-as</pre>	<p>Removes private AS numbers from outbound route updates to an eBGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.</p> <p>The optional parameters are as follows:</p> <ul style="list-style-type: none"> <li>• <b>no</b> —Disables the command.</li> <li>• <b>default</b> —Moves the command to its default mode.</li> <li>• <b>all</b> —Removes all private-as numbers from the AS-path value.</li> <li>• <b>replace-as</b> —Replaces all private AS numbers with the replace-as AS-path value.</li> </ul> <p>See the <a href="#">Guidelines and Limitations for Advanced BGP, on page 14</a> section for additional information on this command.</p>
<p><b>update-source</b> <i>interface-type number</i></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor)# update-source ethernet 2/1</pre>	<p>Configures the BGP speaker to use the source IP address of the configured interface for BGP sessions to the peer. This command triggers an automatic notification and session reset for the BGP neighbor sessions. Single-hop iBGP peers support fast external fallover when <b>update-source</b> is configured.</p>

To tune BGP, use the following optional commands in neighbor address-family configuration mode:

Command	Purpose
<p><b>allowas in</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # allowas in</pre>	Allows routes that have their own AS in the AS path to be installed in the BRIB.
<p><b>default-originate [route-map map-name]</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # default-originate</pre>	Generates a default route to the BGP peer.
<p><b>disable-peer-as-check</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # disable-peer-as-check</pre>	Disables peer AS-number checking while the device advertises routes learned from one node to another node in the same AS path.
<p><b>filter-list list-name {in   out}</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # filter-list BGPFilter in</pre>	Applies an AS_path filter list to this BGP peer for inbound or outbound route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<p><b>prefix-list list-name {in   out}</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # prefix-list PrefixFilter in</pre>	Applies a prefix list to this BGP peer for inbound or outbound route updates. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<p><b>send-community</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # send-community</pre>	Sends the community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<p><b>send-community extended</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # send-community extended</pre>	Sends the extended community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
<p><b>suppress-inactive</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # suppress-inactive</pre>	Advertises the best (active) routes only to the BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.

Command	Purpose
<p><b>[no   default] as-override</b></p> <p><b>Example:</b></p> <pre>switch(config-router-neighbor-af) # as-override</pre>	<p><b>no</b> - (Optional) Disables the command.</p> <p><b>default</b> - (Optional) Moves the command to its default mode.</p> <p><b>as-override</b> - While sending updates to eBGP peer, replaces in the <i>path</i> attribute all occurrences of the peer's AS number with the local AS number.</p>

## Configuring Policy-Based Administrative Distance

You can configure a distance for external BGP (eBGP) and internal BGP (iBGP) routes that match a policy described in the configured route map. The distance configured in the route map is downloaded to the unicast RIB along with the matching routes. BGP uses the best path to determine the administrative distance when downloading next hops in the unicast RIB table. If there is no match or a deny clause in the policy, BGP uses the distance configured in the distance command or the default distance for routes.

The policy-based administrative distance feature is useful when there are two or more different routes to the same destination from two different routing protocols.

### Before you begin

You must enable BGP.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	switch# <b>configure terminal</b>	Enters global configuration mode.
<b>Step 2</b>	switch(config)# <b>ip prefix-list</b> <i>name seq number permit prefix-length</i>	Creates a prefix list to match IP packets or routes with the permit keyword.
<b>Step 3</b>	switch(config)# <b>route-map</b> <i>map-tag permit sequence-number</i>	Creates a route map and enters route-map configuration mode with the permit keyword. If the match criteria for the route is met in the policy, the packet is policy routed.
<b>Step 4</b>	switch(config-route-map)# <b>match ip address prefix-list</b> <i>prefix-list-name</i>	Matches IPv4 network routes based on a prefix list. The prefix-list name can be any alphanumeric string up to 63 characters.
<b>Step 5</b>	switch(config-route-map)# <b>set distance</b> <i>value1 value2 value3</i>	Specifies the administrative distance for interior BGP (iBGP) or exterior BGP (eBGP) routes and BGP routes originated in the local autonomous system. The range is from 1 to 255.  After you enter the value for the external administrative distance, you must enter the value for the administrative distance for the

	Command or Action	Purpose
		internal routes or/and the value for the administrative distance for the local routes depending on your requirement; so that the internal/local routes are also considered in the route administration.
<b>Step 6</b>	switch(config-route-map)# <b>exit</b>	Exits route-map configuration mode.
<b>Step 7</b>	switch(config)# <b>router bgp</b> <i>as-number</i>	Enters BGP mode and assigns the AS number to the local BGP speaker.
<b>Step 8</b>	switch(config-router)# <b>address-family</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } <b>unicast</b>	Enters address family configuration mode.
<b>Step 9</b>	switch(config-router-af)# <b>table-map</b> <i>map-name</i>	Configures the selective administrative distance for a route map for BGP routes before forwarding them to the RIB table. The table-map name can be any alphanumeric string up to 63 characters.  <b>Note</b> You can also configure the <b>table-map</b> command under the VRF address-family configuration mode.
<b>Step 10</b>	(Optional) switch(config-router-af)# <b>show forwarding distribution</b>	Displays forwarding information distribution.
<b>Step 11</b>	(Optional) switch(config)# <b>copy running-config startup-config</b>	Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

## Configuring Multiprotocol BGP

You can configure MP-BGP to support multiple address families, including IPv4 and IPv6 unicast and multicast routes.

### Before you begin

You must enable BGP.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.



	Command or Action	Purpose
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i>  <b>Example:</b> switch(config)# router bgp 65535 switch(config-router)#	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>neighbor</b> <i>ip-address remote-as as-number</i>  <b>Example:</b> switch(config-router)# neighbor 192.168.1.2 remote-as 65534 switch(config-router-neighbor)#	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.
<b>Step 4</b>	<b>address-family</b> { <i>ipv4   ipv6</i> } { <i>unicast   multicast</i> }  <b>Example:</b> switch(config-router-neighbor)# address-family ipv4 multicast switch(config-router-neighbor-af)#	Enters address family configuration mode.
<b>Step 5</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> switch(config-router-neighbor-af)# copy running-config startup-config	Saves this configuration change.

### Example

This example shows how to enable advertising and receiving IPv4 and IPv6 routes for multicast RPF for a neighbor:

```
switch# configure terminal
switch(config)# interface ethernet 2/1
switch(config-if)# ipv6 address 2001:0DB8::1
switch(config-if)# router bgp 65536
switch(config-router)# neighbor 192.168.1.2 remote-as 35537
switch(config-router-neighbor)# address-family ipv4 multicast
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# address-family ipv6 multicast
switch(config-router-neighbor-af)# copy running-config startup-config
```

## Configuring BMP

Beginning with Cisco NX-OS Release 7.0(3)I5(2), you can configure BMP on the device.

### Before you begin

You must enable BGP (see the [Enabling BGP](#) section).

## Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b> <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
<b>Step 2</b>	<b>router bgp as-number</b> <b>Example:</b> switch(config)# router bgp 200 switch(config-router)#	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
<b>Step 3</b>	<b>bmp server server-number</b> <b>Example:</b> switch(config-router)# bmp server 1	Configures the BMP server to which BGP should send information. The server number is used as a key. <b>Note</b> You can configure up to two BMP servers.
<b>Step 4</b>	<b>address ip-address port-number port-number</b> <b>Example:</b> switch(config-router)# address 10.1.1.1 port-number 2000	Configures the IPv4 or IPv6 address of the host and the port number on which the BMP speaker connects to the BMP server.
<b>Step 5</b>	<b>description string</b> <b>Example:</b> switch(config-router)# description BMPserver1	Configures the BMP server description. You can enter up to 256 alphanumeric characters.
<b>Step 6</b>	<b>initial-refresh { skip / delay time }</b> <b>Example:</b> switch(config-router)# initial-refresh delay 100	Configures the option to send a route refresh when BGP is converged and the BMP server connection is established later. The skip option specifies to not send a route refresh if the BMP server connection comes up later. The delay option specifies the time in seconds after which the route refresh should be sent. The range is from 30 to 720 seconds, and the default value is 30 seconds.
<b>Step 7</b>	<b>initial-delay time</b> <b>Example:</b> switch(config-router)# initial-delay 120	Configures the delay after which a connection is attempted to the BMP server. The range is from 30 to 720 seconds, and the default value is 45 seconds.
<b>Step 8</b>	<b>stats-reporting-period time</b> <b>Example:</b>	Configures the time interval in which the BMP server receives the statistics report from BGP

	Command or Action	Purpose
	<code>switch(config-router)# stats-reporting-period 50</code>	neighbors. The range is from 30 to 720 seconds, and the default is disabled.
<b>Step 9</b>	<code>shutdown</code> <b>Example:</b> <code>switch(config-router)# shutdown</code>	Disables the connection to the BMP server.
<b>Step 10</b>	<b>neighbor ip-address</b> <b>Example:</b> <code>switch(config-router)# neighbor 192.168.1.2 switch(config-router-neighbor)#</code>	Enters neighbor configuration mode for BGP routing and configures the neighbor IP address.
<b>Step 11</b>	<b>remote-as <i>as-number</i></b> <b>Example:</b> <code>switch(config-router-neighbor)# remote-as 65535</code>	Configures the AS number for a remote BGP peer.
<b>Step 12</b>	<b>bmp-activate-server <i>server-number</i></b> <b>Example:</b> <code>switch(config-router-neighbor)# bmp-activate-server 1</code>	Configures the BMP server to which a neighbor's information should be sent.
<b>Step 13</b>	(Optional) <b>show bgp bmp server</b> <i>[server-number] [detail]</i> <b>Example:</b> <code>switch(config-router-neighbor)# show bgp bmp server</code>	Displays BMP server information.
<b>Step 14</b>	(Optional) <b>copy running-config startup-config</b> <b>Example:</b> <code>switch(config-router-neighbor)# copy running-config startup-config</code>	Saves this configuration change.

## BGP Local Route Leaking

### About BGP Local Route Leaking

Beginning with release 9.3(1), NX-OS BGP supports leaking imported VPN routes between:

- The VPN route table and default VRF route table
- The VPN route table and VRF-lite route table
- Border leaf (BL) switch route tables for leaf-to-leaf connectivity

This feature enables the propagation of routes between the route tables. You can control route leaking for a VRF by configuring an import or export map, which now includes an option to allow or prevent incoming locally originated routes and specify whether they should be advertised. Local route leaking is bidirectional, so routes that are locally originated can be leaked from a VRF out to a BGP VPN, and routes that are imported from the BGP VPN can be leaked into a VRF.



**Note** NX-OS supports a similar feature called centralized route leaking. For information, see [Configuring Layer 3 Virtualization](#).

## Guidelines and Limitations for BGP Local Route Leaking

The following are the guidelines and limitations for the BGP local route leaking feature:

- The following Cisco hardware supports this feature:
  - Cisco Nexus 9332C, 9364C, 9300-EX, 9300-FX/FXP/FX2/FX3, and 9300-GX platform switches, and Cisco Nexus 9500 platform switches with 9700-EX/FX line cards
  - Cisco Nexus 9500 platform switches with -R line cards
- When using route-targets, the same route-targets might have duplicated paths pointing to the same remote path, which can negatively impact the switch's memory and performance. Be careful when using route targets.
- Be careful when using local route leaking in a leaf-to-leaf case, where border-leaf routers (BLs) are leaking between the same VRFs. This scenario is more prone to routing loops. We recommend using inbound route-maps to exclude the imported routes from other BLs.
- After a remote path gets withdrawn, it can take up to 20 seconds more for BGP to completely clean up the path.

## Configuring Routes Imported from a VPN to Leak into the Default VRF

You can configure a VRF to allow routes that are imported from a BGP VPN to be exported to the default VRF. Use this procedure for a non-default VRF.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>config terminal</b>  <b>Example:</b> <pre>switch-1# config terminal Enter configuration commands, one per line. End with CNTL/Z. switch-1(config)#</pre>	Enters global configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<b>vrf context</b> <i>vrf-name</i>  <b>Example:</b> switch-1(config)# <b>vrf context</b> vpn1 switch-1(config-vrf)#	Creates a new VRF and enters VRF configuration mode. The name can be any case-sensitive, alphanumeric string up to 32 characters.
<b>Step 3</b>	<b>address-family</b> <i>address-family sub family</i>  <b>Example:</b> switch-1(config-vrf)# <b>address-family</b> ipv4 <b>unicast</b> switch-1(config-vrf-af-ipv4)#	
<b>Step 4</b>	<b>export vrf default</b> [ <i>prefix-limit</i> ] <b>map</b> <i>route-map</i> <b>allow-vpn</b>  <b>Example:</b> switch-1(config-vrf-af-ipv4)# <b>export vrf</b> <b>default map</b> vpnmap1 <b>allow-vpn</b> switch-1(config-vrf-af-ipv4)#	Configures the current VRF to allow routes that are imported from a BGP VPN to be exported to the default VRF.

## Configuring Routes Leaked from the Default-VRF to Export to a VPN

You can configure a VRF to allow routes leaked from the default VRF to be exported to a BGP VPN. Use this procedure for a non-default VRF.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>config terminal</b>  <b>Example:</b> switch-1# <b>config</b> terminal Enter configuration commands, one per line. End with CNTL/Z. switch-1(config)#	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context</b> <i>vrf-name</i>  <b>Example:</b> switch-1(config)# <b>vrf context</b> vpn1 switch-1(config-vrf)#	Creates a new VRF and enters VRF configuration mode. The name can be any case-sensitive, alphanumeric string up to 32 characters.
<b>Step 3</b>	<b>address-family</b> <i>address-family sub family</i>  <b>Example:</b> switch-1(config-vrf)# <b>address-family</b> ipv4 <b>unicast</b> switch-1(config-vrf-af-ipv4)#	

	Command or Action	Purpose
<b>Step 4</b>	<b>import vrf default</b> [ <i>prefix-limit</i> ] <b>map route-map advertise-vpn</b> <b>Example:</b> <pre>switch-1(config-vrf-af-ipv4)# import vrf map vpnmap1 advertise-vpn switch-1(config-vrf-af-ipv4)#</pre>	Configures the current VRF to allow routes imported from the default VRF to be exported to a BGP VPN.

## Configuring Routes Imported from a VPN to Export to a VRF

You can configure a VRF to allow VPN imported routes to be exported to another VRF. Use this procedure for non-default VRFs.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>config terminal</b> <b>Example:</b> <pre>switch-1# config terminal Enter configuration commands, one per line. End with CNTL/Z. switch-1(config)#</pre>	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context vrf-name</b> <b>Example:</b> <pre>switch-1(config)# vrf context vpn1 switch-1(config-vrf)#</pre>	Creates a new VRF and enters VRF configuration mode. The name can be any case-sensitive, alphanumeric string up to 32 characters.
<b>Step 3</b>	<b>address-family address-family sub family</b> <b>Example:</b> <pre>switch-1(config-vrf)# address-family ipv4 unicast switch-1(config-vrf-af-ipv4)#</pre>	
<b>Step 4</b>	<b>export vrf allow-vpn</b> <b>Example:</b> <pre>switch-1(config-vrf-af-ipv4)# export vrf allow-vpn nxosv2(config-vrf-af-ipv4)#</pre>	Configures a VRF to allow routes imported from a BGP VPM to be exported to a non-default VRF.

## Configuring Routes Imported from a VRF to Export to a VPN

You can configure a VRF to allow routes imported from another VRF to be exported to a BGP VPN. Use this procedure for non-default VRFs.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>config terminal</b>  <b>Example:</b> <pre>switch-1# config terminal Enter configuration commands, one per line. End with CNTL/Z. switch-1 (config) #</pre>	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context</b> <i>vrf-name</i>  <b>Example:</b> <pre>switch-1 (config) # vrf context vpn1 switch-1 (config-vrf) #</pre>	Creates a new VRF and enters VRF configuration mode. The name can be any case-sensitive, alphanumeric string up to 32 characters.
<b>Step 3</b>	<b>address-family</b> <i>address-family sub family</i>  <b>Example:</b> <pre>switch-1 (config-vrf) # address-family ipv4 unicast switch-1 (config-vrf-af-ipv4) #</pre>	
<b>Step 4</b>	<b>import vrf advertise-vpn</b>  <b>Example:</b> <pre>switch-1 (config-vrf-af-ipv4) # import vrf advertise-vpn nxosv2 (config-vrf-af-ipv4) #</pre>	Configures the current VRF to allow routes that are imported from another VRF to be exported to a BGP VPN.

## Configuration Examples

The following show sample configurations for the BGP local route leaking feature.

### Configuring BGP VPN to Default VRF Reachability

In this example, the configuration enables route re-importation through an intermediate VRF, called VRF\_A, which is between the VPN and the default VRF.

```
vrf context VRF_A
  address-family ipv4 unicast
  route-target both auto evpn
  import vrf default map MAP_1 advertise-vpn
  export vrf default map MAP_1 allow-vpn
```

Route re-importation is enabled by using the **advertise-vpn** option to control importing routes from the VPN into VRF\_A, and **allow-vpn** for the export map to control exporting VPN-imported routes from VRF\_A out to the default VRF. Configuration occurs on the intermediate VRF.

### Configuring VPN to VRF-Lite Reachability

In this example, the VPN connects to a tenant VRF, called VRF\_A. VRF\_A connects a VRF-Lite, called VRF-B. The configuration enables VPN imported routes to be leaked from VRF\_A to VRF\_B.

```
vrf context VRF_A
  address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn
  route-target import 3:3
  route-target export 2:2
  import vrf advertise-vpn
  export vrf allow-vpn
vrf context VRF_B
  address-family ipv4 unicast
  route-target both 1:1
  route-target import 2:2
  route-target export 3:3
```

Route leaking between the two is enabled by using the **allow-vpn** in an export map configured in VRF\_A (tenant). The export map in VRF\_A allows route imported from the VPN to be leaked into the VRF\_B. Routes processed by the export map have the **route-mapexport** and **export-map** attributes added to the route's set of route targets. The import map uses **advertise-vpn** which enables routes that are imported from the VRF-Lite for be exported out to the VPN.

After a route leaks between the VRFs, it is reoriginated and its route targets are replaced by the route target export and export map attributes specified by the new VRF's configuration.

### Leaf-to-Leaf Reachability

In this example, two VPNs exist and two VRFs exist. VPN\_1 is connected to VRF\_A and VPN\_2 is connected to VRF\_B. Both VRFs are route distinguishers (RDs).

```
vrf context VRF_A
  address-family ipv4 unicast
  route-target both auto
  route-target both auto evpn
  route-target import 3:3
  route-target export 2:2
  import vrf advertise-vpn
  export vrf allow-vpn
vrf context VRF_B
  address-family ipv4 unicast
  route-target both 1:1
  route-target import 2:2
  route-target export 3:3
  import vrf advertise-vpn
  export vrf allow-vpn
```

Route leaking between the two is enabled by **allow-vpn** in an export map configured in VRF\_A and VRF\_B. VPN imported routes have **route-mapexport** and **export-map** attributes added to the route's set of route targets. An import map map uses the advertise-vpn option which enables routes that are imported from each VRF to be exported out to the VPN.

After a route leaks between the VRFs, it is reoriginated and its route targets are replaced by the route target export and export map attributes specified by the new VRF's configuration.



### Leaf-to-Leaf with Loop Prevention

In the leaf-to-leaf configuration, you can inadvertently cause loops between the BLs that are leaking between the same VRFs unless you are careful with your route maps:

- You can use an inbound route map in each BL to deny updates from every other BL.
- If a BL originates a route, a standard community can be applied, which enables other BLs to accept the routes. This community is then stripped in the receiving BL.

In the following example, VTEPs 3.3.3.3, 4.4.4.4 and 5.5.5.5 are the BLs.

```
ip prefix-list BL_PREFIX_LIST seq 5 permit 3.3.3.3/32
ip prefix-list BL_PREFIX_LIST seq 10 permit 4.4.4.4/32
ip prefix-list BL_PREFIX_LIST seq 20 permit 5.5.5.5/32
ip community-list standard BL_COMMUNITY seq 10 permit 123:123
route-map INBOUND_MAP permit 5
  match community BL_COMMUNITY
  set community none
route-map INBOUND_MAP deny 10
  match ip next-hop prefix-list BL_PREFIX_LIST
route-map INBOUND_MAP permit 20
route-map OUTBOUND_SET_COMM permit 10
  match evpn route-type 2 mac-ip
  set community 123:123
route-map SET_COMM permit 10
  set community 123:123
route-map allow permit 10

vrf context vni100
vni 100
address-family ipv4 unicast
  route-target import 2:2
  route-target export 1:1
  route-target both auto
  route-target both auto evpn
import vrf advertise-vpn
export vrf allow-vpn

vrf context vni200
vni 200
address-family ipv4 unicast
  route-target import 1:1
  route-target export 2:2
  route-target both auto
  route-target both auto evpn
import vrf advertise-vpn
export vrf allow-vpn

router bgp 100
template peer rr
remote-as 100
update-source loopback0
address-family l2vpn evpn
  send-community
  send-community extended
  route-map INBOUND_MAP in
  route-map OUTBOUND_SET_COMM out
neighbor 101.101.101.101
  inherit peer rr
neighbor 102.102.102.102
  inherit peer rr
vrf vni100
  address-family ipv4 unicast
```

```

network 3.3.3.100/32 route-map SET_COMM
vrf vni200
address-family ipv4 unicast
network 3.3.3.200/32 route-map SET_COMM

```

In this example, the tenant VRFs for the border leaf (BL) router can leak traffic by enabling extra import export flows, and the route targets in the route maps determine where the routes are imported from or exported to.

### Multipath in a VRF

In this example, a VPN has multiple incoming paths. This configuration enables route leaking through an intermediate VRF, called VRF\_A, which is between the VPN and another VRF, named VRF\_B. Assume that multipathing is enabled in VRF\_A.

```

vrf context VRF_A
address-family ipv4 unicast
route-target both auto evpn
route-target export 3:3
export vrf allow-vpn
vrf context VRF_B
address-family ipv4 unicast
route-target import 3:3

```

Route leaking is enabled by **allow-vpn** in the export map configured in VRF\_A. When two paths for a given prefix are learnt from a VPN and imported into VRF\_A, two different paths exist in VRF\_B with the same source RD (VRF\_A's local RD). Each route is distinguished by the original source RD (remote RD).

### Path Duplication

In this example, the configuration enables a single VPN path to be imported into both VRF\_A and VRF\_B. Because VRF\_A is configured with **export vrf allow-vpn**, VRF\_A also leaks its routes into VRF\_B. VRF\_B then has two paths with same source RD (VRF\_A's local RD), each one distinguished by the original source RD (remote RD).

```

vrf context VRF_A
address-family ipv4 unicast
route-target import 1:1 evpn
route-target export 1:1 evpn
route-target export 2:2
export vrf allow-vpn
vrf context VRF_B
address-family ipv4 unicast
route-target import 1:1 evpn
route-target import 2:2

```

This configuration creates a situation in which multipathing does not exist.

## Displaying BGP Local Route Leaking Information

The following show commands contain information for the BGP local route leaking feature.

Command	Action
<b>show bgp vrf <i>vrf-name</i> process</b>	For a default or non-default VRF, shows the enabled state (Yes or No) of the <b>import advertise-vpn</b> and <b>export allow-vpn</b> options.

Command	Action
<code>show bgp vrf <i>vrf-name</i> ipv4 unicast <i>prefix</i></code>	Shows information about imported paths, including a list of destinations a route has been imported from.

# BGP Graceful Shutdown

## About BGP Graceful Shutdown

Beginning with release 9.3(1), BGP supports the graceful shutdown feature. This BGP feature works with the BGP **shutdown** command to:

- Dramatically decrease the network convergence time when a router or link is taken offline.
- Reduce or eliminate dropped packets that are in transit when a router or link is taken offline.

Despite the name, BGP graceful shutdown does not actually cause a shutdown. Instead, it alerts connected routers that a router or link will be going down soon.

The graceful shutdown feature uses the GRACEFUL\_SHUTDOWN well-known community (0xFFFF0000 or 65535:0), which is identified by IANA and the IETF through RFC 8326. This well-known community can be attached to any routes, and it is processed like any other attribute of a route.

Because this feature announces that a router or link will be going down, the feature is useful in preparation of maintenance windows or planned outages. Use this feature before shutting down BGP to limit the impact on traffic.

## Graceful Shutdown Aware and Activate

BGP routers can control the preference of all routes with the GRACEFUL\_SHUTDOWN community through the concept of GRACEFUL SHUTDOWN awareness. Graceful shutdown awareness is enabled by default, which enables the receiving peers to deprefer incoming routes carrying the GRACEFUL\_SHUTDOWN community. Although not a typical use case, you can disable and reenable graceful shutdown awareness through the **graceful-shutdown aware** command.

Graceful shutdown aware is applicable only at the BGP global context. For information about contexts, see [Graceful Shutdown Contexts, on page 73](#). The aware option operates with another option, the **activate** option, which you can assign to a route map for more granular control over graceful shutdown routes.

### Interaction of the Graceful Shutdown Aware and Activate Options

When a graceful shutdown is activated, the GRACEFUL\_SHUTDOWN community is appended to route updates only when you specify the **activate** keyword. At this point, new route updates that contain the community are generated and transmitted. When the **graceful-shutdown aware** command is configured, all routers that receive the community then deprefer (lower the route preference of) the routes in the update. Without the **graceful-shutdown aware** command, BGP does not deprefer routes with the GRACEFUL\_SHUTDOWN community.

After the feature is activated and the routers are aware of graceful shutdown, BGP still considers the routes with the GRACEFUL\_SHUTDOWN community as valid. However, those routes are given the lowest priority

in the best-path calculation. If alternate paths are available, new best paths are chosen, and convergence occurs to accommodate the router or link that will soon go down.

## Graceful Shutdown Contexts

BGP graceful shutdown feature has two contexts that determine what the feature affects and what functionality is available.

Context	Affects	Commands
Global	The entire switch and all routes processed by it. For example, readvertise all routes with the GRACEFUL_SHUTDOWN community.	<b>graceful-shutdown activate</b> <b>[route-map route-map]</b> <b>graceful-shutdown aware</b>
Peer	A BGP peer or a link between neighbors. For example, advertise only one link between peers with GRACEFUL_SHUTDOWN community.	<b>graceful-shutdown activate</b> <b>[route-map route-map]</b>

## Graceful Shutdown with Route Maps

Graceful shutdown works with the route policy manager (RPM) feature to control how the switch's BGP router transmits and receives routes with the GRACEFUL\_SHUTDOWN community. Route maps can process route updates with the community in the inbound and outbound directions. Typically, route maps are not required. However, if needed, you can use them to customize the control of graceful shutdown routes.

### Normal Inbound Route Maps

Normal inbound route maps affect routes that are incoming to the BGP router. Normal inbound route maps are not commonly used with the graceful shutdown feature because routers are aware of graceful shutdown by default.

Cisco Nexus switches running Cisco NX-OS Release 9.3(1) and later do not require an inbound route map for the graceful shutdown feature. Cisco NX-OS Release 9.3(1) and later have implicit inbound route maps that automatically deprefer any routes that have the GRACEFUL\_SHUTDOWN community if the BGP router is graceful shutdown aware.

Normal inbound route maps can be configured to match against the well-known GRACEFUL\_SHUTDOWN community. Although these inbound route maps are not common, there are some cases where they are used:

- If switches are running a Cisco NX-OS release earlier than 9.3(1), they do not have the implicit inbound route map present in NX-OS 9.3(1). To use the graceful shutdown feature on these switches, you must create a graceful shutdown inbound route map. The route map must match inbound routes with the well-known GRACEFUL\_SHUTDOWN community, permit them, and deprefer them. If an inbound route map is needed, create it on the BGP peer that is running a version of NX-OS earlier than 9.3(1) and is receiving the graceful shutdown routes.

- If you want to disable graceful shutdown aware, but still want the router to act on incoming routes with GRACEFUL\_SHUTDOWN community from some BGP neighbors, you can configure an inbound route map under the respective peers.

### Normal Outbound Route Maps

Normal outbound route maps control forwarding the routes that a BGP router sends. Normal outbound route maps can affect the graceful shutdown feature. For example, you can configure an outbound route map to match on the GRACEFUL\_SHUTDOWN community and set attributes, and it takes precedence over any graceful shutdown outbound route maps.

### Graceful Shutdown Outbound Route Maps

Outbound Graceful shutdown route maps are specific type of outbound route map for the graceful shutdown feature. They are optional, but they are useful when you already have a community list that is associated with a route map. The typical graceful shutdown outbound route map contains only `set` clauses to set or modify certain attributes.

You can use outbound route maps in the following ways:

- For customers that already have existing outbound route maps, you can add a new entry with a higher sequence number, match on the GRACEFUL\_SHUTDOWN well-known community, and add any attributes that you want.
- You can also use a graceful shutdown outbound route map with the **graceful-shutdown activate route-map *name*** option. This is the typical use case.

This route map requires no match clauses, so the route map matches on all routes being sent to the neighbor.

### Route Map Precedence

When multiple route maps are present on the same router, the following order of precedence is applied to determine how routes with the community are processed: Consider the following example. Assume you have a standard outbound route map name Red that sets a local-preference of 60. Also, assume you have a peer graceful-shutdown route map that is named Blue that sets local-pref to 30. When the route update is processed, the local preference will be set to 60 because Red overwrites Blue.

- Normal outbound route maps take precedence over peer graceful shutdown maps.
- Peer graceful shutdown maps take precedence over global graceful shutdown maps.

## Guidelines and Limitations

The following are limitations and guidelines for BGP global shutdown:

- Graceful shutdown feature can only help avoid traffic loss when alternative routes exist in the network for the affected routers. If the router has no alternate routes, routes carrying the GRACEFUL\_SHUTDOWN community are the only ones available, and therefore, are used in the best-path calculation. This situation defeats the purpose of the feature.
- Configuring a BGP send community is required to send the GRACEFUL\_SHUTDOWN community.
- For route maps:

- When global route maps and neighbor route maps are configured, the per-neighbor route maps take precedence.
- Outbound route maps take precedence over any global route maps configured for graceful shutdown.
- Outbound route maps take precedence over any peer route maps configured for graceful shutdown.
- To add the graceful shutdown functionality to legacy (existing) inbound route maps, follow this order:
  1. Add the graceful shutdown match clause to the top of the route map by setting a low sequence number for the clause (for example, sequence number 0).
  2. Add a continue statement after the graceful shutdown clause. If you omit the continue statement, route-map processing stops when it matches the graceful shutdown clause, any other clauses with higher sequence numbers (for example, 1 and higher) are not processed.

## Graceful Shutdown Task Overview

To use the graceful shutdown feature, you typically enable graceful-shutdown aware on all Cisco Nexus switches and leave the feature enabled. When a BGP router must be taken offline, you configure graceful-shutdown activate on it.

The following details document the best practice for using the graceful shutdown feature.

To bring the router or link down:

1. Configure the Graceful Shutdown feature.
2. Watch the neighbor for the best path.
3. When the best path is recalculated, issue the **shutdown** command to disable BGP.
4. Perform the work that required you to shut down the router or link.

To bring the router or link back online:

1. When you finish the work that required the shutdown, reenable BGP (**no shutdown**).
2. Disable the graceful shutdown feature (**no graceful-shutdown activate** in config router mode).

## Configuring Graceful Shutdown on a Link

This task enables you to configure graceful shutdown on a specific link between two BGP routers.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
Step 1	<b>config terminal</b> <b>Example:</b>	Enters global configuration mode.

	Command or Action	Purpose
	switch-1# <b>configure terminal</b> switch-1(config)#	
<b>Step 2</b>	<b>router bgp</b> <i>autonomous-system-number</i>  <b>Example:</b> switch-1(config)# <b>router bgp 110</b> switch-1(config-router)#	Enters router configuration mode to create or configure a BGP routing process.
<b>Step 3</b>	<b>neighbor</b> { <i>ipv4-address ipv6-address</i> } <b>remote-as</b> <i>as-number</i>  <b>Example:</b> switch-1(config-router)# <b>neighbor</b> <b>10.0.0.3 remote-as 200</b> switch-1(config-router-neighbor)#	Configures the autonomous system (AS) to which the neighbor belongs.
<b>Step 4</b>	<b>graceful-shutdown activate</b> [ <b>route-map</b> <i>map-name</i> ]  <b>Example:</b> switch-1(config-router-neighbor)# <b>graceful-shutdown activate route-map</b> <b>gshutPeer</b> switch-1(config-router-neighbor)#	Configures graceful shutdown on the link to the neighbor. Also, advertises the routes with the well-known GRACEFUL_SHUTDOWN community and applies the route map to the outbound route updates.  The routes are advertised with the graceful-shutdown community by default. In this example, routes are advertised to the neighbor with the Graceful-shutdown community with a route-map named gshutPeer.  The devices receiving the gshut community look at the communities of the route and optionally use the communities to apply routing policy.

## Filtering BGP Routes and Setting Local Preference Based On GRACEFUL\_SHUTDOWN Communities

Switches that are not yet running 9.3(1) do not have an inbound route map that matches against the GRACEFUL\_SHUTDOWN community name. Therefore, they have no way of identifying and deprefering the correct routes.

For switches running a release of NX-OS that is earlier than 9.3(1), you must configure an inbound route map that matches on the community value for graceful shutdown (65535:0) and deprefers routes.

If your switch is running 9.3(1) or later, you do not need to configure an inbound route map.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b>	Enters global configuration mode.

	Command or Action	Purpose
	switch-1# <b>configure terminal</b> switch-1(config)#	
<b>Step 2</b>	<b>ip community list standard</b> <i>community-list-name seq sequence-number { permit   deny } value</i>  <b>Example:</b> switch-1(config)# <b>ip community-list standard GSHUT seq 10 permit 65535:0</b> switch-1(config)#	Configures a community list and permits or denies routes that have the well-known graceful shutdown community value.
<b>Step 3</b>	<b>route map map-tag {deny   permit} sequence-number</b>  <b>Example:</b> switch-1(config)# <b>route-map RM_GSHUT permit 10</b> switch-1(config-route-map)#	Configures a route map as sequence 10 and permits routes that have the GRACEFUL_SHUTDOWN community.
<b>Step 4</b>	<b>match community community-list-name</b>  <b>Example:</b> switch-1(config-route-map)# <b>match community GSHUT</b> switch-1(config-route-map)#	Configures that routes that match the IP community list GSHUT are processed by Route Policy Manager (RPM).
<b>Step 5</b>	<b>set local-preference local-pref-value</b>  <b>Example:</b> switch-1(config-route-map)# <b>set local-preference 10</b> switch-1(config-route-map)#	Configures that the routes that match the IP community list GSHUT will be given a specified local preference.
<b>Step 6</b>	<b>exit</b>  <b>Example:</b> switch-1(config-route-map)# <b>exit</b> switch-1(config)#	Leaves route map configuration and returns to global configuration mode.
<b>Step 7</b>	<b>router bgp community-list-name</b>  <b>Example:</b> switch-1(config)# <b>router bgp 100</b> switch-1(config-router)#	Enters router configuration mode and creates a BGP instance.
<b>Step 8</b>	<b>neighbor { ipv4-address ipv6-address }</b>  <b>Example:</b> switch-1(config-router)# <b>neighbor 10.0.0.3</b> switch-1(config-router-neighbor)#	Enters route BGP neighbor mode for a specified neighbor.
<b>Step 9</b>	<b>address-family { address-family sub family }</b>  <b>Example:</b>	Puts the neighbor into address family (AF) configuration mode.



	Command or Action	Purpose
	<pre>nxosv2 (config-router-neighbor) # address-family ipv4 unicast nxosv2 (config-router-neighbor-af) #</pre>	
<b>Step 10</b>	<p><b>send community</b></p> <p><b>Example:</b></p> <pre>nxosv2 (config-router-neighbor-af) # send-community nxosv2 (config-router-neighbor-af) #</pre>	Enables BGP community exchange with the neighbor.
<b>Step 11</b>	<p><b>route map <i>map-tag</i> in</b></p> <p><b>Example:</b></p> <pre>nxosv2 (config-router-neighbor-af) # route-map RM_GSHUT in nxosv2 (config-router-neighbor-af) #</pre>	Applies the route map to incoming routes from the neighbor. In this example, the route map that is named RM_GSHUT permits routes with the GRACEFUL_SHUTDOWN community from the neighbor.

## Configuring Graceful Shutdown for All BGP Neighbors

You can manually apply the GRACEFUL\_SHUTDOWN well-known community to all the neighbors of a graceful shutdown initiator.

You can configure graceful shutdown at the global level for all BGP neighbors.

### Before you begin

If you have not already enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>configure terminal</b></p> <p><b>Example:</b></p> <pre>switch-1# configure terminal switch-1 (config) #</pre>	Enters global configuration mode.
<b>Step 2</b>	<p><b>router bgp <i>autonomous-system-number</i></b></p> <p><b>Example:</b></p> <pre>switch-1 (config) # router bgp 110 switch-1 (config-router) #</pre>	Enters router configuration mode to create or configure a BGP routing process.
<b>Step 3</b>	<p><b>graceful-shutdown activate [<i>route-map map-name</i>]</b></p> <p><b>Example:</b></p> <pre>switch-1 (config-router-neighbor) # graceful-shutdown activate route-map gshutPeer switch-1 (config-router-neighbor) #</pre>	<p>Configures graceful shutdown route map for the links to all neighbors. Also, advertises all routes with the well-known GRACEFUL_SHUTDOWN community and applies the route map to the outbound route updates.</p> <p>The routes are advertised with the GRACEFUL_SHUTDOWN community by</p>

	Command or Action	Purpose
		<p>default. In this example, routes are advertised to all neighbors with the community with a route-map named gshutPeer. The route map should contain only set clauses.</p> <p>The devices receiving the GRACEFUL_SHUTDOWN community look at the communities of the route and optionally use the communities to apply routing policy.</p>

## Controlling the Preference for All Routes with the GRACEFUL\_SHUTDOWN Community

Cisco NX-OS enables lowering the preference of incoming routes that have the GRACEFUL\_SHUTDOWN community. When **graceful shutdown aware** is enabled, BGP considers routes carrying the community as the lowest preference during best path calculation. By default, lowering the preference is enabled, but you can selectively disable this option.

Whenever you enable or disable this option, you trigger a BGP best-path calculation. This option gives you the flexibility to control the behavior of the BGP best-path calculation for the graceful shutdown well-known community.

### Before you begin

If you have not enabled BGP, enable it now (**feature bgp**).

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<p><b>configure terminal</b></p> <p><b>Example:</b></p> <pre>switch-1(config)# <b>config terminal</b> switch-1(config)#</pre>	Enters global configuration mode.
<b>Step 2</b>	<p><b>router bgp <i>autonomous-system</i></b></p> <p><b>Example:</b></p> <pre>switch-1(config)# <b>router bgp 100</b> switch-1(config-router)#</pre>	Enters router configuration mode and configures a BGP routing process.
<b>Step 3</b>	<p>(Optional) <b>no graceful-shutdown aware</b></p> <p><b>Example:</b></p> <pre>switch-1(config-router)# <b>no graceful-shutdown aware</b> switch-1(config-router)#</pre>	For this BGP router, do not give lower preference for all routes that have the GRACEFUL_SHUTDOWN community. The default action is to deprefer routes when the graceful shutdown aware feature is disabled, so using the <b>no</b> form of the command is optional for not deprefering graceful shutdown routes.

## Preventing Sending the GRACEFUL\_SHUTDOWN Community to a Peer

If you no longer need the GRACEFUL\_SHUTDOWN community that is appended as a route attribute to outbound route updates, you can remove the community, which no longer sends it to a specified neighbor. One use case would be when a router is at an autonomous system boundary, and you do not want the graceful shutdown functionality to propagate outside of an autonomous system boundary.

To prevent sending the GRACEFUL\_SHUTDOWN to a peer, you can disable the send community option or strip the community from the outbound route map.

Choose either of the following methods:

- Disable the send-community in the running config.

### Example:

```
nxosv2(config-router-neighbor-af)# no send-community standard
nxosv2(config-router-neighbor-af)#
```

If you use this option, the GRACEFUL\_SHUTDOWN community is still received by the switch, but it is not sent to the downstream neighbor through the outbound route map. All standard communities are not sent either.

- Delete the GRACEFUL\_SHUTDOWN community through an outbound route map by following these steps:
  1. Create an IP community list matches the GRACEFUL\_SHUTDOWN community.
  2. Create an outbound route map to match against the GRACEFUL\_SHUTDOWN community.
  3. Use a **set community-list delete** clause to strip GRACEFUL\_SHUTDOWN community.

If you use this option, the community list matches and permits the GRACEFUL\_SHUTDOWN community, then the outbound route map matches against the community and then deletes it from the outbound route map. All other communities pass through the outbound route map without issue.

## Displaying Graceful Shutdown Information

Information about the graceful shutdown feature is available through the following **show** commands.

Command	Action
<b>show ip bgp community-list graceful-shutdown</b>	Shows all entries in the BGP routing table that have the GRACEFUL_SHUTDOWN community.
<b>show running-config bgp</b>	Shows the running BGP configuration.
<b>show running-config bgp all</b>	Shows all information for the running BGP configuration including information about the graceful shutdown feature.

Command	Action
<code>show bgp address-family neighbors neighbor-address</code>	<p>When the feature is configured for the peer, shows the following:</p> <ul style="list-style-type: none"> <li>• The state of the graceful-shutdown-activate feature for the specified neighbor</li> <li>• The name of any graceful shutdown route map configured for the specified neighbor</li> </ul>
<code>show bgp process</code>	<p>Shows different information depending on the context.</p> <p>When the graceful-shutdown-activate option is configured in peer context, shows the enabled or disabled state for the feature through <code>graceful-shutdown-active</code>.</p> <p>When the graceful-shutdown-activate option is configured in global context and has a graceful-shutdown route map, shows the enabled state of the feature through the following:</p> <ul style="list-style-type: none"> <li>• <code>graceful-shutdown-active</code></li> <li>• <code>graceful-shutdown-aware</code></li> <li>• <code>graceful-shutdown route-map</code></li> </ul>
<code>show ip bgp address</code>	<p>For the specified address, shows the BGP routing table information, including the following:</p> <ul style="list-style-type: none"> <li>• The state of the specified address as the best path</li> <li>• Whether the specified address is part of the GRACEFUL_SHUTDOWN community</li> </ul>

## Graceful Shutdown Configuration Examples

These examples show some configurations for using the graceful shutdown feature.

### Configuring Graceful Shutdown for a BGP Link

The following example shows how to configure graceful shutdown while setting a local preference and a community:

- Configuring graceful shutdown activate for the link to the specified neighbor
- Adding the GRACEFUL\_SHUTDOWN community to the routes
- Setting a route map named `gshutPeer` with only set clauses for outbound routes with the community.

```
router bgp 100
  neighbor 20.0.0.3 remote-as 200
    graceful-shutdown activate route-map gshutPeer
  address-family ipv4 unicast
```

```

        send-community

route-map gshutPeer permit 10
    set local-preference 0
    set community 200:30

```

### Configuring Graceful Shutdown for All-Neighbor BGP Links

The following example shows:

- Configuring graceful shutdown activate for all the links connecting the local router and all its neighbors.
- Adding the GRACEFUL\_SHUTDOWN community to the routes.
- Setting a route map that is named gshutAall with only set clauses for all outbound routes.

```

router bgp 200
    graceful-shutdown activate route-map gshutAll

route-map gshutAll permit 10
    set as-path prepend 10 100 110
    set community 100:80

route-map Red permit 10
    set local-pref 20

router bgp 100
    graceful-shutdown activate route-map gshutAll
    router-id 2.2.2.2
    address-family ipv4 unicast
    network 2.2.2.2/32
    neighbor 1.1.1.1 remote-as 100
    update-source loopback0
    address-family ipv4 unicast
    send-community
    neighbor 20.0.0.3 remote-as 200
    address-family ipv4 unicast
    send-community
    route-map Red out

```

In this example, the `gshutAll` route-map takes effect for neighbor 1.1.1.1, but not neighbor 20.0.0.3, because the outbound route-map `Red` configured under neighbor 20.0.0.3 takes precedence instead.

### Configuring Graceful Shutdown Under a Peer-Template

This example configures the graceful shutdown feature under a peer-session template, which is inherited by a neighbor.

```

router bgp 200
    template peer-session p1
        graceful-shutdown activate route-map gshut_out
    neighbor 1.1.1.1 remote-as 100
    inherit peer-session p1
    address-family ipv4 unicast
    send-community

```

### Filtering BGP Routes and Setting Local Preference Based on GRACEFUL\_SHUTDOWN Community Using and Inbound Route Map

This example shows how to use a community list to filter the incoming routes that have the GRACEFUL\_SHUTDOWN community. This configuration is useful for legacy switches that are not running Cisco NX-OS 9.3(1) as a minimum version.

The following example shows:

- An IP Community List that permits routes that have the GRACEFUL\_SHUTDOWN community.
- A route map that is named RM\_GSHUT that permits routes based on a standard community list named GSHUT.
- The route map also sets the preference for the routes it processes to 0 so that those routes are given lower preference for best path calculation when the router goes offline. The route map is applied to incoming IPv4 routes from the neighbor (20.0.0.2).

```
ip community-list standard GSHUT permit 65535:0

route-map RM_GSHUT permit 10
  match community GSHUT
  set local-preference 0

router bgp 200
  neighbor 20.0.0.2 remote-as 100
  address-family ipv4 unicast
    send-community
  route-map RM_GSHUT in
```

## Configuring a Graceful Restart

You can configure a graceful restart and enable the graceful restart helper feature for BGP.



**Note** Cisco NX-OS Release 10.1(1) supports a higher number of BFD sessions. If BGP sessions are associated with BFD, the BGP **restart-time** may need to be increased to maintain peer connection during ISSU.

### Before you begin

You must enable BGP (see the "Enabling BGP" section).

Create the VRFs.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters configuration mode.

	Command or Action	Purpose
<b>Step 2</b>	<b>router bgp</b> <i>as-number</i> <b>Example:</b> <pre>switch(config)# router bgp 65535 switch(config-router)#</pre>	Creates a new BGP process with the configured autonomous system number.
<b>Step 3</b>	(Optional) <b>timers prefix-peer-timeout</b> <i>timeout</i> <b>Example:</b> <pre>switch(config-router)# timers prefix-peer-timeout 20</pre>	Configures the timeout value (in seconds) for BGP prefix peers. The default value is 90 seconds.  <b>Note</b> This command is supported beginning with Cisco NX-OS Release 9.3(3).
<b>Step 4</b>	<b>graceful-restart</b> <b>Example:</b> <pre>switch(config-router)# graceful-restart</pre>	Enables a graceful restart and the graceful restart helper functionality. This command is enabled by default.  This command triggers an automatic notification and session reset for the BGP neighbor sessions.
<b>Step 5</b>	<b>graceful-restart {restart-time</b> <i>time stalepath-time time</i> <b>Example:</b> <pre>switch(config-router)# graceful-restart restart-time 300</pre>	Configures the graceful restart timers.  The optional parameters are as follows: <ul style="list-style-type: none"> <li>• <b>restart-time</b>—Maximum time for a restart sent to the BGP peer. The range is from 1 to 3600 seconds. The default is 120.</li> </ul> <b>Note</b> Cisco NX-OS Release 10.1(1) supports a higher number of BFD sessions. If BGP sessions are associated with BFD, the BGP <b>restart-time</b> may need to be increased to maintain peer connection during ISSU. <ul style="list-style-type: none"> <li>• <b>stalepath-time</b>—Maximum time that BGP keeps the stale routes from the restarting BGP peer. The range is from 1 to 3600 seconds. The default is 300.</li> </ul> This command triggers an automatic notification and session reset for the BGP neighbor sessions.
<b>Step 6</b>	<b>graceful-restart-helper</b> <b>Example:</b> <pre>switch(config-router)# graceful-restart restart-time 300</pre>	Enables the graceful restart helper functionality. Use this command if you have disabled graceful restart but you still want to enable graceful restart helper functionality. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

	Command or Action	Purpose
<b>Step 7</b>	(Optional) <b>show running-config bgp</b>  <b>Example:</b> switch(config-router)# show running-config bgp	Displays the BGP configuration.
<b>Step 8</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> switch(config-router)# copy running-config startup-config	Saves this configuration change.

### Example

This example shows how to enable a graceful restart:

```
switch# configure terminal
switch(config)# router bgp 65536
switch(config-router)# graceful-restart
switch(config-router)# graceful-restart restart-time 300
switch(config-router)# copy running-config startup-config
```

## Configuring Virtualization

You can configure one BGP process, create multiple VRFs, and use the same BGP process in each VRF.

### Before you begin

You must enable BGP.

### Procedure

	Command or Action	Purpose
<b>Step 1</b>	<b>configure terminal</b>  <b>Example:</b> switch# configure terminal switch(config)#	Enters global configuration mode.
<b>Step 2</b>	<b>vrf context vrf-name</b>  <b>Example:</b> switch(config)# vrf context RemoteOfficeVRF switch(config-vrf)#	Creates a new VRF and enters VRF configuration mode.
<b>Step 3</b>	<b>exit</b>  <b>Example:</b>	Exits VRF configuration mode.



	Command or Action	Purpose
	<code>switch(config-vrf)# exit</code> <code>switch(config)#</code>	
<b>Step 4</b>	<b>router bgp <i>as-number</i></b>  <b>Example:</b> <code>switch(config)# router bgp 65535</code> <code>switch(config-router)#</code>	Creates a new BGP process with the configured autonomous system number.
<b>Step 5</b>	<b>vrf <i>vrf-name</i></b>  <b>Example:</b> <code>switch(config-router)# vrf</code> <code>RemoteOfficeVRF</code> <code>switch(config-router-vrf)#</code>	Enters the router VRF configuration mode and associates this BGP instance with a VRF.
<b>Step 6</b>	<b>neighbor <i>ip-address</i> remote-as <i>as-number</i></b>  <b>Example:</b> <code>switch(config-router-vrf)# neighbor</code> <code>209.165.201.1 remote-as 65535</code> <code>switch(config-router--vrf-neighbor)#</code>	Configures the IP address and AS number for a remote BGP peer.
<b>Step 7</b>	(Optional) <b>copy running-config startup-config</b>  <b>Example:</b> <code>switch(config-router-vrf-neighbor)# copy</code> <code>running-config startup-config</code>	Saves this configuration change.

### Example

This example shows how to create a VRF and configure the router ID in the VRF:

```
switch# configure terminal
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router bgp 65536
switch(config-router)# vrf NewVRF
switch(config-router-vrf)# neighbor 209.165.201.1 remote-as 65536
switch(config-router-vrf-neighbor)# copy running-config startup-config
```

## Verifying the Advanced BGP Configuration

To display the BGP configuration, perform one of the following tasks:

Command	Purpose
<code>show bgp all [summary] [vrf <i>vrf-name</i>]</code>	Displays the BGP information for all address families.
<code>show bgp convergence [vrf <i>vrf-name</i>]</code>	Displays the BGP information for all address families.

Command	Purpose
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>community</b> { <b>regex</b> <i>expression</i>   [ <b>community</b> ] [ <b>no-advertise</b> ] [ <b>no-export</b> ] [ <b>no-export-subconfed</b> ]} [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match a BGP community.
<b>show bgp</b> [ <b>vrf</b> <i>vrf-name</i> ] { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>community-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match a BGP community list.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>extcommunity</b> { <b>regex</b> <i>expression</i>   <b>generic</b> [ <b>non-transitive</b>   <b>transitive</b> ] <i>aa4:nn</i> [ <b>exact-match</b> ]} [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match a BGP extended community.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>extcommunity-list</b> <i>list-name</i> [ <b>exact-match</b> ]} [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match a BGP extended community list.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>extcommunity-list</b> <i>list-name</i> [ <b>exact-match</b> ]} [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information for BGP route dampening. Use the <b>clear bgp dampening</b> command to clear the route flap dampening information.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] { <b>dampening</b> <b>dampened-paths</b> [ <b>regex</b> <i>expression</i> ]} [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP route history paths.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>filter-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information for the BGP filter list.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b>   <b>vpn4</b>   <b>vpn6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>neighbors</b> [ <i>ip-address</i>   <i>ipv6-prefix</i> ] [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information for BGP peers. Use the <b>clear bgp neighbors</b> command to clear these neighbors.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] { <b>nexthop</b>   <b>nexthop-database</b> } [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information for the BGP route next hop.
<b>show bgp paths</b>	Displays the BGP path information.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>policy</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP policy information. Use the <b>clear bgp policy</b> command to clear the policy information.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>prefix-list</b> <i>list-name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match the prefix list.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>received-paths</b> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP paths stored for soft reconfiguration.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>regex</b> <i>expression</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match the AS_path regular expression.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } { <b>unicast</b>   <b>multicast</b> } [ <i>ip-address</i>   <i>ipv6-prefix</i> ] <b>route-map</b> <i>map-name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the BGP routes that match the route map.

Command	Purpose
<b>show bgp peer-policy</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information about BGP peer policies.
<b>show bgp peer-session</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information about BGP peer sessions.
<b>show bgp peer-template</b> <i>name</i> [ <b>vrf</b> <i>vrf-name</i> ]	Displays the information about BGP peer templates. Use the <b>clear bgp peer-template</b> command to clear all neighbors in a peer template.
<b>show bgp process</b>	Displays the BGP process information.
<b>show bgp</b> { <b>ipv4</b>   <b>ipv6</b> } <b>unicast neighbors</b> <i>interface</i>	Displays information about BGP peers for the specified interface.
<b>show ip bgp neighbors</b> <i>interface-name</i>	Displays the interface used as a BGP peer.
<b>show ip route</b> <i>ip-address</i> <b>detail vrf all</b>   <b>i bw</b>	Displays the link bandwidth EXTCOMM fields. <i>bw:xx</i> (such as <i>bw:40</i> ) in the output indicates that BGP peers are sending BGP extended attributes with the bandwidth (for weighted ECMP).
<b>show</b> { <b>ipv4</b>   <b>ipv6</b> } <b>bgp options</b>	Displays the BGP status and configuration information.
<b>show</b> { <b>ipv4</b>   <b>ipv6</b> } <b>mbgp options</b>	Displays the BGP status and configuration information.
<b>show ipv6 routers interface</b> <i>interface</i>	Displays the link-local address of remote IPv6 routers, which is learned through IPv6 ICMP router advertisement.
<b>show running-configuration bgp</b>	Displays the current running BGP configuration.

## Monitoring BGP Statistics

To display BGP statistics, use the following commands:

Command	Purpose

<b>show bgp {ipv4   ipv6} {unicast   multicast} [ip-address   ipv6-prefix] flap-statistics [vrf vrf-name]</b>	Displays the BGP route flap statistics. Use the <b>clear bgp flap-statistics</b> command to clear these statistics.
<b>show bgp {ipv4   ipv6} unicast injected-routes</b>	Displays injected routes in the routing table.
<b>show bgp sessions [vrf vrf-name]</b>	Displays the BGP sessions for all peers. Use the <b>clear bgp sessions</b> command to clear these statistics.
<b>show bgp statistics</b>	Displays the BGP statistics.

## Configuration Examples

This example shows how to enable BFD for individual BGP neighbors:

```
router bgp 400
  router-id 2.2.2.2
  neighbor 172.16.2.3
    bfd
    remote-as 400
    update-source Vlan1002
    address-family ipv4 unicast
```

This example shows how to enable BFD for BGP prefix peers:

```
router bgp 400
  router-id 1.1.1.1
  neighbor 172.16.2.0/24
    bfd
    remote-as 400
    update-source Vlan1002
    address-family ipv4 unicast
```

This example shows how to configure MD5 authentication for prefix-based neighbors:

```
template peer BasePeer-V6
  description BasePeer-V6
  password 3 f4200cfc725bbd28
  transport connection-mode passive
  address-family ipv6 unicast
template peer BasePeer-V4
  bfd
  description BasePeer-V4
  password 3 f4200cfc725bbd28
  address-family ipv4 unicast
--
neighbor fc00::10:3:11:0/127 remote-as 65006
  inherit peer BasePeer-V6
neighbor 10.3.11.0/31 remote-as 65006
  inherit peer BasePeer-V4
```

This example shows how to enable neighbor status change messages globally and suppress them for a specific neighbor:

```
router bgp 65100
  log-neighbor-changes
  neighbor 209.165.201.1 remote-as 65535
    description test
    address-family ipv4 unicast
```

```
soft-reconfiguration inbound
disable log-neighbor-changes
```

## Related Topics

The following topics can give more information on BGP:

- [Configuring Basic BGP](#)
- [Configuring Route Policy Manager](#)

## Additional References

For additional information related to implementing BGP, see the following sections:

### MIBs

MIBs	MIBs Link
MIBs related to BGP	To locate and download supported MIBs, go to the following URL:  <a href="ftp://ftp.cisco.com/pub/mibs/supportlists/nexus9000/Nexus9000MIBSupportList.html">ftp://ftp.cisco.com/pub/mibs/supportlists/nexus9000/Nexus9000MIBSupportList.html</a>