Configuring Advanced BGP

This chapter contains the following sections:

- Advanced BGP, on page 2
- Prerequisites for Advanced BGP, on page 15
- Guidelines and Limitations for Advanced BGP, on page 15
- Default Settings, on page 19
- Configuring Advanced BGP, on page 19
- Configure BGP Additional Paths, on page 37
- Configuring eBGP, on page 41
- Configure AS Confederations, on page 46
- Configure Route Reflector, on page 46
- Configure Next-Hops on Reflected Routes Using an Outbound Route-Map, on page 48
- Configure Route Dampening, on page 51
- Configure Load Sharing and ECMP, on page 51
- Configure Maximum Prefixes, on page 52
- Configure DSCP, on page 52
- Configure Dynamic Capability, on page 53
- Configure Aggregate Addresses, on page 53
- Suppress BGP Routes, on page 54
- Configure BGP Conditional Advertisement, on page 55
- Configure Route Redistribution, on page 57
- Advertise the Default Route, on page 58
- Configure BGP Attribute Filtering and Error Handling, on page 60
- Tune BGP, on page 62
- Configure Policy-Based Administrative Distance, on page 67
- Configure Multiprotocol BGP, on page 68
- Configure BMP, on page 69
- BGP Local Route Leaking, on page 72
- BGP Graceful Shutdown, on page 79
- Configure a Graceful Restart, on page 91
- Configure Virtualization, on page 93
- Verify the Advanced BGP Configuration, on page 95
- Monitor BGP Statistics, on page 97
- Configuration Examples, on page 97

• Related Topics, on page 98
• Additional References, on page 98

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.

Beginning with Cisco NX-OS Release 10.5(1)F, Configuring Basic BGP and Configuring Advanced BGP chapters are merged to create Configuring BGP chapter.

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: This 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: This 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: This 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

A route policy is a configuration mechanism that you associate with a BGP peer to control and modify the routes that BGP recognizes and processes. Route policies use route maps to selectively filter, accept, deny, or modify routes during BGP route updates. You can apply route policies to inbound or outbound route updates, allowing granular control over routing granular control over routing behavior.

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.

Undefined Policies

Currently, in BGP (Border Gateway Protocol) configurations, if a route-map is referenced but has not been defined, the system's default behavior is to deny the route. This can lead to unintended route filtering or route loss, especially during configuration transitions or "Config-Replace" operations where a route-map might be temporarily unavailable while the system processes changes.

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) is a type of Border Gateway Protocol (BGP) that

- is used to exchange routing information between BGP peers within the same autonomous system (AS)
- used within an AS to distribute external routes learned from other ASes (via eBGP) to all routers within that same AS, typically in multihomed BGP networks that have multiple connections to the same external AS
- ensures all routers have a consistent view of external networks, and
- enforces routing policies for traffic leaving the AS.

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

Figure 1: iBGP Network

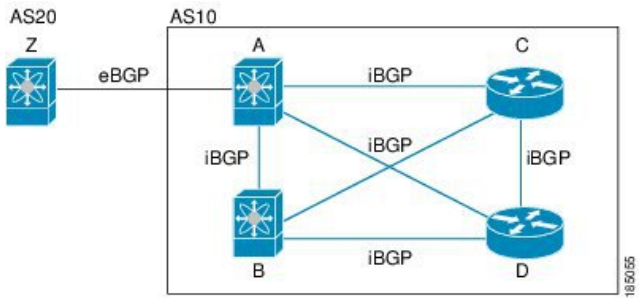

iBGP networks are fully meshed. Each iBGP peer has a direct connection to all other iBGP peers to prevent network loops.

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 41 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

An AS confederation is a BGP architectural technique that divides a large autonomous system (AS) into multiple smaller subautonomous systems, called sub-ASes, which are grouped together to appear as a single AS to external networks. This approach reduces the complexity of a fully meshed iBGP network by limiting the number of iBGP peer connections required within the overall AS.

You can reduce the iBGP mesh by dividing the autonomous system into multiple subautonomous systems and grouping them into a single confederation. A confederation 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

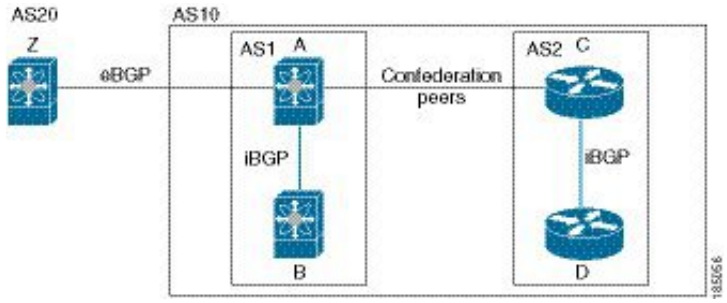

In this example, AS10 is split into two subautonomous systems, AS1 and AS2. Each subautonomous system is fully meshed, but there is only one link between the subautonomous systems. By using AS confederations, you can reduce the number of links compared to the fully meshed autonomous system.

Route reflectors

Route reflectors are a BGP architectural technique that reduces the complexity of a fully meshed iBGP network by allowing certain routers, called route reflectors, to redistribute routes learned from one iBGP peer to other iBGP peers. This eliminates the need for every iBGP peer to be directly connected to every other iBGP peer.

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 sp iBGP peers do not need to be fully meshed.

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

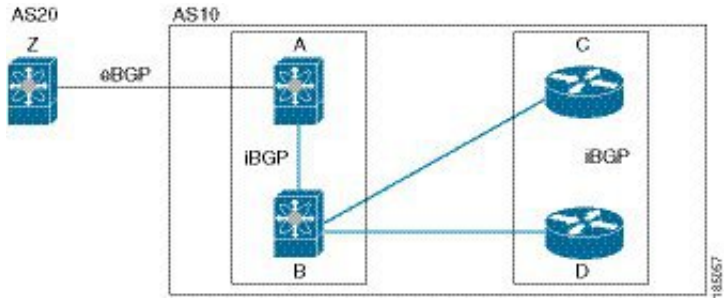

The route reflector and its client peers form a cluster. You do not have to configure all iBGP peers to act as client peers of the route reflector. You must configure any nonclient peer as fully meshed to guarantee that complete BGP updates reach all peers.

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. It works by assigning a penalty to routes that frequently alternate between available and unavailable states (flapping). When the penalty exceeds a configured suppression limit, the route is suppressed (dampened) and not advertised to neighbors until the penalty decays below a reuse threshold.

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.

Note When leaking routes between VRFs with BGP, you may experience that route changes are not reflected in the destination VRF until several minutes later. This can happen due to BGP dampening mechanisms which ignore route changes based on metric changes to avoid cases in which an IGP is flapping (link issues, for example) and you do not want these flaps to be propagated beyond. If you suffer from this condition and you want the routes to be immediately updated on the destination VRF as well, see Configure dampen-igp-metric in BGP on Cisco Nexus Switches.

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.

To configure `as-path multipath-relax` command effectively, configure the command per VRF under BGP. Also, configure as-path multipath-relax command under the custom VRF so that multiple routers get installed in the custom VRF Route-Target (RT).

The BGP best-path algorithm considers the paths as equal-cost paths if these 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

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. Figure Figure 4: BGP Route Advertisement with the Additional Paths Capability, on page 8 illustrates the BGP additional paths capability.

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.

Figure 4: BGP Route Advertisement with the Additional Paths Capability

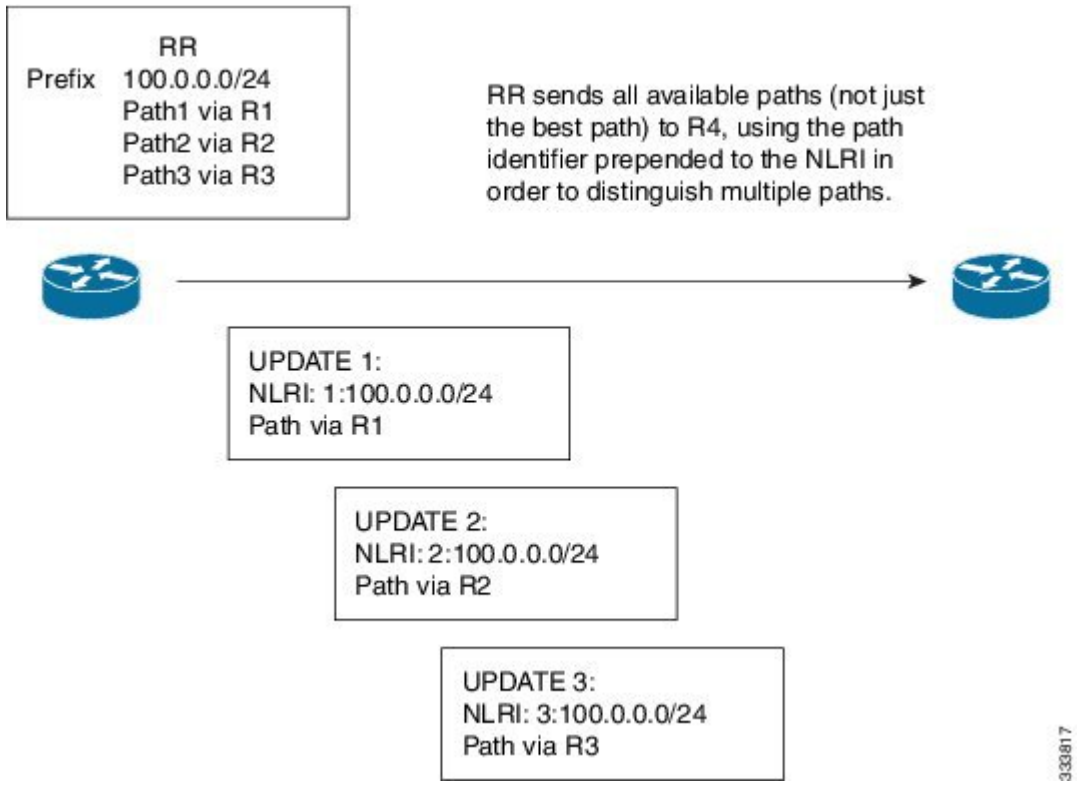

For information on configuring BGP additional paths, see the Configure BGP Additional Paths, on page 37section.

Route aggregation

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 these 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 Configuring BGP Next-Hop Address Tracking for more information.

Route redistribution

Route redistributions are configurations in BGP that allow you to import routes from static routes or other routing protocols into BGP. To control which routes are redistributed, you must use a route map. A route map acts as a filter and can match routes based on attributes such as destination prefix, originating protocol, route type, route tag, and more. This filtering helps ensure that only desired routes are passed into BGP. See Configuring Route Policy Manager, for more information.

You can also use route maps to override the default behavior of route redistribution. However, caution is necessary because incorrect use of route maps can cause network loops, which can destabilize the network.

Theses 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

Bidirectional forwarding detection (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. Advanced BGP supports BFD for IPv4 and IPv6.

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.

Tune BGP

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

BGP timers

BGP timers are protocol timers used in the Border Gateway Protocol (BGP) that

- control the timing of periodic keepalive messages sent between BGP peers,
- determine the timeout period for detecting lost neighbor sessions when keepalives are not received, and
- manage other specific protocol events and features.

These timers are typically configured in seconds and include a random adjustment to prevent synchronization of timer events across different BGP peers.

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. The timers include a random adjustment so that the same timers on different BGP peers trigger at different times.

Tune 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.

Multicast BGP does not propagate multicast state information, so 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

Cisco NX-OS RFC 6368 enables Internal Border Gateway Protocol (iBGP) support between Provider Edge (PE) and Customer Edge (CE) feature is implemented in Cisco NX-OS.

In current deployments, when BGP is used as the PE or CE 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.

Framework

Beginning with Cisco NX-OS Release 10.1(2), you can use iBGP PE-CE to do these:

- 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

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.

BGP needs to converge before graceful-restart timer expires. BGP graceful-restart timer needs to be increased in high route scale network accordingly in order to avoid temporary traffic loss. If BGP itself provides the reachability to open other BGP sessions, then stalepath-time should also be increased to accommodate for the extra time needed to converge the overlay session after the initial underlay session has already converged.

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 low memory handling is a mechanism that manages BGP peer sessions and memory usage when the router experiences low memory conditions. It helps maintain stability and performance by controlling peer establishment and shutdown based on the severity of the memory alert.

BGP reacts to low memory for theses 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.

- If you are writing JSON payload, use the standard JSON syntax defined on RFC 8259: The JavaScript Object Notation (JSON) Data Interchange Format.
- 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.
- Edge Services Gateway (ESG) route-inject into BGP is assigned weight 0.
- 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 applies only to routers in a public AS only. The workaround to this restriction would be to apply the **neighbor local-as** command on a per-neighbor basis, with the local AS number being a public AS number.
  - 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.
- In case the AS-Path is missing from the matching route for the next-hop in the BGP table or there is no matching route for the nexthop in the BGP table then the AS number is not printed in the traceroute output along with the next hop.
- 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.
  - While configuring BGP interface peering using IP6 Link-Local static IPv6 address, ensure the subnet is matching with the peer for BGP to work.

- 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 "ing-sup" from 512 to 768 except for Cisco Nexus 34XX-S platform, where default carving is sufficient.
- The command **[maximum-paths eibgp]** is supported only in MPLS environments.
- Route-map deletion feature adds a mechanism to block the deletion of entire route-map that is associated with the BGP. With the route-map deletion blocked, the modifications to the route-map statement are still allowed.
- If there are more than one sequence in the route-map, user can still delete any route map sequence until there is at least one sequence available.
- Users can have the forward reference case for route-map from client. However, once route-map is created and associated, the deletion of route-map is blocked.
- Blocking deletion functionality is configurable dynamically using the knob.
- It is allowed to delete the BGP association to the route-map and deletion of route-map itself in a single transaction payload.

- It is allowed to add the BGP association to the route-map and an error must be thrown for deletion of route-map.
- The following is the list of the dual stage related behaviors:
  - If knob and deletion occur together, dual stage has to verify and throw an error without commit.
  - If knob already exists and route-map deletion occurs in dual stage, it must throw an error.
  - If route-map and CLI knob is single commit with different order, it must throw an error.
  - If knob is not enabled and route-map deletion occurs in dual stage, it has to execute successfully.
  - In a single verify, if "cli knob is disabled AND route-map deletion" is executed, the route-map deletion is allowed.
- If the route-map used by BGP template is not inherited by any of the BGP neighbors, the entire route-map deletion will still be blocked.
- There are few commands under vrf context that are owned by BGP, but are not part of bgpInst.
- Cloudscale IPv6 link-local BGP support requires carving > 512 ing-sup TCAM region (this requires a reload to take effect).
- As the VPN address family (L3VPN and EVPN) is not supported, the routes received from confederate peers are not advertised in the VPN address family.
- Encryption decrypt type6 is not supported for BGP passwords and keychain.

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

Enable 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
                        switch(config)#
```

Enters global configuration mode.

Step 2 **interface** *type slot/port*

Example:

```
switch(config)# interface ethernet 1/2
                        switch(config-if)#
```

Enters interface configuration mode.

Step 3 **ip forward**

Example:

```
switch(config-if)# ip forward
```

Allows IPv4 traffic on the interface even when there is no IP address configuration on that interface.

Step 4 (Optional) **copy running-config startup-config**

Example:

```
switch(config-if)# copy running-config startup-config
```

Saves this configuration change.

Configure 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).

- 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.
- When using BGP Peer Template, there is no check for the commands used inside template to verify if that command applies to iBGP/eBGP peer or not. For example if you create a template and add a command "Remove-private-as" inside a template and then assign this tempate to iBGP peer, then no error will be printed saying this command "Remove-private-as" does not apply to iBGP peer.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 **router bgp***autonomous-system-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enables BGP and assigns the autonomous system number to the local BGP speaker.

Step 3 **template peer-session***template-name*

Example:

```
switch(config-router)# template
peer-session BaseSession
switch(config-router-stmp)#
```

Enters peer-session template configuration mode.

Step 4 (Optional) **password***number password*

Example:

```
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) **timers***keepalive hold*

Example:

```
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 **exit**

Example:

```
switch(config-router-stmp)# exit
switch(config-router)#
```

Exits peer-session template configuration mode.

Step 7 **neighbor***ip-address***remote-as***as-number*

Example:

```
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 **inherit peer-session***template-name*

Example:

```
switch(config-router-neighbor)# inherit peer-session
BaseSession
switch(config-router-neighbor)#
```

Applies a peer-session template to the peer.

Step 9 (Optional) **description***text*

Example:

```
switch(config-router-neighbor)#
description Peer Router A
switch(config-router-neighbor)#
```

Adds a description for the neighbor.

Step 10 (Optional) **show bgp peer-session***template-name*

Example:

```
switch(config-router-neighbor)# show bgp
peer-session BaseSession
```

Displays the peer-policy template.

Step 11 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor)# copy
running-config startup-config
```

Saves this configuration change.

Use the **show bgp neighbor** command to see the template applied.

Example

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

```
switch# <userinput>configure terminal</userinput>
switch(config)# <userinput>router bgp 65536</userinput>
switch(config-router)# <userinput>template peer-session BaseSession</userinput>
switch(config-router-stmp)# <userinput>timers 30 90</userinput>
switch(config-router-stmp)# <userinput>exit</userinput>
```

```
switch(config-router)# neighbor <userinput>192.168.1.2 remote-as 65536</userinput>
switch(config-router-neighbor)# <userinput>inherit peer-session BaseSession</userinput>
switch(config-router-neighbor)# <userinput>description Peer Router A</userinput>
switch(config-router-neighbor)# <userinput>address-family ipv4 unicast</userinput>
switch(config-router-neighbor-af)# <userinput>copy running-config startup-config</userinput>
```

Configure 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.

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).

- 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
```

Enters configuration mode.

Step 2 **router bgp***autonomous-system-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enables BGP and assigns the autonomous system number to the local BGP speaker.

Step 3 **template peer-session***template-name*

Example:

```
switch(config-router)# template
peer-policy BasePolicy
switch(config-router-ptmp)#
```

Creates a peer-policy template.

Step 4 (Optional) **advertise-active-only**

Example:

```
switch(config-router-ptmp)#
advertise-active-only
```

Advertises only active routes to the peer.

Step 5 (Optional) **maximum-prefix***number*

Example:

```
switch(config-router-ptmp)#
maximum-prefix 20
```

Sets the maximum number of prefixes allowed from this peer.

Step 6 **exit**

Example:

```
switch(config-router-ptmp)# exit
switch(config-router)#
```

Exits peer-policy template configuration mode.

Step 7 **neighbor***ip-address***remote-as***as-number*

Example:

```
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 **address-family** {**ipv4** | **ipv6**} {**multicast** | **unicast**}

Example:

```
switch(config-router-neighbor)#
address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

Enters global address family configuration mode for the address family specified.

Step 9 **inherit peer-policy***template-name preference*

Example:

```
switch(config-router-neighbor-af)#
inherit peer-policy BasePolicy 1
```

Applies a peer-policy template to the peer address family configuration and assigns the preference value for this peer policy.

Step 10 (Optional) **show bgp peer-policy***template-name*

Example:

```
switch(config-router-neighbor-af)# show
bgp peer-policy BasePolicy
```

Displays the peer-policy template.

Step 11 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor-af)# copy
running-config startup-config
```

Saves this configuration change.

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

Configure 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).

- 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
```

Enters global configuration mode.

Step 2 **router bgp***autonomous-system-number*

Example:

```
switch(config)# router bgp 65535
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **template peer***template-name*

Example:

```
switch(config-router)# template peer
BasePeer
```

Enters peer template configuration mode.

Step 4 (Optional) **inherit peer-session***template-name*

Example:

```
switch(config-router-neighbor)# inherit
peer-session BaseSession
```

Adds a peer-session template to the peer template.

Step 5 (Optional) **address-family** {**ipv4**|**ipv6**} {**multicast**|**unicast**}

Example:

```
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.

Step 6 (Optional) **inherit peer-policy***template-name*

Example:

```
switch(config-router-neighbor-af)#
inherit peer-policy BasePolicy 1
```

Applies a peer-policy template to the neighbor address family configuration.

Step 7 **exit**

Example:

```
switch(config-router-neighbor-af)# exit
```

Exits BGP neighbor address family configuration mode.

Step 8 (Optional) **timers** *keepalive hold*

Example:

```
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.

Step 9 **exit**

Example:

```
switch(config-router-neighbor)# exit
```

Exits BGP neighbor configuration mode.

Step 10 **neighbor** *ip-address* **remote-as** *as-number*

Example:

```
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.

Step 11 **inherit peer** *template-name*

Example:

```
switch(config-router-neighbor)# inherit
peer BasePeer
```

Inherits the peer template.

Step 12 (Optional) **timers** *keepalive hold*

Example:

```
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.

Step 13 (Optional) **show bgp peer-template** *template-name*

Example:

```
switch(config-router-neighbor)# show
bgp peer-template BasePeer
```

Displays the peer template.

Step 14 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor)# copy
running-config startup-config
```

Saves this configuration change.

Use the **show bgp neighbor** 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# <userinput>configure terminal</userinput>
switch(config)# <userinput>router bgp 65536</userinput>
switch(config-router)# <userinput>template peer BasePeer</userinput>
switch(config-router-neighbor)# <userinput>inherit peer-session BaseSession</userinput>
switch(config-router-neighbor)# <userinput>address-family ipv4 unicast</userinput>
switch(config-router-neighbor-af)# <userinput>inherit peer-policy BasePolicy 1</userinput>
switch(config-router-neighbor-af)# <userinput>exit</userinput>
switch(config-router-neighbor)# <userinput>exit</userinput>
switch(config-router)# <userinput>neighbor 192.168.1.2 remote-as 65536</userinput>
```

```
switch(config-router-neighbor)# <userinput>inherit peer BasePeer</userinput>
switch(config-router-neighbor)# <userinput>copy running-config startup-config</userinput>
```

Configure 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

Step 1 **timers prefix-peer-timeout** *value*

Example:

```
switch(config-router-neighbor)# timers
prefix-peer-timeout 120
```

Configures the BGP prefix peering timeout value in router configuration mode. The range is from 0 to 1200 seconds. The default value is 30.

Note
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.

Step 2 **maximum-peers** *value*

Example:

```
switch(config-router-neighbor)#
maximum-peers 120
```

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

Configure 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 (see the Enabling BGP section).

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
```

Enters configuration mode.

Step 2 **router bgp** *autonomous-system-number*

Example:

```
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.

Step 3 **neighbor** *interface-name* **remote-as** {*as-number* | **route-map** *map-name*}

Example:

```
switch(config-router)# neighbor
                        Ethernet1/1 remote-as 65535
                        switch(config-router-neighbor)#
```

Places the router in the neighbor configuration mode for BGP routing and configures the interface for BGP peering.

Note
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 *Dynamic AS Numbers for Prefix Peers and Interface Peers* for more information about using dynamic AS numbers.

Step 4 **inherit peer** *template-name*

Example:

```
switch(config-router-neighbor)# inherit peer PEER
```

Inherits the peer template.

Step 5 **address-family** {**ipv4** | **ipv6**} **unicast**

Example:

```
switch(config-router-neighbor)#
                        address-family ipv4 unicast
                        switch(config-router-neighbor-af)#
```

Enters global address family configuration mode for the address family specified.

Step 6 (Optional) **show bgp** {**ipv4** | **ipv6**} **unicast neighbors** *interface*

Example:

```
switch(config-router-neighbor-af)# show
                        bgp ipv4 unicast neighbors e1/25
```

Example:

```
switch(config-router-neighbor-af)# show
                        bgp ipv6 unicast neighbors 3FFE:700:20:1::11
```

Displays information about BGP peers.

Step 7 (Optional) **show ip bgp neighbors** *interface-name*

Example:

```
switch(config-router-neighbor-af)# show ip bgp neighbors Ethernet1/1
```

Displays the interface used as a BGP peer.

Step 8 (Optional) **show ipv6 routers** [**interface** *interface*]

Example:

```
switch(config-router-neighbor-af)# show ipv6 routers interface Ethernet1/1
```

Displays the link-local address of remote IPv6 routers, which is learned through IPv6 ICMP router advertisement.

Step 9 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor-af)# copy
                        running-config startup-config
```

Saves this configuration change.

Example

This example shows how to configure BGP Interface Peering via IPv6 Link-Local for IPv4 and IPv6 Address Families.

iBGP Interface Peering Configuration for Leaf 1:

```
switch# configure terminal
                switch(config)# router bgp 65000
                switch(config-router)# neighbor Ethernet1/1 remote-as 65000
```

```
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

BGP version 4, remote router ID 5.5.5.5
Neighbor previous state = OpenConfirm
BGP state = Established, up for 2d16h
Neighbor vrf: default

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.

Configure 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:

Before you begin

- Ensure the primary-key is configured using the **key config-key ascii** <*primary_key*> command on Cisco NX-OS switches.
- For Type-6 encryption to function properly, ensure **feature password encryption aes** is enabled on Cisco NX-OS switches.

Procedure

password {**0** | **3** | **7**} *string*

Configures a Type-6 encryption 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:

SUMMARY STEPS

1. **soft-reconfiguration inbound**
2. (Optional) **clear bgp** {**ipv4** | **ipv6** } {**unicast** | **multicast** *ip-address* **soft** {**in** | **out**}
3. **clear bgp** {**ipv4** | **ipv6**} {**unicast** | **multicast**} *ip-address* **soft** (**in** | **out**)

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	**soft-reconfiguration inbound** **Example:** `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.
Step 2	(Optional) **clear bgp** {**ipv4**	**ipv6** } {**unicast**	**multicast** *ip-address* **soft** {**in**	**out**} **Example:** `switch# clear bgp ip unicast 192.0.2.1 soft in`	Resets the BGP session without tearing down the TCP session.
Step 3	**clear bgp** {**ipv4**	**ipv6**} {**unicast**	**multicast**} *ip-address* **soft** (**in**	**out**) **Example:** `switch# clear bgp ip unicast 192.0.2.1 soft in`	Resets the BGP session without tearing down the TCP session.

Modify 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

Step 1 **next-hop-self**

Example:

```
switch(config-router-neighbor-af)# next-hop-self
```

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.

Step 2 **next-hop-third-party**

Example:

```
switch(config-router-neighbor-af)#
next-hop-third-party
```

Sets the next-hop address as a third-party address. Use this command for single-hop eBGP peers that do not have **next-hop-self** configured.

Configure 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

nexthop trigger-delay {**critical** | **non-critical**} *milliseconds*

Example:

```
switch(config-router-af)# nexthop
trigger-delay critical 5000
```

Specifies the next-hop address tracking delay timer for critical next-hop reachability routes and for noncritical routes. The range is from 1 to 4294967295 milliseconds. The critical timer default is 3000. The noncritical timer default is 10000.

Configure 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

nexthop trigger-delay {**critical** | **non-critical**} *milliseconds*

Example:

```
switch(config-router-af)# nexthop
trigger-delay critical 5000
```

Specifies the next-hop address tracking delay timer for critical next-hop reachability routes and for noncritical routes. The range is from 1 to 4294967295 milliseconds. The critical timer default is 3000. The noncritical timer default is 10000.

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:

SUMMARY STEPS

1. [**no**] **nexthop suppress-default-resolution**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	[**no**] **nexthop suppress-default-resolution** **Example:** `switch(config-router)# nexthop suppress-default-resolution`	Prevents resolution of BGP next hop through the IP default route. When this command is enabled: • The output of the **show bgp process detail** command includes the following line: Use default route for nexthop resolution: No • The output of the **show routing clients bgp** command includes the following line: Owned rnh will never resolve to 0.0.0.0/0

Control 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
next-hop-self [all] **Example**: `switch(config-router-af)# next-hop-self all`	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.

Shrink 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

neighbor-down fib-accelerate

Example:

```
switch(config-router)# neighbor-down
fib-accelerate
```

Withdraws the corresponding next hop from all next-hop groups (ECMP groups and single next-hop routes) whenever a BGP session goes down.

Note
This command applies to both IPv4 and IPv6 routes.

Disable 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

dont-capability-negotiate

Example:

```
switch(config-router-neighbor)#
dont-capability-negotiate
```

Disables capabilities negotiation. You must manually reset the BGP sessions after configuring this command.

Disable 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

disable-policy-batching

Example:

```
switch(config-router)#
disable-policy-batching
```

Disables the batching evaluation of prefix advertisements to all peers.

Configure BGP Additional Paths

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

Advertise 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

Step 1 [**no**] **capability additional-paths send** [**disable**]

Example:

```
switch(config-router-neighbor-af)#
capability addtional-paths send
```

Advertises the capability to send additional paths to the BGP peer. The **disable** option disables the advertising capability of sending additional paths.

The **no** form of this command disables the capability of sending additional paths.

Step 2 [**no**] **capability additional-paths receive** [**disable**]

Example:

```
switch(config-router-neighbor-af)#
capability addtional-paths receive
```

Advertises the capability to receive additional paths from the BGP peer. The **disable** option disables the advertising capability of receiving additional paths.

The **no** form of this command disables the capability of receiving additional paths.

Step 3 **show bgp neighbor**

Example:

```
switch(config-router-neighbor-af)# show
bgp neighbor
```

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

Configure 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

Step 1 [**no**] **additional-paths send**

Example:

```
switch(config-router-af)# additional-paths
send
```

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 **no** form of this command disables the send capability.

Step 2 [**no**] **additional-paths receive**

Example:

```
switch(config-router-af)# additional-paths
receive
```

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 **no** form of this command disables the receive capability.

Step 3 **show bgp neighbor**

Example:

```
switch(config-router-af)# show bgp
neighbor
```

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

Configure 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

Step 1 [**no**] **set ip next-hop unchanged**

Example:

```
switch(config-route-map)# set ip next-hop
unchanged
```

Specifies and unchanged next-hop IP address.

Step 2 [**no**] **set path-selection { all | backup | best2 | multipaths}** | advertise

Example:

```
switch(config-route-map)# set
path-selection all advertise
```

Specifies that all paths be advertised for a given prefix. You can use one of the following options:

- all—Advertises all available valid paths.

- backup—Advertises paths marked as backup paths. This option requires that backup paths be enabled using the additional-path install backup command.
- best2—Advertises the second best path, which is the best path of the remaining available paths, except the already calculated best path.
- multipaths—Advertises all multipaths. This option requires that multipaths be enabled using the maximum-paths command.

Note
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.

The **no** form of this command specifies that only the best path be advertised.

Step 3 **show bgp** {**ipv4** | **ipv6**} **unicast** [*ip-address* | *ipv6-prefix*] [**vrf** *vrf-name*]

Example:

```
switch(config-route-map)# show bgp ipv4
unicast
```

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
```

Configure 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

Step 1 [**no**] **additional-paths selection route-map** *map-name*

Example:

```
switch(config-router-af)# additional paths
selection route-map map1
```

Configures the capability of selecting additional paths for a prefix.

The **no** form of this command disables the additional paths selection capability.

Step 2 **show bgp** {**ipv4** | **ipv6**} **unicast** [*ip-address* | *ipv6-prefix*] [**vrf** *vrf-name*]

Example:

```
switch(config-route-af)# show bgp ipv4
unicast
```

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

Configuring eBGP

Disable 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

disable-connected-check

Example:

```
switch(config-router-neighbor)#
disable-connected-check
```

Disables checking whether or not a single-hop eBGP peer is directly connected. You must manually reset the BGP sessions after using this command.

Configure TTL Security Hops

Perform this task to allow BGP to establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the BGP neighbor session.

Before you begin

To maximize the effectiveness of the BGP Support for TTL Security Check feature, we recommend that you configure it on each participating router. Enabling this feature secures the eBGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router.

- The **neighbor ebgp-multihop** command is not needed when the BGP Support for TTL Security Check feature is configured for a multihop neighbor session and should be disabled before configuring this feature.
- The effectiveness of the BGP Support for TTL Security Check feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected neighbor sessions to handle the attack.
- This feature is not effective against attacks from a peer that has been compromised inside of the local and remote network. This restriction also includes peers that are on the network segment between the local and remote network.

Procedure

Step 1 **enable**

Example:

```
switch(config)# enable
```

Enables privileged EXEC mode.

Enter your password if prompted.

Step 2 **trace** *[protocol] destination*

Example:

```
switch(config)# trace ip 10.1.1.1
```

Discovers the routes of the specified protocol that packets will actually take when traveling to their destination.

Enter the trace command to determine the number of hops to the specified peer.

Step 3 **configure terminal**

Example:

```
switch(config)# configure terminal
```

Enters global configuration mode.

Step 4 **router bgp** *autonomous-system-number*

Example:

```
switch(config)# router bgp 65000
```

Enters router configuration mode, and creates a BGP routing process.

Step 5 **neighbor** *ip-address*

Example:

```
switch(config)# neighbor 10.1.1.1
```

Configures the neighbor IP address.

Step 6 **ttl-security hops** *hop-count*

Example:

```
switch(config)# ttl-security hops 2
```

Configures the maximum number of hops that separate two peers.

The hop-count argument is set to the number of hops that separate the local and remote peer. If the expected TTL value in the IP packet header is 254, then the number 1 should be configured for the hop-count argument. The range of values is a number from 1 to 254.

When the BGP Support for TTL Security Check feature is enabled, BGP will accept incoming IP packets with a TTL value that is equal to or greater than the expected TTL value. Packets that are not accepted are discarded.

The example configuration sets the expected incoming TTL value to at least 253, which is 255 minus the TTL value of 2, and this is the minimum TTL value expected from the BGP peer. The local router will accept the peering session from the 10.1.1.1 neighbor only if it is one or two hops away.

Step 7 **end**

Example:

```
switch(config)# end
```

Exits router configuration mode and enters privileged EXEC mode.

Step 8 **show running-config**

Example:

```
switch(config)# show running-config | begin bgp
```

(Optional) Displays the contents of the currently running configuration file.

The output of this command displays the configuration of the neighbor ttl-security command for each peer under the BGP configuration section of output. That section includes the neighbor address and the configured hop count.

Note
Only the syntax applicable to this task is used in this example. For more details, see the Cisco IOS IP Routing: BGP Command Reference.

Step 9 **show ip bgp neighbors** *[ip-address]*

Example:

```
switch(config)# show ip bgp neighbors 10.4.9.5
```

(Optional) Displays information about the TCP and BGP connections to neighbors.

This command displays "External BGP neighbor may be up to number hops away" when the BGP Support for TTL Security Check feature is enabled. The number value represents the hop count. It is a number from 1 to 254.

Note
Only the syntax applicable to this task is used in this example. For more details, see the Cisco IOS IP Routing: BGP Command Reference.

Configure 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

ebgp-multihop *ttl-value*

Example:

```
switch(config-router-neighbor)#
ebgp-multihop 5
```

Configures the eBGP TTL value for eBGP multihop. The range is from 2 to 255. You must manually reset the BGP sessions after using this command.

Disable a Fast External Fallover

Be 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

no fast-external-fallover

Example:

```
switch(config-router)# no
fast-external-fallover
```

Disables a fast external fallover for eBGP peers. This command is enabled by default.

Limit 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

maxas-limit *number*

Example:

```
switch(config-router)# maxas-limit 50
```

Discards eBGP routes that have a number of AS-path segments that exceed the specified limit. The range is from 1 to 512.

Configure 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

local-as *number* [**no-prepend** [**replace-as** [**dual-as**]]]

Example:

```
switch(config-router-neighbor)# local-as
1.1
```

Configures eBGP to prepend the local AS *number* to the AS_PATH attribute. 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.

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
```

Configure 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

Step 1 **confederation identifier** *as-number*

Example:

```
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.

Step 2 **bgp confederation peers** *as-number* [*as-number2...*]

Example:

```
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.

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
```

Enters global configuration mode.

Step 2 **router bgp***as-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **cluster-id***cluster-id*

Example:

```
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.

Step 4 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
switch(config-router)# address-family
ipv4 unicast
switch(config-router-af)#
```

Enters router address family configuration mode for the specified address family.

Step 5 (Optional) **client-to-client reflection**

Example:

```
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.

Step 6 **exit**

Example:

```
switch(config-router-af)# exit
switch(config-router)#
```

Exits router address configuration mode.

Step 7 **neighbor***ip-address***remote-as***as-number*

Example:

```
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.

Step 8 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
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.

Step 9 **route-reflector-client**

Example:

```
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.

Step 10 (Optional) **show bgp** {**ipv4** | **ipv6**} {**unicast** | **multicast**} **neighbors**

Example:

```
switch(config-router-neighbor-af)#
show bgp ipv4 unicast neighbors
```

Displays the BGP peers.

Step 11 (Optional) **copy running-config startup-config**

Example:

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

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 200
switch(config-router)#
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **neighbor** *ip-address* **remote-as** *as-number*

Example:

```
switch(config-router)# neighbor
192.0.2.12 remote-as 200
switch(config-router-neighbor)#
```

Configures the IP address and AS number for a remote BGP peer.

Step 4 (Optional) **update-source** *interface number*

Example:

```
switch(config-router-neighbor)#
update-source loopback 300
```

Specifies and updates the source of the BGP session.

Step 5 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
switch(config-router-neighbor)#
address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

Enters router address family configuration mode for the specified address family.

Step 6 **route-reflector-client**

Example:

```
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.

Step 7 **route-map** *map-name* **out**

Example:

```
switch(config-router-neighbor-af)#
route-map setrrnh out
```

Applies the configured BGP policy to outgoing routes.

Step 8 (Optional) **show bgp** {**ipv4** | **ipv6**} {**unicast** | **multicast**} [**ip-address** | **ipv6-prefix**] **route-map** *map-name* [**vrf** *vrf-name*]

Example:

```
switch(config-router-neighbor-af)#
show bgp ipv4 unicast route-map
setrrnh
```

Displays the BGP routes that match the route map.

Step 9 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor-af)#
copy running-config startup-config
```

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)# <userinput>interface loopback 300</userinput>
 switch(config-if)# <userinput>ip address 192.0.2.11/32</userinput>
```

```
switch(config-if)# <userinput>ipv6 address 2001::a0c:1a65/64</userinput>
switch(config-if)# <userinput>ip router ospf 1 area 0.0.0.0</userinput>
switch(config-if)# <userinput>exit</userinput>
switch(config)# <userinput>route-map setrrnh permit 10</userinput>
switch(config-route-map)# <userinput>set ip next-hop peer-address</userinput>
switch(config-route-map)# <userinput>exit</userinput>
switch(config)# <userinput>route-map setrrnhv6 permit 10</userinput>
switch(config-route-map)# <userinput>set ipv6 next-hop peer-address</userinput>
switch(config-route-map)# <userinput>exit</userinput>
switch(config)# <userinput>router bgp 200</userinput>
switch(config-router)# <userinput>neighbor 192.0.2.12 remote-as 200</userinput>
switch(config-router-neighbor)# <userinput>update-source loopback 300</userinput>
switch(config-router-neighbor)# <userinput>address-family ipv4 unicast</userinput>
switch(config-router-neighbor-af)# <userinput>route-reflector-client</userinput>
switch(config-router-neighbor-af)# <userinput>route-map setrrnh out</userinput>
switch(config-router-neighbor-af)# <userinput>exit</userinput>
switch(config-router-neighbor)# <userinput>address-family ipv6 unicast</userinput>
switch(config-router-neighbor-af)# <userinput>route-reflector-client</userinput>
switch(config-router-neighbor-af)# <userinput>route-map setrrnhv6 out</userinput>
```

Configure 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

dampening [{*half-life reuse-limit suppress-limit max-suppress-time* | **route-map** *map-name*}]

Example:

```
switch(config-router-af)# dampening
route-map bgpDamp
```

Disables capabilities negotiation. The parameter values are as follows:

- *half-life*—The range is from 1 to 45.
- *resuse-limit*—The range is from 1 to 20000.
- *suppress-limit*—The range is from 1 to 20000.
- *max-suppress-time*—The range is from 1 to 255.

Configure 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

maximum-paths [**ibgp**] *maxpaths*

Example:

```
switch(config-router-af)# maximum-paths 8
```

Configures the maximum number of equal-cost paths for load sharing. The default is 1.

Configure 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

maximum-prefix *maximum* [*threshold*] [**restart** *time* | **warning-only**]

Example:

```
switch(config-router-neighbor-af)#
maximum-prefix 12
```

Configures the maximum number of prefixes from a peer. The parameter ranges are as follows:

- *maximum*—The range is from 1 to 300000.
- *threshold*—The range is from 1 to 100 percent. The default is 75 percent.
- *time*—The range is from 1 to 65535 minutes.

This command triggers an automatic notification and session reset for the BGP neighbor sessions if the prefix is exceeded.

Configure 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

dscp *dscp_value*

Example:

```
switch(config-router-neighbor)# dscp
                        63
```

Below is an example of the corresponding **show** command:

```
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
                        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
```

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: **ef**, **af11**, **af12**, **af13**, **af21**, **af22**, **af23**, **af31**, **af32**, **af33**, **af41**, **af42**, **af43**, **cs1**, **cs2**, **cs3**, **cs4**, **cs5**, **cs6**, or **cs7**.

The default value is cs6.

Configure Dynamic Capability

You can configure dynamic capability for a BGP peer.

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

Procedure

dynamic-capability

Example:

```
switch(config-router-neighbor)#
dynamic-capability
```

Enables dynamic capability. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

Configure 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

aggregate-address *ip-prefix/length* [**as-set**] [**summary-only**] [**advertise-map** *map-name*] [**attribute-map** *map-name*] [**suppress-map** *map-name*]

Example:

```
switch(config-router-af)#
aggregate-address 192.0.2.0/8 as-set
```

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:

- The **as-set** keyword generates autonomous system set path information and community information from contributing paths.
- The **summary-only** keyword filters all more specific routes from updates.
- The **advertise-map** keyword and argument specify the route map used to select attribute information from selected routes.
- The **attribute-map** keyword and argument specify the route map used to select attribute information from the aggregate.
- The **suppress-map** keyword and argument conditionally filter more specific routes. If you specify the **suppress-map** 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.
- The **suppress-map** keyword and argument conditionally filter more specific routes. If you specify the **suppress-map** 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 suppress-map 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.

Suppress 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

suppress-fib-pending

Example:

```
switch(config-router)#
suppress-fib-pending
```

Suppresses newly learned BGP routes (IPv4 or IPv6) from being advertised to downstream BGP neighbors until the routes have been programmed in the hardware.

Configure 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.
- Nexus does not support any other BGP Attribute change operation (example prepend AS Path) with Conditional Route Advertisements. It is used to control which routes are advertised based on exist/non-exist map configuration.

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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **neighbor** *ip-address* **remote-as** *as-number*

Example:

```
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.

Step 4 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
switch(config-router-neighbor)#
address-family ipv4 multicast
switch(config-router-neighbor-af)#
```

Enters address family configuration mode.

Step 5 **advertise-map** *adv-map* {**exist-map** *exist-rmap*|**non-exist-map** *nonexist-rmap*}

Example:

```
switch(config-router-neighbor-af)#
advertise-map advertise exist-map exist
```

Configures BGP to conditionally advertise routes based on the two configured route maps:

- *adv-map*—Specifies a route map with **match** statements that the route must pass before BGP passes the route to the next route map. The *adv-map* is a case-sensitive, alphanumeric string up to 63 characters.
- *exist-rmap*—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 *exist-rmap* is a case-sensitive, alphanumeric string up to 63 characters.
- *nonexist-rmap*—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 *nonexist-rmap* is a case-sensitive, alphanumeric string up to 63 characters.

Note
For BGP conditional advertisement feature, ensure that the "le" or "ge" statements are not used on prefix-list when associated to exist or nonexist map.

Step 6 (Optional) **show bgp** {**ipv4** | **ipv6**} {**unicast** | **multicast**} **neighbors**

Example:

```
switch(config-router-neighbor-af)# show
ip bgp neighbor
```

Displays information about BGP and the configured conditional advertisement route maps.

Step 7 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor-af)# copy
running-config startup-config
```

Saves this configuration change.

Example

This example shows how to configure BGP conditional advertisement:

```
switch# <userinput>configure terminal</userinput>
switch(config)# <userinput>router bgp 65536</userinput>
switch(config-router)# <userinput>neighbor 192.0.2.2 remote-as 65537</userinput>
switch(config-router-neighbor)# <userinput>address-family ipv4 unicast</userinput>
switch(config-router-neighbor-af)# <userinput>advertise-map advertise exist-map
exist</userinput>
switch(config-router-neighbor-af)# <userinput>exit</userinput>
switch(config-router-neighbor)# <userinput>exit</userinput>
switch(config-router)# <userinput>exit</userinput>
switch(config)# <userinput>route-map advertise</userinput>
switch(config-route-map)# <userinput>match as-path pathList</userinput>
switch(config-route-map)# <userinput>exit</userinput>
switch(config)# <userinput>route-map exit</userinput>
switch(config-route-map)# <userinput>match ip address prefix-list plist</userinput>
switch(config-route-map)# <userinput>exit</userinput>
switch(config)# <userinput>ip prefix-list plist permit 209.165.201.0/27</userinput>
```

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **address-family** {**ipv4** | **ipv6** } {**unicast** | **multicast**}

Example:

```
switch(config-router)# address-family
vpnv4 unicast
switch(config-router-af)#
```

Enters address family configuration mode.

Step 4 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
switch(config-router)# address-family
ipv4 unicast
switch(config-router-af)#
```

Enters address-family configuration mode.

Step 5 **redistribute** {**direct** | {**eigrp** | **isis** | **ospf** | **ospfv3** | **rip**} *instance-tag* | **static**} **route-map** *map-name*

Example:

```
switch(config-router-af)# redistribute
eigrp 201 route-map Eigrpmap
```

Redistributes routes from other protocols into BGP.

Step 6 (Optional) **default-metric** *value*

Example:

```
switch(config-router-af)# default-metric
33
```

Generates a default route into BGP.

Step 7 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-af)# copy
running-config startup-config
```

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
```

Advertise 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters global configuration mode.

Step 2 **route-map allow permit**

Example:

```
switch(config)# route-map allow permit
switch(config-route-map)#
```

Enters router map configuration mode and defines the conditions for redistributing routes.

Step 3 **exit**

Example:

```
switch(config-route-map)# exit
switch(config)#
```

Exits router map configuration mode.

Step 4 **ip route** *ip-address network-mask* **null** *null-interface-number*

Example:

```
switch(config)# ip route 192.0.2.1 255.255.255.0 null 0
```

Configures the IP address.

Step 5 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Enters BGP mode and assigns the AS number to the local BGP speaker.

Step 6 **address-family** {**ipv4** | **ipv6**} **unicast**

Example:

```
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)#
```

Enters address-family configuration mode.

Step 7 **default-information originate**

Example:

```
switch(config-router-af)# default-information originate
```

Advertises the default route.

Step 8 **redistribute static route-map allow**

Example:

```
switch(config-router-af)# redistribute static route-map allow
```

Redistributes the default route.

Step 9 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-af)# copy running-config startup-config
```

Saves this configuration change.

Configure 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.

Treat 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

[**no**] **path-attribute treat-as-withdraw** [*value* | **range** *start end*] **in**

Example:

```
switch#(config-router)# neighbor 10.20.30.40
            switch(config-router-neighbor)# path-attribute treat-as-withdraw 100 in
```

Example:

```
switch#(config-router)# neighbor 10.20.30.40
            switch(config-router-neighbor)# path-attribute treat-as-withdraw range 21 255 in
```

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.

This command is also supported for BGP template peers and BGP template peer sessions.

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

Step 1	[**no**] **path-attribute discard** [*value* | **range** *start end*] **in**
	Example:
	`switch#(config-router)# neighbor 10.20.30.40`
	`switch(config-router-neighbor)# path-attribute discard 100 in`
	Example:
	`switch#(config-router)# neighbor 10.20.30.40`
	`switch(config-router-neighbor)# path-attribute discard range 100 255 in`
		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.
		This command is also supported for BGP template peers and BGP template peer sessions.
		Note
		When the same path attribute is configured for both discard and treat-as-withdaw, treat-as-withdraw has a higher priority.

Enable or Disable 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

[**no**] **enhanced-error**

Example:

```
switch(config)# router bgp 1000
switch(config-router)# enhanced-error
```

Enables or disables BGP enhanced attribute error handling.

Display Discarded or Unknown Path Attributes

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

Command	Purpose
show bgp {**ipv4** | **ipv6**} **unicast path-attribute discard**]	Displays all prefixes for which an attribute has been discarded.
show bgp {**ipv4** | **ipv6**} **unicast path-attribute unknown**]	Displays all prefixes that have an unknown attribute.
show bgp {**ipv4** | **ipv6**} **unicast** *ip-address*	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
```

Tune 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
bestpath [**always-compare-med** | **as-pathmultipath-relax** | **compare-routerid** |**cost-community ignore** | **igp-metric ignore** |**med** {**confed** |**missing-as-worst**| **non-deterministic**}] **Example:** `switch(config-router)# bestpath always-compare-med`	Modifies the best-path algorithm. The optional parameters are as follows: • **always-compare-med** —Compares MED on paths from different autonomous systems. • **as-path multipath-relax** — 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. When configured, BGP selects the best path with the lowest MED among potential multipaths, even if they come from different ASNs. • **compare-routerid** —Compares the router IDs for identical eBGP paths. • **cost-community ignore** —Ignores the cost community for BGP best-path calculations. • **igp-metric ignore** —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). • **med confed** —Forces bestpath to do a MED comparison only between paths originated within a confederation. • **med missing-as-worst** —Treats a missing MED as the highest MED. • **med non-deterministic** —Does not always pick the best MED path from among the paths from the same autonomous system.
enforce-first-as **Example:** `switch(config-router)# enforce-first-as`	Enforces the neighbor autonomous system to be the first AS number listed in the AS_path attribute for eBGP.
log-neighbor-changes **Example:** `switch(config-router)# log-neighbor-changes`	Generates a system message when any neighbor changes state. **Note** To suppress neighbor status change messages for a specific neighbor, you can use the **log-neighbor-changes disable** command in router address-family configuration mode.

Command	Purpose
router-id *id* **Example:** `switch(config-router)# router-id 10.165.20.1`	Manually configures the router ID for this BGP speaker.
timers [**bestpath-delay** *delay* | *bgpkeepalive holdtime* | **prefix-peer-timeout** *timeout* | **bestpath-limit** *bestpath-timeout*] **Example:** `switch(config-router)# timers bestpath-limit 300`	Sets BGP timer values. The optional parameters are as follows: • *delay* —Initial best-path timeout value after a restart. The range is from 0 to 3600 seconds. The default value is 300. • *keepalive* —BGP session keepalive time. The range is from 0 to 3600 seconds. The default value is 60. • *holdtime* —BGP session hold time. The range is from 0 to 3600 seconds. The default value is180. • *timeout* —Prefix peer timeout value. The range is from 0 to 1200 seconds. The default value is 30. • *bestpath-timeout* —Bestpath timeout in seconds. The default value is 300. When a high-scale BGP setup is expected, the timeout value needs to be set to 480. You must manually reset the BGP sessions after configuring this command.

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

Command	Purpose
distance *ebgp-distance ibgp-distance local-distance* **Example:** `switch(config-router-af)# distance 20 100 200`	Sets the administrative distance for BGP. The range is from 1 to 255. The defaults are as follows: • *ebgp-distance* —20. • *ibgp-distance* —200. • *local-distance* —220. Local-distance is the administrative distance used for aggregate discard routes when they are installed in the RIB. 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.

Command	Purpose
log-neighbor-changes [**disable**] **Example:** `switch(config-router-af)# log-neighbor-changes disable`	Generates a system message when this specific neighbor changes state. The **disable** option suppresses neighbor status changes messages for this specific neighbor.

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

Command	Purpose
description *string* **Example:** `switch(config-router-neighbor)# description main site`	Sets a descriptive string for this BGP peer. The string can be up to 80 alphanumeric characters.
low-memory exempt **Example:** `switch(config-router-neighbor)# low-memory exempt`	Exempts this BGP neighbor from a possible shutdown due to a low memory condition.
transport connection-mode passive **Example:** `switch(config-router-neighbor)# transport connection-mode passive`	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.
[**no** | **default**] **remove-private-as** [**all** | **replace-as**] **Example:** `switch(config-router-neighbor)# remove-private-as`	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. The optional parameters are as follows: • **no** —Disables the command. • **default** —Moves the command to its default mode. • **all** —Removes all private-as numbers from the AS-path value. • **replace-as** —Replaces all private AS numbers with the replace-as AS-path value. See the *Guidelines and Limitations for BGP* section for additional information on this command.

Command	Purpose
update-source *interface-type number* **Example:** `switch(config-router-neighbor)# update-source ethernet 2/1`	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 **update-source** is configured.

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

Command	Purpose
allowas in **Example:** `switch(config-router-neighbor-af)# allowas in`	Allows routes that have their own AS in the AS path to be installed in the BRIB.
default-originate [**route-map** *map-name*] **Example:** `switch(config-router-neighbor-af)# default-originate`	Generates a default route to the BGP peer.
disable-peer-as-check **Example:** `switch(config-router-neighbor-af)# disable-peer-as-check`	Disables peer AS-number checking while the device advertises routes learned from one node to another node in the same AS path.
filter-list *list-name* {**in** | **out**} **Example:** `switch(config-router-neighbor-af)# filter-list BGPFilter in`	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.
prefix-list *list-name* {**in** | **out**} **Example:** `switch(config-router-neighbor-af)# prefix-list PrefixFilter in`	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.
send-community **Example:** `switch(config-router-neighbor-af)# send-community`	Sends the community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
send-community extended **Example:** `switch(config-router-neighbor-af)# send-community extended`	Sends the extended community attribute to this BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.

Command	Purpose
suppress-inactive **Example:** `switch(config-router-neighbor-af)# suppress-inactive`	Advertises the best (active) routes only to the BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
[no | default] as-override **Example:** `switch(config-router-neighbor-af)# as-override`	**no** - (Optional) Disables the command. **default** - (Optional) Moves the command to its default mode. **as-override** - While sending updates to eBGP peer, replaces in the *path* attribute all occurrences of the peer's AS number with the local AS number.

Configure 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

Step 1 switch# **configure terminal**

Enters global configuration mode.

Step 2 switch(config)# **ip prefix-list** *name* **seq** *number* **permit** *prefix-length*

Creates a prefix list to match IP packets or routes with the permit keyword.

Step 3 switch(config)# **route-map** *map-tag* **permit** *sequence-number*

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.

Step 4 switch(config-route-map)# **match ip address prefix-list** *prefix-list-name*

Matches IPv4 network routes based on a prefix list. The prefix-list name can be any alphanumeric string up to 63 characters.

Step 5 switch(config-route-map)# **set distance** *value1 value2 value3*

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 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.

Step 6 switch(config-route-map)# **exit**

Exits route-map configuration mode.

Step 7 switch(config)# **router bgp** *as-number*

Enters BGP mode and assigns the AS number to the local BGP speaker.

Step 8 switch(config-router)# **address-family** {**ipv4** | **ipv6** | **vpnv4** | **vpnv6**} **unicast**

Enters address family configuration mode.

Step 9 switch(config-router-af)# **table-map** *map-name*

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.

Note
You can also configure the **table-map** command under the VRF address-family configuration mode.

Step 10 (Optional) switch(config-router-af)# **show forwarding distribution**

Displays forwarding information distribution.

Step 11 (Optional) switch(config)# **copy running-config startup-config**

Saves the change persistently through reboots and restarts by copying the running configuration to the startup configuration.

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
  switch(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
  switch(config-router)#
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **neighbor** *ip-address* **remote-as** *as-number*

Example:

```
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.

Step 4 **address-family** {**ipv4** | **ipv6**} {**unicast** | **multicast**}

Example:

```
switch(config-router-neighbor)#
  address-family ipv4 multicast
  switch(config-router-neighbor-af)#
```

Enters address family configuration mode.

Step 5 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor-af)# copy
  running-config startup-config
```

Saves this configuration change.

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
```

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
```

Enters global configuration mode.

Step 2 **router bgp as-number**

Example:

```
switch(config)# router bgp 200
```

Enters BGP mode and assigns the autonomous system number to the local BGP speaker.

Step 3 **bmp server** *server-number*

Example:

```
switch(config-router-bmp)# bmp-server 1
```

Configures the BMP server to which BGP should send information. The server number is used as a key.

Note
You can configure up to two BMP servers.

Step 4 **address ip-address port-number port-number**

Example:

```
switch(config-router-bmp)# 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.

Step 5 **description** *string*

Example:

```
switch(config-router-bmp)# description BMPserver1
```

Configures the BMP server description. You can enter up to 256 alphanumeric characters.

Step 6 **initial-refresh** { *skip* | *delay time*}

Example:

```
switch(config-router-bmp)# 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.

Step 7 **initial-delay** *time*

Example:

```
switch(config-router-bmp)# 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.

Step 8 **stats-reporting-period** *time*

Example:

```
switch(config-router-bmp)# stats-reporting-period 50
```

Configures the time interval in which the BMP server receives the statistics report from BGP neighbors. The range is from 30 to 720 seconds, and the default is disabled.

Step 9 shutdown

Example:

```
switch(config-router-bmp)# shutdown
```

Disables the connection to the BMP server.

Step 10 **vrf** *vrf-name*

Example:

```
switch(config-router-bmp)# vrf BMP
```

Selects vrf in which BMP server is reachable.

Step 11 **update-source** <*interface-name*>

Example:

```
switch(config-router-bmp)# update-source ethernet4/2
```

Selects local interface to be used for establishing BMP server connection.

Step 12 **neighbor ip-address**

Example:

```
switch(config-router-bmp)# neighbor 192.168.1.2
```

Enters neighbor configuration mode for BGP routing and configures the neighbor IP address.

Step 13 **remote-as** *as-number*

Example:

```
switch(config-router-neighbor)# remote-as 65535
```

Configures the AS number for a remote BGP peer.

Step 14 **bmp-activate-server** *server-number*

Example:

```
switch(config-router-neighbor)# bmp-activate-server 1
```

Configures the BMP server to which a neighbor's information should be sent.

Step 15 (Optional) **show bgp bmp** *server [server-number] [detail]*

Example:

```
switch(config-router-neighbor)# show bgp bmp server
```

Displays BMP server information.

Step 16 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-neighbor)# copy running-config startup-config
```

Saves this configuration change.

BGP Local Route Leaking

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 Configure 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.

Configure 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

Step 1 **configure terminal**

Example:

```
switch-1# config terminal
                        Enter configuration commands, one per line. End with CNTL/Z.
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **vrf context** *vrf-name*

Example:

```
switch-1(config)# vrf context 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.

Step 3 **address-family** *address-family sub family*

Example:

```
switch-1(config-vrf)# address-family ipv4 unicast
                        switch-1(config-vrf-af-ipv4)#
```

Step 4 **export vrf default [***prefix-limit***] map***route-map* **allow-vpn**

Example:

```
switch-1(config-vrf-af-ipv4)# export vrf default map vpnmap1 allow-vpn
                        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.

Configure Routes Leaked from the Default-VRF to Export to a VPN

Before you begin

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

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.

Procedure

Step 1 **config terminal**

Example:

```
switch-1# config terminal
                        Enter configuration commands, one per line. End with CNTL/Z.
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **vrf context** *vrf-name*

Example:

```
switch-1(config)# vrf context 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.

Step 3 **address-family** *address-family sub family*

Example:

```
switch-1(config-vrf)# address-family ipv4 unicast
                        switch-1(config-vrf-af-ipv4)#
```

Step 4 **import vrf default [***prefix-limt***] map***route-map* **advertise-vpn**

Example:

```
switch-1(config-vrf-af-ipv4)# import vrf map vpnmap1 advertise-vpn
                        switch-1(config-vrf-af-ipv4)#
```

Configures the current VRF to allow routes imported from the default VRF to be exported to a BGP VPN.

Configure 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

Step 1 **config terminal**

Example:

```
switch-1# config terminal
                        Enter configuration commands, one per line. End with CNTL/Z.
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **vrf context** *vrf-name*

Example:

```
switch-1(config)# vrf context 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.

Step 3 **address-family** *address-family sub family*

Example:

```
switch-1(config-vrf)# address-family ipv4 unicast
                        switch-1(config-vrf-af-ipv4)#
```

Step 4 **export vrf allow-vpn**

Example:

```
switch-1(config-vrf-af-ipv4)# export vrf allow-vpn
                        nxosv2(config-vrf-af-ipv4)#
```

Configures a VRF to allow routes imported from a BGP VPM to be exported to a non-default VRF.

Configure 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

Step 1 **config terminal**

Example:

```
switch-1# config terminal
                        Enter configuration commands, one per line. End with CNTL/Z.
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **vrf context** *vrf-name*

Example:

```
switch-1(config)# vrf context 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.

Step 3 **address-family** *address-family sub family*

Example:

```
switch-1(config-vrf)# address-family ipv4 unicast
                      switch-1(config-vrf-af-ipv4)#
```

Step 4 **import vrf advertise-vpn**

Example:

```
switch-1(config-vrf-af-ipv4)# import vrf advertise-vpn
                      nxosv2(config-vrf-af-ipv4)#
```

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.

Configure 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.

Display BGP Local Route Leaking Information

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

Command	Action
show bgp vrf *vrf-name* **process**	For a default or non-default VRF, shows the enabled state (Yes or No) of the **import advertise-vpn** and **export allow-vpn** options.
show bgp vrf *vrf-name* **ipv4 unicast** *prefix*	Shows information about imported paths, including a list of destinations a route has been imported from.

BGP Graceful Shutdown

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 80.
- 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.

- Global context: Affects the entire switch and all routes processed by it. For example, readvertise all routes with the GRACEFUL_SHUTDOWN community.
- Peer context: Affects a BGP peer or a link between neighbors. For example, advertise only one link between peers with GRACEFUL_SHUTDOWN community.

Graceful Shutdown Contexts: Comparison of Contexts

Table 1: Comparison of Graceful Shutdown Contexts

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

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).

Configure 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

Step 1 **config terminal**

Example:

```
switch-1# configure terminal
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *autonomous-system-number*

Example:

```
switch-1(config)# router bgp 110
                        switch-1(config-router)#
```

Enters router configuration mode to create or configure a BGP routing process.

Step 3 **neighbor {** *ipv4-address*|*ipv6-address* **} remote-as** *as-number*

Example:

```
switch-1(config-router)# neighbor 10.0.0.3 remote-as 200
                        switch-1(config-router-neighbor)#
```

Configures the autonomous system (AS) to which the neighbor belongs.

Step 4 **graceful-shutdown activate [route-map** *map-name***]**

Example:

```
switch-1(config-router-neighbor)# graceful-shutdown activate route-map gshutPeer
                        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.

Filter 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 depreferring 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

Step 1 **configure terminal**

Example:

```
switch-1# configure terminal
                         switch-1<config)#
```

Enters global configuration mode.

Step 2 **ip community list standard** *community-list-name* **seq** *sequence-number* **{ permit | deny }** *value*

Example:

```
switch-1(config)# ip community-list standard GSHUT seq 10 permit 65535:0
                        switch-1(config)#
```

Configures a community list and permits or denies routes that have the well-known graceful shutdown community value.

Step 3 **route map** *map-tag* **{deny | permit}** *sequence-number*

Example:

```
switch-1(config)# route-map RM_GSHUT permit 10
                        switch-1(config-route-map)#
```

Configures a route map as sequence 10 and permits routes that have the GRACEFUL_SHUTDOWN community.

Step 4 **match community** *community-list-name*

Example:

```
switch-1(config-route-map)# match community GSHUT
                        switch-1(config-route-map)#
```

Configures that routes that match the IP community list GSHUT are processed by Route Policy Manager (RPM).

Step 5 **set local-preference** *local-pref-value*

Example:

```
switch-1(config-route-map)# set local-preference 10
                        switch-1(config-route-map)#
```

Configures that the routes that match the IP community list GSHUT will be given a specified local preference.

Step 6 **exit**

Example:

```
switch-1(config-route-map)# exit
                        switch-1(config)#
```

Leaves route map configuration and returns to global configuration mode.

Step 7 **router bgp** *community-list-name*

Example:

```
switch-1(config)# router bgp 100
                        switch-1(config-router)#
```

Enters router configuration mode and creates a BGP instance.

Step 8 **neighbor {** *ipv4-address*|*ipv6-address* **}**

Example:

```
switch-1(config-router)# neighbor 10.0.0.3
                        switch-1(config-router-neighbor)#
```

Enters route BGP neighbor mode for a specified neighbor.

Step 9 **address-family {** *address-family sub family* **}**

Example:

```
nxosv2(config-router-neighbor)# address-family ipv4 unicast
                          nxosv2(config-router-neighbor-af)#
```

Puts the neighbor into address family (AF) configuration mode.

Step 10 **send community**

Example:

```
nxosv2(config-router-neighbor-af)# send-community
                          nxosv2(config-router-neighbor-af)#
```

Enables BGP community exchange with the neighbor.

Step 11 **route map** *map-tag* **in**

Example:

```
nxosv2(config-router-neighbor-af)# route-map RM_GSHUT in
                          nxosv2(config-router-neighbor-af)#
```

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.

Configure 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

Step 1 **configure terminal**

Example:

```
switch-1# configure terminal
                          switch-1(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *autonomous-system-number*

Example:

```
switch-1(config)# router bgp 110
                          switch-1(config-router)#
```

Enters router configuration mode to create or configure a BGP routing process.

Step 3 **graceful-shutdown activate [route-map** *map-name***]**

Example:

```
switch-1(config-router-neighbor)# graceful-shutdown activate route-map gshutPeer
                        switch-1(config-router-neighbor)#
```

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.

The routes are advertised with the GRACEFUL_SHUTDOWN community by 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.

The devices receiving the GRACEFUL_SHUTDOWN community look at the communities of the route and optionally use the communities to apply routing policy.

Control the Preference for All Routes with the GRACEFUL_SHUTDOWN Community

Before you begin

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

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.

Procedure

Step 1 **configure terminal**

Example:

```
switch-1(config)# config terminal
                        switch-1(config)#
```

Enters global configuration mode.

Step 2 **router bgp** *autonoums-system*

Example:

```
switch-1(config)# router bgp 100
                        switch-1(config-router)#
```

Enters router configuration mode and configures a BGP routing process.

Step 3 (Optional) **no graceful-shutdown aware**

Example:

```
switch-1(config-router)# no graceful-shutdown aware
                        switch-1(config-router)#
```

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 **no** form of the command is optional for not depreferring graceful shutdown routes.

Prevent 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.

Procedure

Step 1 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.

Step 2 Delete the GRACEFUL_SHUTDOWN community through an outbound route map by following these steps:

a. Create an IP community list matches the GRACEFUL_SHUTDOWN community.

b. Create an outbound route map to match against the GRACEFUL_SHUTDOWN community.

c. 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.

Display Graceful Shutdown Information

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

Command	Action
show ip bgp community-list graceful-shutdown	Shows all entires in the BGP routing table that have the GRACEFUL_SHUTDOWN community.

Command	Action
show running-config bgp	Shows the running BGP configuration.
show running-config bgp all	Shows all information for the running BGP configuration including information about the graceful shutdown feature.
show bgp *address-family* **neighbors** *neighbor-address*	When the feature is configured for the peer, shows the following: • The state of the graceful-shutdown-activate feature for the specified neighbor • The name of any graceful shutdown route map configured for the specified neighbor
show bgp process	Shows different information depending on the context. When the graceful-shutdown-activate option is configured in peer context, shows the enabled or disabled state for the feature through `graceful-shutdown-active`. 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: • `graceful-shutdown-active` • `graceful-shutdown-aware` • `graceful-shutdown route-map`
show ip bgp *address*	For the specified address, shows the BGP routing table information, including the following: • The state of the specified address as the best path • Whether the specified address is part of the GRACEFUL_SHUTDOWN community

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
```

Configure 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.

Note From the perspective of BGP Graceful Restart, if there are idle peers during a node restart, they can potentially cause traffic loss during an ISSU because they may delay the establishment of the first best-path. It is recommended to either bring all these idle neighbors up, or configure 'shutdown' under each of them, or remove them entirely from the configuration.

Before you begin

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

Create the VRFs.

Procedure

Step 1 **configure terminal**

Example:

```
switch# configure terminal
switch(config)#
```

Enters configuration mode.

Step 2 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
switch(config-router)#
```

Creates a new BGP process with the configured autonomous system number.

Step 3 (Optional) **timers prefix-peer-timeout** *timeout*

Example:

```
switch(config-router)# timers prefix-peer-timeout 20
```

Configures the timeout value (in seconds) for BGP prefix peers. The default value is 90 seconds.

Note
This command is supported beginning with Cisco NX-OS Release 9.3(3).

Step 4 **graceful-restart**

Example:

```
switch(config-router)# graceful-restart
```

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.

Step 5 **graceful-restart** {**restart-time** *time*|**stalepath-time** *time*}

Example:

```
switch(config-router)# graceful-restart
restart-time 300
```

Configures the graceful restart timers.

The optional parameters are as follows:

- **restart-time** —Maximum time for a restart sent to the BGP peer. The range is from 1 to 3600 seconds. The default is 120.

  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.

- **stalepath-time** —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.

In NX-OS software release 10.2(1), a manual reset of a BGP session is needed for the BGP session to advertise Graceful Restart capabilities. For NX-OS software releases 10.2(2) and later, BGP sessions dynamically advertise Graceful Restart capabilities without needing to restart the BGP sessions when this command is enabled.

Step 6 **graceful-restart-helper**

Example:

```
switch(config-router)# graceful-restart
restart-time 300
```

With BGP GR disabled, the N9K itself will not necessarily preserve its own forwarding state during certain GR-capable events like SSO, BGP process restart, etc. occurring locally on the N9K. However, as a GR helper, it will support a peer that has advertised its GR capability and is restarting. This means, when the N9K detects the peering has gone down (other than a holdtimer expiration or receipt of a Notification message), the N9K will stale the routes pointing to the peer and will wait for the peer's EOR (or stalepath timeout). When the peer restarts and re-establishes its peering with the N9K, it will re-advertise all its own routes and the N9K will refresh them in its BGP and routing tables. On receipt of the EOR from the peer or the stalepath timeout (whichever occurs first), the N9K will flush any remaining stale routes from that peer. In the absence of helper mode, the N9K would instantly clear out the routes learnt from the remote peer that was restarting which could lead to traffic loss.

Step 7 (Optional) **show running-config bgp**

Example:

```
switch(config-router)# show
running-config bgp
```

Displays the BGP configuration.

Step 8 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router)# copy
running-config startup-config
```

Saves this configuration change.

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
```

Configure 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

Step 1 **configure terminal**

Example:

```
switch# configure terminal
  switch(config)#
```

Enters global configuration mode.

Step 2 **vrf context** *vrf-name*

Example:

```
switch(config)# vrf context
  RemoteOfficeVRF
  switch(config-vrf)#
```

Creates a new VRF and enters VRF configuration mode.

Step 3 **exit**

Example:

```
switch(config-vrf)# exit
  switch(config)#
```

Exits VRF configuration mode.

Step 4 **router bgp** *as-number*

Example:

```
switch(config)# router bgp 65535
  switch(config-router)#
```

Creates a new BGP process with the configured autonomous system number.

Step 5 **vrf** *vrf-name*

Example:

```
switch(config-router)# vrf
  RemoteOfficeVRF
  switch(config-router-vrf)#
```

Enters the router VRF configuration mode and associates this BGP instance with a VRF.

Step 6 **neighbor** *ip-address* **remote-as** *as-number*

Example:

```
switch(config-router-vrf)# neighbor
  209.165.201.1 remote-as 65535
  switch(config-router--vrf-neighbor)#
```

Configures the IP address and AS number for a remote BGP peer.

Step 7 (Optional) **copy running-config startup-config**

Example:

```
switch(config-router-vrf-neighbor)# copy
  running-config startup-config
```

Saves this configuration change.

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
```

Verify the Advanced BGP Configuration

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

Command	Purpose
show bgp all [**summary**] [**vrf** *vrf-name*]	Displays the BGP information for all address families.
show bgp convergence [vrf *vrf-name]*	Displays the BGP information for all address families.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **community** {**regexp** *expression* | [**community**] [**no-advertise**] [**no-export**] [**no-export-subconfed**]} [**vrf** *vrf-name*]	Displays the BGP routes that match a BGP community.
show bgp [**vrf** *vrf-name*] {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **community-list** *list-name* [**vrf** *vrf-name*]	Displays the BGP routes that match a BGP community list.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **extcommunity** {**regexp** *expression* | **generic** [**non-transitive** | **transitive**] *aa4:nn* [**exact-match**]} [**vrf** *vrf-name*]	Displays the BGP routes that match a BGP extended community.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **extcommunity-list** *list-name* [**exact-match**]} [**vrf** vrf-name]	Displays the BGP routes that match a BGP extended community list.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **extcommunity-list** *list-name* [**exact-match**]} [**vrf** vrf-name]	Displays the information for BGP route dampening. Use the **clear bgp dampening** command to clear the route flap dampening information.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] {**dampening dampened-paths** [**regexp** *expression*]} [**vrf** vrf-name]	Displays the BGP route history paths.
show bgp {**ipv4** | **ipv6** | **vpnv4** | **vpnv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **filter-list** *list-name* [**vrf** *vrf-name*]	Displays the information for the BGP filter list.

Command	Purpose
show bgp {**ipv4** | **ipv6** | **vpnv4** | **vpnv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **neighbors** [*ip-address* | *ipv6-prefix*] [**vrf** *vrf-name*]	Displays the information for BGP peers. Use the **clear bgp neighbors** command to clear these neighbors.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] {**nexthop** | **nexthop-database**} [**vrf** *vrf-name*]	Displays the information for the BGP route next hop.
show bgp paths	Displays the BGP path information.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **policy** *name* [**vrf** *vrf-name*]	Displays the BGP policy information. Use the **clear bgp policy** command to clear the policy information.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **prefix-list** *list-name* [**vrf** *vrf-name*]	Displays the BGP routes that match the prefix list.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **received-paths** [**vrf** *vrf-name*]	Displays the BGP paths stored for soft reconfiguration.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **regexp** *expression* [**vrf** *vrf-name*]	Displays the BGP routes that match the AS_path regular expression.
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **route-map** *map-name* [**vrf** *vrf-name*]	Displays the BGP routes that match the route map.
show bgp peer-policy *name* [**vrf** *vrf-name*]	Displays the information about BGP peer policies.
show bgp peer-session *name* [**vrf** *vrf-name*]	Displays the information about BGP peer sessions.
show bgp peer-template *name* [**vrf** *vrf-name*]	Displays the information about BGP peer templates. Use the **clear bgp peer-template** command to clear all neighbors in a peer template.
show bgp process	Displays the BGP process information.
show bgp {**ipv4** | **ipv6**} **unicast neighbors** *interface*	Displays information about BGP peers for the specified interface.
show ip bgp neighbors *interface-name*	Displays the interface used as a BGP peer.

Command	Purpose
show ip route *ip-address* **detail vrf all | i bw**	Displays the link bandwidth EXTCOMM fields. bw:xx (such as bw:40) in the output indicates that BGP peers are sending BGP extended attributes with the bandwidth (for weighted ECMP).
show {**ipv4** | **ipv6**} **bgp** *options*	Displays the BGP status and configuration information.
show {**ipv4** | **ipv6**} **mbgp** *options*	Displays the BGP status and configuration information.
show ipv6 routers interface *interface*	Displays the link-local address of remote IPv6 routers, which is learned through IPv6 ICMP router advertisement.
show running-configuration bgp	Displays the current running BGP configuration.

Monitor BGP Statistics

To display BGP statistics, use the following commands:

Command	Purpose
show bgp {**ipv4** | **ipv6**} {**unicast** | **multicast**} [*ip-address* | *ipv6-prefix*] **flap-statistics** [**vrf** *vrf-name*]	Displays the BGP route flap statistics. Use the **clear bgp flap-statistics** command to clear these statistics.
show bgp {**ipv4** | **ipv6**} **unicast injected-routes**	Displays injected routes in the routing table.
show bgp sessions [**vrf** *vrf-name*]	Displays the BGP sessions for all peers. Use the **clear bgp sessions** command to clear these statistics.
show bgp statistics	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: https://cisco.github.io/cisco-mibs/supportlists/nexus9000/Nexus9000MIBSupportList.html