



CHAPTER 10

Configuring Advanced BGP

This chapter describes how to configure advanced features of the Border Gateway Protocol (BGP).

This chapter includes the following sections:

- [Information About Advanced BGP, page 10-1](#)
- [Licensing Requirements for Advanced BGP, page 10-10](#)
- [Prerequisites for BGP, page 10-10](#)
- [Guidelines and Limitations for BGP, page 10-10](#)
- [Configuring Advanced BGP, page 10-10](#)
- [Verifying Advanced BGP Configuration, page 10-34](#)
- [Displaying BGP Statistics, page 10-35](#)
- [Related Topics, page 10-36](#)
- [Default Settings, page 10-36](#)
- [Default Settings, page 10-36](#)
- [Additional References, page 10-36](#)

Information About Advanced BGP

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



Note

Cisco NX-OS does not support IPv6 for BGP.

This section includes the following topics:

- [Peer Templates, page 10-2](#)
- [Authentication, page 10-2](#)
- [Route Policies and Resetting BGP Sessions, page 10-3](#)

Send document comments to nexus7k-docfeedback@cisco.com.

- eBGP, page 10-3
- iBGP, page 10-3
- Capabilities Negotiation, page 10-6
- AS Confederations, page 10-4
- Router Reflector, page 10-5
- Route Dampening, page 10-6
- Load Sharing and Multipath, page 10-6
- Route Aggregation, page 10-7
- Route Redistribution, page 10-7
- Tuning BGP, page 10-7
- Multiprotocol BGP, page 10-8
- Graceful Restart and High Availability, page 10-8
- ISSU, page 10-9

Peer Templates

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

Cisco NX-OS implements three types of peer templates:

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

Authentication

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

**Note**

The MD5 password must be identical between BGP peers.

Send document comments to nexus7k-docfeedback@cisco.com.

Route Policies and Resetting BGP Sessions

You can associate a route policy to a BGP peer. Route policies use route maps to control or modify the routes that BGP recognizes. You can configure a route policy for inbound or outbound route updates. The route policies can match on different criteria, such as a prefix or AS_path attribute, and selectively accept or deny the routes. Route policies can also modify the path attributes. See [Chapter 16, “Configuring Policy-Based Routing”](#) for more information on route policies.

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

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



Note

BGP also uses route maps for route redistribution, route aggregation, route dampening, and other features. See [Chapter 15, “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.

You should use loopback interfaces for establishing eBGP 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” section on page 10-21](#) for information on multihop, fast external failover and support for the General Time-To-Live Security Mechanism.


iBGP

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

Send document comments to nexus7k-docfeedback@cisco.com.

Figure 10-1 shows an iBGP network within a larger BGP network.

Figure 10-1 iBGP Network



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



Note

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

This section includes the following topics:

- [AS Confederations, page 10-4](#)
- [Router Reflector, page 10-5](#)

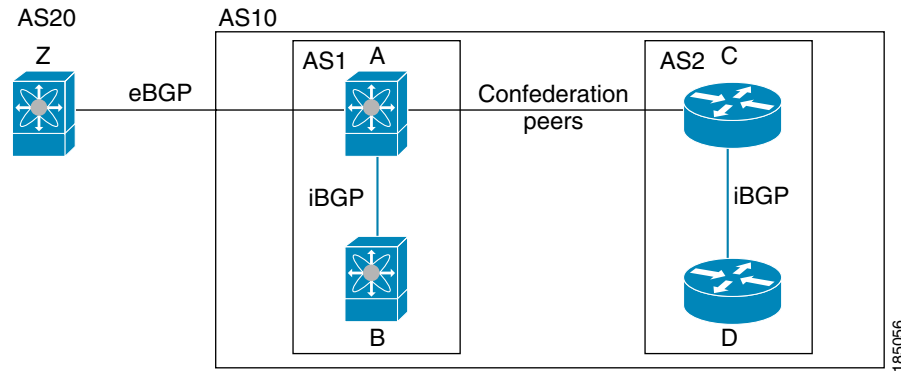
AS Confederations

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

Send document comments to nexus7k-docfeedback@cisco.com.

Figure 10-2 shows the BGP network from Figure 10-1, split into two subautonomous systems and one confederation.

Figure 10-2 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 in Figure 10-1.

Router Reflector


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

Figure 10-1 shows a simple iBGP configuration with four meshed iBGP speakers (router A, B, C, and D). Without router reflectors, when router A receives a route from an external neighbor, it advertises the route to all three iBGP neighbors.

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.

In Figure 10-3, router B is the router reflector. When the router reflector receives routes advertised from router A, it advertises (reflects) the routes to router C and D. Router A no longer has to advertise to both router C and D.

Figure 10-3 router reflector



Send document comments to nexus7k-docfeedback@cisco.com.

The router reflector and its client peers form a cluster. You do not have to configure all iBGP peers to act as client peers of the router 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 a BGP extensions 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 will attempt 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 capability negotiation.

Route Dampening

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

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

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

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



Note

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

Load Sharing and Multipath

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

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

- Weight
- Local preference
- AS_path

Send document comments to nexus7k-docfeedback@cisco.com.

- 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.



Note

Paths 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 router reflector for iBGP multipath, and the router reflector advertises the selected best-path to its peers, the next hop for the path is not modified.

Route Aggregation

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

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

Route Redistribution

You can configure BGP to redistribute static routes or routes from other protocols. You configure a route policy with the redistribution to control which routes are passed into BGP. A route policy allows you to filter routes based on attributes such as the destination, origination protocol, route type, route tag, and so on. See [Chapter 15, “Configuring Route Policy Manager,”](#) for more information.

Tuning BGP

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

This section includes the following topics:

- [BGP Timers, page 10-8](#)
- [Tuning the Best-Path Algorithm, page 10-8](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

BGP Timers

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

Tuning the Best-Path Algorithm

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

Multiprotocol BGP

BGP on Cisco NX-OS supports multiple address families. Multiprotocol BGP (MBGP) carries a 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 routes.

Use the router address-family and neighbor address-family configuration modes to support multiprotocol BGP configurations.

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.

See the *Cisco Nexus 7000 Series NX-OS Multicast Routing Configuration Guide, Release 4.0* for multicast configuration examples using MBGP.

Graceful Restart and High Availability

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

You can use nonstop forwarding 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 non-graceful 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. Before the switchover occurs, BGP announces that a graceful restart is starting and that BGP will be unavailable for some time. During the switchover, BGP uses nonstop forwarding to forward traffic based on the information in the FIB, and the system is not taken out of the network topology. The router that is restarted marks these routes from its peers as stale.

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.

Send document comments to nexus7k-docfeedback@cisco.com.

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

ISSU

Cisco NX-OS supports in-service software upgrades (ISSU). ISSU allows you to upgrade software without impacting forwarding.



Note

You must enable graceful restart to support ISSU for BGP.

BGP uses a peer hold timer to tear down sessions for peers that have become inactive and stopped responding. As part of the ISSU process, BGP control packets might not be received or transmitted during the switchover and peers may notice loss of keepalive messages. However, as long as the hold time is greater than the switchover time, the peers should not tear down sessions with the local router.

Once switchover occurs, the peers receive TCP connection resets from the new active TCP on the local router. If you enabled graceful restart, the peers treat the resets as an indication that the router restarted and initiate the graceful-restart helper procedures.



Note

Cisco NX-OS cannot guarantee ISSU if the negotiated hold time between BGP peers is less than the system switchover time (approximately 29 seconds).

BGP supports ISSU in the following ways:

- If you disable graceful restart, Cisco NX-OS issues a warning that ISSU cannot be supported with this configuration.
- If you configure the hold time to be less than the system switchover time, Cisco NX-OS issues a similar warning. If the peer negotiates a shorter hold time, Cisco NX-OS logs a message.
- When Cisco NX-OS starts an ISSU, BGP checks both the graceful restart status and thenegotiated hold time for all active peers. Cisco NX-OS issues appropriate warnings and ends the ISSU process if graceful restart is disabled or the negotiated hold times for the active peers is less than the system switchover time.

If the negotiated hold time is less than the system switchover time, you must reconfigure the hold timers on the BGP peers to be greater than the system switchover time and restart the BGP sessions before you can proceed with the ISSU.

Virtualization Support

Cisco NX-OS supports multiple instances of the BGP protocol that run on the same system. BGP supports Virtual Routing and Forwarding instances (VRFs) which exist within virtual device contexts (VDCs). You can configure one BGP instance in a VDC, but you can have multiple VDCs on the system.

By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. See the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0* and [Chapter 14, “Configuring Layer 3 Virtualization.”](#)

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Licensing Requirements for Advanced BGP

The following table shows the licensing requirements for this feature:

Product	License Requirement
NX-OS	BGP requires an Enterprise Services license. For a complete explanation of the NX-OS licensing scheme and how to obtain and apply licenses, see the <i>Cisco Nexus 7000 Series NX-OS Licensing Guide, Release 4.0</i>

Prerequisites for BGP

BGP has the following prerequisites:

- You must enable the BGP feature (see the [“Enabling the BGP Feature”](#) section on page 9-9).
- 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 IGP, static route or 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 BGP

BGP has the following guidelines and limitations:

- 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 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 redistribution.
- Configure the BGP router ID within a VRF.
- If you decrease the keepalive and hold timer values, the network might experience session flaps.
- If you configure VDCs, install the Advanced Services license and enter the desired VDC (see the *Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0*).

Configuring Advanced BGP

This section describes how to configure advanced BGP and includes the following topics:

- [Configuring BGP Session Templates, page 10-11](#)
- [Configuring BGP Peer-Policy Templates, page 10-14](#)
- [Configuring BGP Peer Templates, page 10-16](#)

Send document comments to nexus7k-docfeedback@cisco.com.

- [Configuring Prefix Peering, page 10-18](#)
- [Configuring BGP Authentication, page 10-19](#)
- [Resetting a BGP Session, page 10-20](#)
- [Modifying the Next-Hop Address, page 10-20](#)
- [Disabling Capabilities Negotiation, page 10-21](#)
- [Configuring eBGP, page 10-21](#)
- [Configuring AS Confederations, page 10-22](#)
- [Configuring Router Reflector, page 10-23](#)
- [Configuring Route Dampening, page 10-25](#)
- [Configuring Load Sharing and ECMP, page 10-25](#)
- [Configuring Maximum Prefixes, page 10-25](#)
- [Configuring Dynamic Capability, page 10-26](#)
- [Configuring Aggregate Addresses, page 10-26](#)
- [Configuring Route Redistribution, page 10-27](#)
- [Tuning BGP, page 10-28](#)
- [Configuring a Graceful Restart, page 10-31](#)
- [Configuring Virtualization, page 10-33](#)

**Note**

If you are familiar with the Cisco IOS CLI, be aware that the Cisco NX-OS commands for this feature might differ from the Cisco IOS commands that you would use.

Configuring BGP Session Templates

You can use BGP session templates to simplify 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

Ensure that you have enabled the BGP feature (see the [“Enabling the BGP Feature”](#) section on page 9-9).

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

Send document comments to nexus7k-docfeedback@cisco.com.

**Note**

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

SUMMARY STEPS

1. **config t**
2. **router bgp** *autonomous-system-number*
3. **template peer-session** *template-name*
4. Add appropriate attributes to the session template.
5. **exit**
6. **neighbor** *ip-address* **remote-as** *as-number*
7. **inherit peer-session** *template-name*
8. Add appropriate neighbor attributes.
9. **show bgp peer-session** *template-name*
10. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router bgp <i>autonomous-system-number</i> Example: switch(config)# router bgp 45000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	template peer-session <i>template-name</i> Example: switch(config-router)# template peer-session BaseSession switch(config-router-stmp)#	Enters peer-session template configuration mode.
Step 4	password <i>number password</i> Example: switch(config-router-stmp)# password 0 test	(Optional) Adds the clear text password <i>test</i> to the neighbor. The password is stored and displayed in type 3 encrypted form (3DES).
Step 5	timers <i>keepalive hold</i> Example: switch(config-router-stmp)# timers 30 90	(Optional) Adds the BGP keepalive and holdtimer values to the peer-session template. The default keepalive interval is 60. The default hold time is 180.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
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 40000 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	description text Example: switch(config-router-neighbor)# description Peer Router A switch(config-router-neighbor)	(Optional) Adds a description for the neighbor.
Step 10	show bgp peer-session template-name Example: switch(config-router-neighbor)# show bgp peer-session BaseSession	(Optional) Displays the peer-policy template.
Step 11	copy running-config startup-config Example: switch(config-router-neighbor)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0* for details on all commands available in the template.

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

```
switch# config t
switch(config)# router bgp 45000
switch(config-router)# template peer-session BaseSession
switch(config-router-stmp)# timers 30 90
switch(config-router-stmp)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 40000
switch(config-router-neighbor)# inherit peer-session BaseSession
switch(config-router-neighbor)# description Peer Router A
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor)# copy running-config startup-config
```

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring BGP Peer-Policy Templates

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

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

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the “Enabling the BGP Feature” section on page 9-9).

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



Note

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

SUMMARY STEPS

1. **config t**
2. **router bgp** *autonomous-system-number*
3. **template peer-policy** *template-name*
4. Add appropriate attributes to the policy template.
5. **exit**
6. **neighbor ip-address remote-as** *as-number*
7. **address-family** {*ipv4* | *ipv6*} {**multicast** | **unicast**}
8. **inherit peer-policy** *template-name preference*
9. **show bgp peer-policy** *template-name*
10. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router bgp <i>autonomous-system-number</i> Example: switch(config)# router bgp 45000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 3	template peer-policy <i>template-name</i> Example: <pre>switch(config-router)# template peer-policy BasePolicy switch(config-router-ptmp)#</pre>	Creates a peer-policy template.
Step 4	advertise-active-only Example: <pre>switch(config-router-ptmp)# advertise-active-only</pre>	(Optional) Advertises only active routes to the peer.
Step 5	maximum-prefix <i>number</i> Example: <pre>switch(config-router-ptmp)# maximum-prefix 20</pre>	(Optional) Sets the maximum number of prefixes allowed from this peer.
Step 6	exit Example: <pre>switch(config-router-ptmp)# exit switch(config-router)#</pre>	Exits peer-policy template configuration mode.
Step 7	neighbor ip-address remote-as as-number Example: <pre>switch(config-router)# neighbor 192.168.1.2 remote-as 40000 switch(config-router-neighbor)#</pre>	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.
Step 8	address-family {ipv4 ipv6} {multicast unicast} Example: <pre>switch(config-router-neighbor)# address-family ipv4 unicast switch(config-router-neighbor-af)#</pre>	Enters global address family configuration mode for the IPv4 address family.
Step 9	inherit peer-policy <i>template-name</i> <i>preference</i> Example: <pre>switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1</pre>	Applies a peer-policy template to the peer address family configuration and assigns the preference value for this peer policy.
Step 10	show bgp peer-policy <i>template-name</i> Example: <pre>switch(config-router-neighbor-af)# show bgp peer-policy BasePolicy</pre>	(Optional) Displays the peer-policy template.
Step 11	copy running-config startup-config Example: <pre>switch(config-router-neighbor-af)# copy running-config startup-config</pre>	(Optional) Saves this configuration change.

Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0* for details on all commands available in the template.

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

Send document comments to nexus7k-docfeedback@cisco.com.

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

```
switch# config t
switch(config)# router bgp 40000
switch(config-router)# template peer-session BasePolicy
switch(config-router-ptmp)# maximum-prefix 20
switch(config-router-ptmp)# exit
switch(config-router)# neighbor 192.168.1.1 remote-as 45000
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy
switch(config-router-neighbor-af)# copy running-config startup-config
```

Configuring BGP Peer Templates

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

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the “[Enabling the BGP Feature](#)” section on page 9-9).

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



Note

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

SUMMARY STEPS

1. **config t**
2. **router bgp** *autonomous-system-number*
3. **template peer** *template-name*
4. Add appropriate attributes to the peer template.
5. **exit**
6. **neighbor** *ip-address*
7. **inherit peer** *template-name*
8. Add appropriate neighbor attributes.
9. **show bgp peer-template** *template-name*
10. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router bgp <i>autonomous-system-number</i> Example: switch(config)# router bgp 45000	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
Step 3	template peer <i>template-name</i> Example: switch(config-router)# template peer BasePeer switch(config-router-neighbor)#	Enters peer template configuration mode.
Step 4	inherit peer-session <i>template-name</i> Example: switch(config-router-neighbor)# inherit peer-session BaseSession	(Optional) Inherits a peer-session template in the peer template.
Step 5	address-family { ipv4 ipv6 } { multicast unicast } Example: switch(config-router-neighbor)# address-family ipv4 unicast switch(config-router-neighbor-af)#	(Optional) Configures the global address family configuration mode for the IPv4 address family.
Step 6	inherit peer-policy <i>template-name</i> <i>preference</i> Example: switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1	(Optional) Applies a peer-policy template to the neighbor address family configuration and assigns the preference value for this peer policy.
Step 7	exit Example: switch(config-router-neighbor-af)# exit switch(config-router-neighbor)#	Exits BGP neighbor address family configuration mode.
Step 8	timers <i>keepalive hold</i> Example: switch(config-router-neighbor)# timers 45 100	(Optional) 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 switch(config-router)#	Exits BGP peer template configuration mode.
Step 10	neighbor <i>ip-address</i> remote-as <i>as-number</i> Example: switch(config-router)# neighbor 192.168.1.2 remote-as 40000 switch(config-router-neighbor)#	Places the router in neighbor configuration mode for BGP routing and configures the neighbor IP address.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 11	inherit peer <i>template-name</i> Example: switch(config-router-neighbor)# inherit peer BasePeer	Inherits the peer template.
Step 12	timers <i>keepalive hold</i> Example: switch(config-router-neighbor)# timers 60 120	(Optional) 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	show bgp peer-template <i>template-name</i> Example: switch(config-router-neighbor-af)# show bgp peer-template BasePeer	(Optional) Displays the peer template.
Step 14	copy running-config startup-config Example: switch(config-router-neighbor-af)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **show bgp neighbor** command to see the template applied. See the *Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0* for details on all commands available in the template.

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

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

```
switch# config t
switch(config)# router bgp 45000
switch(config-router)# template peer BasePeer
switch(config-router-neighbor)# inherit peer-session BaseSession
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# inherit peer-policy BasePolicy 1
switch(config-router-neighbor-af)# exit
switch(config-router-neighbor)# exit
switch(config-router)# neighbor 192.168.1.2 remote-as 40000
switch(config-router-neighbor)# inherit peer BasePeer
switch(config-router-neighbor)# copy running-config startup-config
```

Configuring Prefix Peering

BGP supports the definition of a set of peers using a prefix for both IPv4 and IPv6. This simplifies the configuration even further than using templates because you do not need to add each neighbor to the configuration.

When defining a prefix peering, you must specify the remote AS number along with the prefix. BGP accepts any peer connecting from that prefix and autonomous system as long as 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 defined prefix peer time-out value. This helps network stability by allowing an established peer to reset and reconnect without danger of being blocked because other peers have consumed all slots for that prefix peering. The default setting for prefix-peer-time-out is 30 seconds.

Send document comments to nexus7k-docfeedback@cisco.com.

To configure BGP prefix peering time-out value, use the following command in router configuration mode:

Command	Purpose
timers prefix-peer-timeout <i>value</i> Example: switch(config-router-neighbor)# timers prefix-peer-timeout 120	<ul style="list-style-type: none"> Configures the time-out value for prefix peering. The range is from 0 to 1200 seconds. The default value is 30.

To configure the maximum number of peers, use the following commands in neighbor configuration mode:

Command	Purpose
maximum-peers <i>value</i> Example: switch(config-router-neighbor)# timers prefix-peer-timeout 120	Configures the maximum number of peers for this prefix peering. The range is from 1 to 1000.

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

```
switch(config)# router bgp 1
switch(config-router)# timers prefix-peer-timeout 120
switch(config-router)# neighbor 10.100.200.0/24 remote-as 1
switch(config-router-neighbor)# maximum-peers 10
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

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

Configuring BGP Authentication

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

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

Command	Purpose
password [0 3 7] <i>string</i> Example: switch(config-router-neighbor)# password BGPpassword	Configures an MD5 password for BGP neighbor sessions.

Send document comments to nexus7k-docfeedback@cisco.com.

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:

Command	Purpose
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.

To reset a BGP neighbor session, use the following command in any mode:

Command	Purpose
clear bgp {ip 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.

Modifying the Next-Hop Address

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

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

You can modify the next-hop address by configuring the following parameters in neighbor address-family configuration mode:

Command	Purpose
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.
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

Send document comments to nexus7k-docfeedback@cisco.com.

Disabling Capabilities Negotiation

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

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

Command	Purpose
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.

Configuring eBGP

This section includes the following topics:

- [Disabling eBGP Single-Hop Checking, page 10-21](#)
- [Configuring eBGP Multihop, page 10-21](#)
- [Disabling a Fast External Failover, page 10-22](#)
- [Configuring AS Confederations, page 10-22](#)

Disabling eBGP Single-Hop Checking

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

To disable checking whether a single-hop eBGP peer is directly connected, use the following command in neighbor configuration mode:

Command	Purpose
disable-connected-check Example: switch(config-router-neighbor)# soft-reconfiguration inbound	Disables checking whether a single-hop eBGP peer is directly connected. You must manually reset the BGP sessions after configuring this command.

Configuring eBGP Multihop

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

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

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
ebgp-multihop <i>ttl-value</i> 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 configuring this command.

Disabling a Fast External Failover

Typically, when a BGP router loses connectivity to a directly connected eBGP peer, BGP triggers a fast external failover by resetting the eBGP session to the peer. You can disable this fast external failover to limit the instability caused by link flaps.

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

Command	Purpose
no fast-external-failover Example: switch(config-router)# no fast-external-failover	Disables a fast external failover for eBGP peers. Enabled by default.

Configuring AS Confederations

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

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

Command	Purpose
confederation identifier <i>as-number</i> Example: switch(config-router)# confederation identifier 4000	Configures a confederation identifier for an AS confederation. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

To configure the autonomous systems that belong to the AS confederation, use the following command in router configuration mode:

Command	Purpose
bgp confederation peers <i>as-number</i> <i>[as-number2...]</i> Example: switch(config-router)# bgp confederation peers 5 33 44	Specifies a list of autonomous systems that belong to the confederation. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring Router Reflector

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the “[Enabling the BGP Feature](#)” section on page 9-9).

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

SUMMARY STEPS

1. `config t`
2. `router bgp as-number`
3. `cluster-id cluster-id`
4. `address-family {ipv4 | ipv6} {unicast | multicast}`
5. `client-to-client reflection`
6. `exit`
7. `neighbor ip-address remote-as as-number`
8. `address-family {ipv4 | ipv6} {unicast | multicast}`
9. `route-reflector-client`
10. `show bgp {ip | ipv6} {unicast | multicast} as-number`
11. `copy running-config startup-config`

DETAILED STEPS

	Command or Action	Purpose
Step 1	<code>config t</code> Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	<code>router bgp as-number</code> Example: switch(config)# router bgp 45000 switch(config-router)#	Enters BGP mode and assigns the autonomous system number to the local BGP speaker.
Step 3	<code>cluster-id cluster-id</code> Example: switch(config-router)# cluster-id 192.0.2.1	Configures the local router as one of the router 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.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command or Action	Purpose
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	client-to-client reflection Example: switch(config-router-af)# client-to-client reflection	(Optional) 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-neighbor)# exit switch(config-router)#	Exits router address configuration mode.
Step 7	neighbor <i>ip-address</i> remote-as <i>as-number</i> Example: switch(config-router)# neighbor 192.0.2.10 remote-as 40000 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 switch as a BGP router reflector and configures the neighbor as its client. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
Step 10	show bgp { ip ipv6 } { unicast multicast } neighbors Example: switch(config-router-neighbor-af)# show bgp ip unicast neighbors	(Optional) Displays the BGP peers.
Step 11	copy running-config startup-config Example: switch(config-router-neighbor-af)# copy running-config startup-config	(Optional) Saves this configuration change.

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

```
switch(config)# router bgp 45000
switch(config-router)# neighbor 192.0.2.10 remote-as 40000
switch(config-router-neighbor)# address-family ip unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)# copy running-config startup-config
```

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring Route Dampening

You can configure route dampening to minimize route flaps propagating through your iBGP network. To configure route dampening, use the following command in address-family or VRF address family configuration mode:

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

Configuring Load Sharing and ECMP

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

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

Command	Purpose
maximum-paths [<i>ibgp</i>] <i>maxpaths</i> Example: switch(config-router-af)# maximum-paths 12	Configures the maximum number of equal-cost paths for load sharing. The range is from 1 to 16. The default is 8.

Configuring Maximum Prefixes

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

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

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
maximum-prefix <i>maximum</i> [<i>threshold</i>] [<i>restart time</i> <i>warming-only</i>] Example: switch(config-router-neighbor-af)# maximum-paths 12	Configures the maximum number of prefixes from a peer. The parameter ranges are as follows: <ul style="list-style-type: none"> • <i>maximum</i>—The range is from 1 to 300000. • <i>Threshold</i>—The range is from 1 to 100 percent. The default is 75 percent. • <i>time</i>—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 limit is exceeded.

Configuring Dynamic Capability

You can configure dynamic capability for a BGP peer.

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

Command	Purpose
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. Disabled by default.

Configuring Aggregate Addresses

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

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

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
<pre>aggregate-address ip-prefix/length [as-set] [summary-only] [advertise-map map-name] [attribute-map map-name] [suppress-map map-name]</pre> <p>Example:</p> <pre>switch(config-router-af)# aggregate-address 192.0.2.0/8 as-set</pre>	<p>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:</p> <ul style="list-style-type: none"> • 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 filters more specific routes.

Configuring Route Redistribution

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the [“Enabling the BGP Feature”](#) section on page 9-9).

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

SUMMARY STEPS

1. `config t`
2. `router bgp as-number`
3. `address-family {ipv4 | ipv6} {unicast | multicast}`
4. `redistribute {direct | eigrp as | isis id | ospf id | ospfv3 id | rip id | static route-map map-name}`
5. `default-metric value`
6. `exit`
7. `copy running-config startup-config`

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router bgp as-number Example: switch(config)# router bgp 45000 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 ipv4 unicast switch(config-router-af)#	Enters address family configuration mode.
Step 4	redistribute {direct eigrp as isis id ospf id ospfv3 id rip id static direct} route-map map-name Example: switch(config-router-af)# redistribute eigrp 201 route-map Eigrpmap	Redistributes routes from other protocols into BGP. See the “ Configuring Route Maps ” section on page 15-9 for more information about route maps.
Step 5	default-metric value Example: switch(config-router-af)# default-metric 33	(Optional) Generates a default route into BGP.
Step 6	copy running-config startup-config Example: switch(config-router-af)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to redistribute EIGRP into BGP:

```
switch# config t
switch(config)# router bgpEnterprise
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
```

Tuning BGP

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

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

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
<pre>bestpath [always-compare-med compare-routerid med {missing-as-worst non-deterministic}]</pre> <p>Example: switch(config-router)# bestpath always-compare-med</p>	<p>Modifies the bestpath algorithm. The optional parameters are as follows:</p> <ul style="list-style-type: none"> • always-compare-med—Compares MED on paths from different autonomous systems. • compare-routerid—Compares the router IDs for identical eBGP paths. • 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.
<pre>enforce-first-as</pre> <p>Example: switch(config-router)# enforce-first-as</p>	<p>Enforces the neighbor autonomous system to be the first AS number listed in the AS_path attribute for eBGP.</p>
<pre>log-neighbor-changes</pre> <p>Example: switch(config-router)# log-neighbor-changes</p>	<p>Generates a system message when a neighbor changes state.</p>
<pre>router-id id</pre> <p>Example: switch(config-router)# router-id 209.165.20.1</p>	<p>Manually configures the router ID for this BGP speaker.</p>
<pre>timers [bestpath-delay delay bgp keepalive holdtime prefix-peer-timeout timeout]</pre> <p>Example: switch(config-router)# timers bgp 90 270</p>	<p>Sets the BGP timer values. The optional parameters are as follows:</p> <ul style="list-style-type: none"> • <i>delay</i>—Initial best-path timeout value after a restart. The range is from 0 to 3600 seconds. The default value is 300. • <i>keepalive</i>—BGP session keepalive time. The range is from 0 to 3600 seconds. The default value is 60. • <i>holdtime</i>—BGP session hold time. The range is from 0 to 3600 seconds. The default value is 180. • <i>timeout</i>—Prefix peer timeout value. The range is from 0 to 1200 seconds. The default value is 30. <p>You must manually reset the BGP sessions after configuring this command.</p>

Send document comments to nexus7k-docfeedback@cisco.com.

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

Command	Purpose
distance <i>ebgp-distance ibgp distance local-distance</i> Example: <pre>switch(config-router-af)# distance 20 100 200</pre>	Sets the administrative distance for BGP. The range is from 1 to 255. The defaults are as follows: <ul style="list-style-type: none"> • 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.

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

Command	Purpose
description <i>string</i> Example: <pre>switch(config-router-neighbor)# description main site</pre>	Sets a descriptive string for this BGP peer. The string can be up to 80 alphanumeric characters.
transport connection-mode passive Example: <pre>switch(config-router-neighbor)# transport connection-mode passive</pre>	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.
remove-private-as Example: <pre>switch(config-router-neighbor)# remove-private-as</pre>	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.
update-source <i>interface-type number</i> Example: <pre>switch(config-router-neighbor)# update-source ethernet 2/1</pre>	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.

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

Command	Purpose
suppress-inactive Example: <pre>switch(config-router-neighbor-af)# suppress-inactive</pre>	Advertises the best (active) routes only to the BGP peer. This command triggers an automatic soft clear or refresh of BGP neighbor sessions.
default-originate [route-map <i>map-name</i>] Example: <pre>switch(config-router-neighbor-af)# default-originate</pre>	Generates a default route to the BGP peer.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
filter-list <i>list-name</i> { 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 <i>list-name</i> { 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.

Configuring a Graceful Restart

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the [“Enabling the BGP Feature”](#) section on page 9-9).

Create the VDCs and VRFs.

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

SUMMARY STEPS

1. **config t**
2. **router bgp** *as-number*
3. **graceful-restart**
4. **graceful-restart** [**restart-time** *time* | **stalepath-time** *time*]
5. **graceful-restart-helper**
6. **show running-config bgp**
7. **copy running-config startup-config**

Send document comments to nexus7k-docfeedback@cisco.com.

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	router bgp as-number Example: switch(config)# router bgp 201 switch(config-router)#	Creates a new BGP process with the configured autonomous system number.
Step 3	graceful-restart Example: switch(config-router)# graceful-restart	Enables a graceful restart and the graceful restart helper functionality. Enabled by default. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
Step 4	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: <ul style="list-style-type: none"> • restart-time—Maximum time for a restart sent to the BGP peer. The range is from 1 to 3600 seconds. The default is 120. • stalepath-time—Maximum time that BGP will keep the stale routes from the restarting BGP peer. The range is from 1 to 3600 seconds. The default is 300. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
Step 5	graceful-restart-helper Example: switch(config-router)# graceful-restart-helper	Enables the graceful restart helper functionality. Use this command if you have disabled graceful restart but you still want to enable graceful restart helper functionality. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
Step 6	show running-config bgp Example: switch(config-router)# show running-config bgp	(Optional) Displays the BGP configuration.
Step 7	copy running-config startup-config Example: switch(config-router)# copy running-config startup-config	(Optional) Saves this configuration change.

The following example shows how to enable a graceful restart:

```
switch# config t
switch(config)# router bgp 201
switch(config-router)# graceful-restart
switch(config-router)# copy running-config startup-config
```

Send document comments to nexus7k-docfeedback@cisco.com.

Configuring Virtualization

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

BEFORE YOU BEGIN

Ensure that you have enabled the BGP feature (see the “[Enabling the BGP Feature](#)” section on page 9-9). Ensure that you are in the correct VDC (or use the **switchto vdc** command).

SUMMARY STEPS

1. **config t**
2. **vrf context** *vrf-name*
3. **exit**
4. **router bgp** *as-number*
5. **vrf** *vrf-name*
6. **neighbor** *ip-address* **remote-as** *as-number*
7. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	vrf context <i>vrf-name</i> 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 <i>as-number</i> Example: switch(config)# router bgp 201 switch(config-router)#	Creates a new BGP process with the configured autonomous system number.

Send document comments to nexus7k-docfeedback@cisco.com.

	Command	Purpose
Step 5	vrf <i>vrf-name</i> 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 <i>ip-address</i> remote-as <i>as-number</i> Example: switch(config-router-vrf)# neighbor 209.165.201.1 remote-as 45000 switch(config-router--vrf-neighbor)#	Configures the IP address and AS number for a remote BGP peer.
Step 7	copy running-config startup-config Example: switch(config-router-vrf-neighbor)# copy running-config startup-config	(Optional) Saves this configuration change.

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

```
switch# config t
switch(config)# vrf context NewVRF
switch(config-vrf)# exit
switch(config)# router bgp 201
switch(config-router)# vrf NewVRF
switch(config-router-vrf)# neighbor 209.165.201.1 remote-as 45000
switch(config-router-vrf-neighbor)# copy running-config startup-config
```

Verifying Advanced BGP Configuration

To verify the BGP configuration, use the following commands:

Command	Purpose
show bgp [<i>vrf vrf-name</i>] all [<i>summary</i>]	Displays the BGP information for all address families.
show bgp [<i>vrf vrf-name</i>] convergence	Displays the BGP information for all address families.
show bgp [<i>vrf vrf-name</i>] {ip ipv6} {unicast multicast} [<i>ip-address ipv6-prefix</i>] community {regexp expression [community] [no-advertise] [no-export] [no-export-subconfed]}	Displays the BGP routes that match a BGP community.
show bgp [<i>vrf vrf-name</i>] {ip ipv6} {unicast multicast} [<i>ip-address ipv6-prefix</i>] community-list <i>list-name</i>	Displays the BGP routes that match a BGP community list.
show bgp [<i>vrf vrf-name</i>] {ip ipv6} {unicast multicast} [<i>ip-address ipv6-prefix</i>] {dampening dampened-paths [regexp expression]}	Displays the information for BGP route dampening. Use the clear bgp dampening command to clear the route flap dampening information.

Send document comments to nexus7k-docfeedback@cisco.com.

Command	Purpose
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] history-paths [<i>regex expression</i>]	Displays the BGP route history paths.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] filter-list <i>list-name</i>	Displays the information for BGP filter list.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] neighbors [<i>ip-address</i> <i>ipv6-prefix</i>]	Displays the information for BGP peers. Use the clear bgp neighbors command to clear these neighbors.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] { nexthop nexthop-database }	Displays the information for the BGP route next-hop.
show bgp paths	Displays the BGP path information.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] policy name	Displays the BGP policy information. Use the clear bgp policy command to clear the policy information.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] prefix-list <i>list-name</i>	Displays the BGP routes that match the prefix list.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] received-paths	Displays the BGP paths stored for soft reconfiguration.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] regex expression	Displays the BGP routes that match the AS_path regular expression.
show bgp [<i>vrf vrf-name</i>] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] route-map <i>map-name</i>	Displays the BGP routes that match the route map.
show bgp [<i>vrf vrf-name</i>] peer-policy name	Displays the information about BGP peer policies.
show bgp [<i>vrf vrf-name</i>] peer-session name	Displays the information about BGP peer sessions.
show bgp [<i>vrf vrf-name</i>] peer-template 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 running-configuration bgp	Displays the current running BGP configuration.

Displaying BGP Statistics

To display BGP statistics, use the following commands:

[Send document comments to nexus7k-docfeedback@cisco.com.](mailto:nexus7k-docfeedback@cisco.com)

Command	Purpose
show bgp [vrf vrf-name] { ip ipv6 } { unicast multicast } [<i>ip-address</i> <i>ipv6-prefix</i>] flap-statistics	Displays the BGP route flap statistics. Use the clear bgp flap-statistics command to clear these statistics.
show bgp [vrf vrf-name] sessions	Displays the BGP sessions for all peers. Use the clear bgp sessions command to clear these statistics.
show bgp [vrf vrf-name] sessions	Displays the BGP sessions for all peers. Use the clear bgp sessions command to clear these statistics.
show bgp statistics	Displays the BGP statistics.

Related Topics

The following topics can give more information on BGP:

- [Chapter 9, “Configuring Basic BGP”](#)
- [Chapter 15, “Configuring Route Policy Manager”](#)

Default Settings

Table 10-1 lists the default settings for BGP parameters.

Table 10-1 Default BGP Parameters

Parameters	Default
BGP feature	disabled
keep alive interval	60 seconds
hold timer	180 seconds

Additional References

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

- [Related Documents, page 10-37](#)
- [MIBs, page 10-37](#)

Send document comments to nexus7k-docfeedback@cisco.com.

Related Documents

Related Topic	Document Title
BGP CLI commands	<i>Cisco Nexus 7000 Series NX-OS Unicast Routing Command Reference, Release 4.0</i>
VDCs and VRFs	<i>Cisco Nexus 7000 Series NX-OS Virtual Device Context Configuration Guide, Release 4.0</i>

MIBs

MIBs	MIBs Link
BGP4-MIB	To locate and download MIBs, go to the following URL:
CISCO-BGP4-MIB	http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml

Send document comments to nexus7k-docfeedback@cisco.com.