



CHAPTER 9

Configuring Basic BGP

This chapter describes how to configure Border Gateway Protocol (BGP) on a Cisco NX-OS device.

This chapter includes the following sections:

- [Information About Basic BGP, page 9-1](#)
- [Licensing Requirements for Basic BGP, page 9-7](#)
- [Prerequisites for BGP, page 9-7](#)
- [Guidelines and Limitations for BGP, page 9-7](#)
- [CLI Configuration Modes, page 9-7](#)
- [Configuring Basic BGP, page 9-9](#)
- [Verifying Basic BGP Configuration, page 9-18](#)
- [Displaying BGP Statistics, page 9-19](#)
- [Basic BGP Example Configuration, page 9-20](#)
- [Related Topics, page 9-20](#)
- [Where to Go Next, page 9-20](#)
- [Default Settings, page 9-21](#)
- [Additional References, page 9-21](#)
- [Feature History for BGP, page 9-21](#)

Information About Basic BGP

Cisco NX-OS supports BGP version 4, which 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.

BGP uses a path-vector routing algorithm to exchange routing information between BGP-enabled networking devices or *BGP speakers*. Based on this information, each BGP speaker determines a path to reach a particular destination while detecting and avoiding paths with routing loops. The routing information includes the actual route prefix for a destination, the path of autonomous systems to the destination, and additional path attributes.

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

BGP selects a single path, by default, as the best path to a destination host or network. Each path carries well-known mandatory, well-known discretionary, and optional transitive attributes that are used in BGP best-path analysis. You can influence BGP path selection by altering some of these attributes by configuring BGP policies. See the “[Route Policies and Resetting BGP Sessions](#)” section on page 10-3 for more information.

BGP also supports load balancing or equal-cost multipath (ECMP). See the “[Load Sharing and Multipath](#)” section on page 10-6 for more information.



Note

Cisco NX-OS does not support IPv6 for BGP.

To deploy and configure basic BGP in your network, you should understand the following concepts:

- [BGP Autonomous Systems](#), page 9-2
- [Administrative Distance](#), page 9-2
- [BGP Peers](#), page 9-3
- [BGP Router Identifier](#), page 9-3
- [BGP Path Selection](#), page 9-3
- [BGP and the Unicast RIB](#), page 9-6
- [BGP Virtualization](#), page 9-6

BGP Autonomous Systems

An *autonomous system* (AS) is a network controlled by a single technical administration entity. An autonomous system forms a routing domain with one or more interior gateway protocols (IGPs) and a consistent set of routing policies. BGP supports 16-bit and 32-bit autonomous system numbers. For more information, see the “[Autonomous Systems](#)” section on page 1-5.

Separate BGP autonomous systems dynamically exchange routing information through external BGP (eBGP) peering sessions. BGP speakers within the same autonomous system can exchange routing information through internal BGP (iBGP) peering sessions.

Administrative Distance

An *administrative distance* is a rating of the trustworthiness of a routing information source. By default, BGP uses the administrative distances shown in [Table 9-1](#).

Table 9-1 BGP Default Administrative Distances

Distance	Default Value	Function
External	20	Applied to routes learned from eBGP.
Internal	200	Applied to routes learned from iBGP.
Local	200	Applied to routes originated by the router.

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

**Note**

The administrative distance does not influence the BGP path selection algorithm, but it does influence whether BGP-learned routes are installed in the IP routing table.

For more information, see the [“Administrative Distance” section on page 1-6](#).

BGP Peers

A BGP speaker does not discover another BGP speaker automatically. You must configure the relationships between BGP speakers. A *BGP peer* is a BGP speaker that has an active TCP connection to another BGP speaker.

BGP uses TCP port 179 to create a TCP session with a peer. When a TCP connection is established between peers, each BGP peer initially exchanges all of its routes—the complete BGP routing table—with the other peer. After this initial exchange, the BGP peers send only incremental updates when a topology change occurs in the network or when a routing policy change occurs. In the periods of inactivity between these updates, peers exchange special messages called *keepalives*. The *hold time* is the maximum time limit that can elapse between receiving consecutive BGP update or keepalive messages.

BGP Router Identifier

To establish BGP sessions between peers, BGP must have a *router ID*. The router ID is sent to BGP peers in the OPEN message when a BGP session is established. The BGP router ID is a 32-bit value that is often represented by an IPv4 address. You can configure the router ID. By default, Cisco NX-OS sets the router ID to the IPv4 address of a loopback interface on the router. If no loopback interface is configured on the router, then the software chooses the highest IPv4 address configured to a physical interface on the router to represent the BGP router ID. The BGP router ID must be unique to the BGP peers in a network.

If BGP does not have a router ID, it cannot establish any peering sessions with BGP peers.

BGP Path Selection

BGP might receive advertisements for the same route from multiple sources. BGP selects only one path as the best path. When the path is selected, BGP puts the selected path in the IP routing table and propagates the path to its peers.

The best-path algorithm runs each time that a path is added or withdrawn for a given network. The best-path algorithm also runs if you change the BGP configuration. BGP selects the best path from the set of valid paths available for a given network.

Cisco NX-OS implements the BGP best-path algorithm in three parts:

- Part 1—Compares two paths to determine which is better (see the [“Comparing Pairs of Paths” section on page 9-4](#)).
- Part 2—Iterates over all paths and determines in which order to compare the paths to select the overall best path (see the [“Order of Comparisons” section on page 9-5](#)).
- Part 3—Determines whether the old and new best paths differ enough that the new best path should be used (see the [“Best-Path Change Suppression” section on page 9-6](#)).

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

**Note**

The order of comparison determined in Part 2 is important. Consider the case where you have three paths, A, B, and C. When Cisco NX-OS compares A and B, it chooses A. When Cisco NX-OS compares B and C, it chooses B. But when Cisco NX-OS compares A and C, it might not choose A because some BGP metrics apply only among paths from the same neighboring autonomous system and not among all paths.

The path selection uses the the BGP AS-path attribute. The AS-path attribute includes the list of autonomous system numbers (AS numbers) traversed in the advertised path. If you subdivide your BGP autonomous system to be a collection or confederation of autonomous systems, the AS-path contains confederation segments that list these locally defined autonomous systems.

Comparing Pairs of Paths

The following sequence describes the basic steps that Cisco NX-OS uses to compare two paths to determine the better path:

1. Cisco NX-OS chooses a valid path for comparison. (For example, a path that has an unreachable next-hop is not valid.)
2. Cisco NX-OS chooses the path with the highest weight.
3. Cisco NX-OS chooses the path with the highest local preference.
4. If one of the paths is locally originated, Cisco NX-OS chooses that path.
5. Cisco NX-OS chooses the path with the shorter AS-path.

**Note**

When calculating the length of the AS-path, Cisco NX-OS ignores confederation segments, and counts AS sets as 1. See the [“AS Confederations” section on page 10-4](#) for more information.

6. Cisco NX-OS chooses the path with the lower origin. Interior Gateway Protocol (IGP) is considered lower than EGP.
7. Cisco NX-OS chooses the path with the lower multi exit discriminator (MED).

You can configure a number of options that affect whether or not this step is performed. In general, Cisco NX-OS compares the MED if both paths were received from peers in the same autonomous system; otherwise, Cisco NX-OS skips the MED comparison.

You can configure Cisco NX-OS to always perform the best-path algorithm MED comparison, regardless of the peer autonomous system in the paths. See the [“Tuning the Best-Path Algorithm” section on page 10-8](#) for more information. Otherwise, the MED comparison depends on the AS-path attributes of the two paths being compared, as follows:

- a. If a path has no AS-path or the AS-path starts with an AS_SET, then the path is internal, and Cisco NX-OS compares the MED to other internal paths.
- b. If the AS-path starts with an AS_SEQUENCE, then the peer autonomous system is the first AS number in the sequence, and Cisco NX-OS compares the MED to other paths that have the same peer autonomous system.
- c. If the AS_Path contains only confederation segments or starts with confederation segments followed by an AS_SET, the path is internal and Cisco NX-OS compares the MED to other internal paths.

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

- d. If the AS-path starts with confederation segments followed by an AS_SEQUENCE, then the peer autonomous system is the first AS number in the AS_SEQUENCE, and Cisco NX-OS compares the MED to other paths that have the same peer autonomous system.



Note If Cisco NX-OS receives no MED attribute with the path, then Cisco NX-OS considers the MED to be 0 unless you configure the best-path algorithm to set a missing MED to the highest possible value. See the [“Tuning the Best-Path Algorithm”](#) section on page 10-8 for more information.

- e. If the nondeterministic MED comparison feature is enabled, the best path algorithm uses the Cisco IOS style of MED comparison. See the [“Tuning the Best-Path Algorithm”](#) section on page 10-8 for more information.
8. If one path is from an internal peer and the other path is from an external peer, then Cisco NX-OS chooses the path from the external peer.
9. If the paths have different IGP metrics to their next-hop addresses, then Cisco NX-OS chooses the path with the lower IGP metric.
10. Cisco NX-OS uses the path that was selected by the best-path algorithm the last time that it was run.

If all path parameters in Step 1 through Step 9 are the same, then you can configure the best path algorithm to compare the router IDs. See the [“Tuning the Best-Path Algorithm”](#) section on page 10-8 for more information. If the path includes an originator attribute, then Cisco NX-OS uses that attribute as the router ID to compare to; otherwise, Cisco NX-OS uses the router ID of the peer that sent the path. If the paths have different router IDs, Cisco NX-OS chooses the path with the lower router ID.



Note When using the originator as the router ID, it is possible that two paths have the same router ID. It is also possible to have two BGP sessions with the same peer router, and therefore you can receive two paths with the same router ID.

11. Cisco NX-OS selects the path with the shorter cluster length is selected. If a path was not received with a cluster list attribute, the cluster length is 0.
12. Cisco NX-OS chooses the path received from the peer with the lower IP address. Locally generated paths (for example, redistributed paths) have a peer IP address of 0.



Note Paths that are equal after step 9 can be used for multipath if you configure multipath. See the [“Load Sharing and Multipath”](#) section on page 10-6 for more information.

Order of Comparisons

The second part of the BGP best-path algorithm implementation determines the order in which the paths should be compared. Cisco NX-OS determines the order of comparison as follows:

1. Cisco NX-OS partitions the paths into groups. Within each group Cisco NX-OS compares the MED among all paths. Cisco NX-OS the same rules as in the [“Comparing Pairs of Paths”](#) section on page 9-4 to determine whether MED can be compared between any two paths. Typically, this comparison results in one group for each neighbor autonomous system. If you configure the **bgp bestpath med always** command, then there is just one group that contains all the paths.

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

2. Cisco NX-OS determines the best path in each group by iterating through all paths in the group and keeping track of the best one so far. Cisco NX-OS compares each path with the temporary best path found so far and if the new path is better, it becomes the new temporary best path and Cisco NX-OS compares it with the next path in the group.
3. Cisco NX-OS forms a set of paths that contain the best path selected from each group in Step 2. Cisco NX-OS selects the overall best path from this set of paths by going through them as in Step 2.

Best-Path Change Suppression

The next part of the implementation is to determine whether to use the new best path. The router can continue to use the existing best path if the new one is identical to the point at which the best-path selection algorithm becomes arbitrary (if the router ID is the same). Continuing to use the existing best path can avoid route changes in the network.

To turn off the suppression, configure the best path algorithm to compare the router IDs. See the [“Tuning the Best-Path Algorithm” section on page 10-8](#) for more information. If you configure this feature, the new best path is always preferred to the existing one.

Otherwise, the best-path change cannot be suppressed if any of the following conditions occur:

- The existing best path is no longer valid.
- Either the existing or new best paths were received from internal (or confederation) peers or were locally generated (for example, by redistribution).
- The paths were received from the same peer (the paths have the same router ID).
- The paths have different weights, local preferences, origins, or IGP metrics to their next-hop addresses.
- The paths have different MEDs.

BGP and the Unicast RIB

BGP communicates with the unicast routing information base (unicast RIB) to store IPv4 routes in the unicast routing table. After selecting the best path, if BGP determines that the best path change needs to be reflected in the routing table, it sends a route update to the unicast RIB.

BGP receives route notifications regarding changes to its routes in the unicast RIB. It also receives route notifications about other protocol routes to support redistribution.

BGP also receives notifications from the unicast RIB regarding next-hop changes. BGP uses these notifications to keep track of the reachability and IGP metric to the next-hop addresses.

Whenever the next-hop reachability or IGP metrics in the unicast RIB change, BGP triggers a best-path recalculation for affected routes.

BGP communicates with the IPv6 unicast RIB to perform these operations for IPv6 routes.

BGP Virtualization

BGP supports Virtual Routing and Forwarding instances (VRFs). VRFs exist within virtual device contexts (VDCs). By default, Cisco NX-OS places you in the default VDC and default VRF unless you specifically configure another VDC and VRF. For more information, see the *CCisco 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.

Licensing Requirements for Basic 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 configure at least one IGP that is capable of recursive next-hop resolution.
- You must configure an address family under a neighbor for the BGP session establishment.

Guidelines and Limitations for BGP

Use the following guidelines and limitations to configure BGP:

- 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 BGP/EBGP multihop sessions.
- Specify a BGP policy if you configure redistribution.
- Define the BGP router ID within a VRF.
- If you decrease the keepalive and hold timer values, you might experience BGP 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*).
- If you configure VRFs, install the Advanced Services license and enter the desired VRF (see [Chapter 14, “Configuring Layer 3 Virtualization”](#)).

CLI Configuration Modes

The following sections show how to enter each of the CLI configuration modes for BGP. From a mode, you can enter the ? command to display the commands available in that mode.

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

Global Configuration Mode

Use global configuration mode to create a BGP process and configure advanced features such as AS confederation and route dampening. For more information, see [Chapter 10, “Configuring Advanced BGP.”](#)

The following example shows how to enter router configuration mode:

```
switch# configuration
switch(config)# router bgp 5
switch(config-router)#
```

BGP supports Virtual Routing and Forwarding (VRF). You can configure BGP within the appropriate VRF if you are using VRFs in your network. See the [“Configuring Virtualization” section on page 10-33](#) for more information.

The following example shows how to enter VRF configuration mode:

```
switch(config)# router bgp 7
switch(config-router)# vrf vrf_A
switch(config-router-vrf)#
```

Address Family Configuration Mode

You can optionally configure the address families that BGP supports. Use the **address-family** command in router configuration mode to configure features for an address family. Use the **address-family** command in neighbor configuration mode to configure the specific address family for the neighbor.

You must configure the address families if you are using route redistribution, address aggregation, load balancing, and other advanced features.

The following example shows how to enter address family configuration mode from the router configuration mode:

```
switch(config)# router bgp 5
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)#
```

The following example shows how to enter VRF address family configuration mode if you are using VRFs:

```
switch(config)# router bgp 7
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# address-family ipv4 unicast
switch(config-router-vrf-af)#
```

Neighbor Configuration Mode

Cisco NX-OS provides the neighbor configuration mode to configure BGP peers. You can use neighbor configuration mode to configure all parameters for peer.

The following example shows how to enter neighbor configuration mode:

```
switch(config)# router bgp 5
switch(config-router)# neighbor 192.0.2.1
switch(config-router-neighbor)#
```

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

The following example shows how to enter VRF neighbor configuration mode:

```
switch(config)# router bgp 7
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# neighbor 192.0.2.1
switch(config-router-vrf-neighbor)#
```

Neighbor Address Family Configuration Mode

An address family configuration submode inside the neighbor configuration submode is available for entering address family-specific neighbor configuration and enabling the address family for the neighbor. Use this mode for advanced features such as limiting the number of prefixes allowed for this neighbor and removing private AS numbers for EBGp.

The following example shows how to enter neighbor address family configuration mode:

```
switch(config)# router bgp 5
switch(config-router)# neighbor 192.0.2.1
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)#
```

The following example shows how to enter VRF neighbor address family configuration mode:

```
switch(config)# router bgp 7
switch(config-router)# vrf vrf_A
switch(config-router-vrf)# neighbor 209.165.201.1
switch(config-router-vrf-neighbor)# address-family ipv4 unicast
switch(config-router-vrf-neighbor-af)#
```

Configuring Basic BGP

To configure a basic BGP, you need to enable BGP and configure a BGP peer. Configuring a basic BGP network consists of a few required tasks and many optional tasks. You must configure a BGP routing process and BGP peers.

This section includes the following topics:

- [Enabling the BGP Feature, page 9-9](#)
- [Creating a BGP Instance, page 9-10](#)
- [Restarting a BGP Instance, page 9-12](#)
- [Configuring BGP Peers, page 9-12](#)
- [Clearing BGP Information, page 9-15](#)



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.

Enabling the BGP Feature

You must enable the BGP feature before you can configure BGP.

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

BEFORE YOU BEGIN

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

SUMMARY STEPS

1. **config t**
2. **feature bgp**
3. **copy running-config startup-config**

DETAILED STEPS

	Command	Purpose
Step 1	config t Example: switch# config t switch(config)#	Enters configuration mode.
Step 2	feature bgp Example: switch(config)# feature bgp	Enables the BGP feature.
Step 3	copy running-config startup-config Example: switch(config)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **no feature bgp** command to disable the BGP feature and remove all associated configuration.

	Command	Purpose
	no feature bgp Example: switch(config)# no feature bgp	Disables the BGP feature and removes all associated configuration.

Creating a BGP Instance

You can create a BGP instance and assign a router ID to the BGP instance. See the [“BGP Router Identifier” section on page 9-3](#). Cisco NX-OS supports 2-byte or 4-byte autonomous system numbers (AS numbers).

BEFORE YOU BEGIN

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

BGP must be able to obtain a router ID (for example, a configured loopback address).

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

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

SUMMARY STEPS

1. **config t**
2. **router bgp** *autonomous-system-number*
3. **router-id** *ip-address*
4. **address-family** {*ipv4* | *ipv6*} {*unicast* | *multicast*}
5. **network** *ip-prefix* [**route-map** *map-name*]
6. **show bgp all**
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	router bgp <i>autonomous-system-number</i> Example: switch(config)# router bgp 40000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker. The autonomous system number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number.lower 16-bit decimal number.
Step 3	router-id <i>ip-address</i> Example: switch(config-router)# router-id 192.0.2.255	(Optional) Configures the BGP router ID. This IP address identifies this BGP speaker. This command triggers an automatic notification and session reset for the BGP neighbor sessions.
Step 4	address-family { <i>ipv4</i> <i>ipv6</i> } { <i>unicast</i> <i>multicast</i> } Example: switch(config-router)# address-family ipv4 unicast switch(config-router-af)#	(Optional) Enters global address family configuration mode for the IPv4 address family. This command triggers an automatic notification and session reset for all BGP neighbors.
Step 5	network <i>ip-prefix</i> [route-map <i>map-name</i>] Example: switch(config-router-af)# network 192.0.2.0	(Optional) Specifies a network as local to this autonomous system and adds it to the BGP routing table. For exterior protocols, the network command controls which networks are advertised. Interior protocols use the network command to determine where to send updates.

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

	Command	Purpose
Step 6	show bgp all Example: switch(config-router-af)# show bgp all	(Optional) Displays information about all BGP address families.
Step 7	copy running-config startup-config Example: switch(config-router-af)# copy running-config startup-config	(Optional) Saves this configuration change.

Use the **no router bgp** command to remove the BGP process and the associated configuration.

Command	Purpose
no router bgp <i>autonomous-system-number</i> Example: switch(config)# no router bgp 201	Deletes the BGP process and the associated configuration.

The following example shows how to enable BGP with the IPv4 unicast address family and manually add one network to advertise:

```
switch# config t
switch(config)# router bgp 40000
switch(config-router)# address-family ipv4 unicast
switch(config-router-af)# network 192.0.2.0
switch(config-router-af)# copy running-config startup-config
```

Restarting a BGP Instance

You can restart a BGP instance. This clears all peer sessions for the instance.

To restart a BGP instance and remove all associated peers, use the following command:

Command	Purpose
restart bgp <i>instance-tag</i> Example: switch(config)# restart bgp 201	Restarts the BGP instance and resets or reestablishes all peering sessions.

Configuring BGP Peers

You can configure a BGP peer within a BGP process. Each BGP peer has an associated keepalive timer and hold timers. You can set these timers either globally or for each BGP peer. A peer configuration overrides a global configuration.

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.

SUMMARY STEPS

1. **config t**
2. **router bgp** *autonomous-system-number*
3. **neighbor** *ip-address remote-as as-number*
4. **description** *text*
5. **timers** *keepalive-time hold-time*
6. **shutdown**
7. **address-family** {*ipv4 | ipv6*} {*unicast | multicast*}
8. **show bgp** {*ipv4 | ipv6*} {*unicast | multicast*} **neighbors**
9. **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 40000 switch(config-router)#	Enables BGP and assigns the autonomous system number to the local BGP speaker.
Step 3	neighbor <i>ip-address remote-as as-number</i> Example: switch(config-router)# neighbor 209.165.201.1 remote-as 45000 switch(config-router-neighbor)#	Configures the IP address and AS number for a remote BGP peer.
Step 4	description <i>text</i> Example: switch(config-router-neighbor)# description Peer Router B switch(config-router-neighbor)#	(Optional) Adds a description for the neighbor. The description is an alphanumeric string up to 80 characters long.
Step 5	timers <i>keepalive-time hold-time</i> Example: switch(config-router-neighbor)# timers 30 90	(Optional) Adds the keepalive and hold time BGP timer values for the neighbor. The range is from 0 to 3600 seconds. The default is 60 seconds for the keepalive time and 180 seconds for the hold time.
Step 6	shutdown Example: switch(config-router-neighbor)# shutdown	(Optional). Administratively shuts down this BGP neighbor. This command triggers an automatic notification and session reset for the BGP neighbor sessions.

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

	Command	Purpose
Step 7	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 8	show bgp {ipv4 ipv6}{unicast multicast} neighbors Example: switch(config-router-neighbor-af)# show bgp ipv4 unicast neighbors	(Optional) Displays information about BGP peers.
Step 9	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 a BGP peer:

```
switch# config t
switch(config)# router bgp 40000
switch(config-router)# neighbor 192.0.2.1 remote-as 45000
switch(config-router-neighbor)# description Peer Router B
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# copy running-config startup-config
```

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

Clearing BGP Information

To clear BGP information, use the following commands:

Command	Purpose
clear bgp all { <i>neighbor</i> * <i>as-number</i> peer-template <i>name</i> <i>prefix</i> } [vrf <i>vrf-name</i>]	Clears one or more neighbors from all address families. * clears all neighbors in all address families. The arguments are as follows: <ul style="list-style-type: none"> <i>neighbor</i>—IPv4 or IPv6 address of neighbor. <i>as-number</i>—The autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number.lower 16-bit decimal number. <i>name</i>—Peer template name. The name can be any case-sensitive alphanumeric string up to 64 characters. <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.
clear bgp all dampening [vrf <i>vrf-name</i>]	Clears route flap dampening networks in all address families. The <i>vrf-name</i> can be any case-sensitive alphanumeric string up to 64 characters.
clear bgp all flap-statistics [vrf <i>vrf-name</i>]	Clears route flap statistics in all address families. The <i>vrf-name</i> can be any case-sensitive alphanumeric string up to 64 characters.
clear bgp {ip ipv6} {unicast multicast} dampening [vrf <i>vrf-name</i>]	Clears route flap dampening networks in the selected address family. The <i>vrf-name</i> can be any case-sensitive alphanumeric string up to 64 characters.
clear bgp {ip ipv6} {unicast multicast} flap-statistics [vrf <i>vrf-name</i>]	Clears route flap statistics in the selected address family. The <i>vrf-name</i> can be any case-sensitive alphanumeric string up to 64 characters.

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

Command	Purpose
clear bgp { ip ipv6 } { unicast multicast } { <i>neighbor</i> * <i>as-number</i> peer-template name <i>prefix</i> } [vrf <i>vrf-name</i>]	Clears one or more neighbors from the selected address family. * clears all neighbors in the address family. The arguments are as follows: <ul style="list-style-type: none"> <i>neighbor</i>—IPv4 or IPv6 address of neighbor. <i>as-number</i>—The autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number.lower 16-bit decimal number. <i>name</i>—Peer template name. The name can be any case-sensitive alphanumeric string up to 64 characters. <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.
clear ip bgp { ip { unicast multicast }} { <i>neighbor</i> * <i>as-number</i> peer-template name <i>prefix</i> } [vrf <i>vrf-name</i>]	Clears one or more neighbors. * clears all neighbors in the address family. The arguments are as follows: <ul style="list-style-type: none"> <i>neighbor</i>—IPv4 or IPv6 address of neighbor. <i>as-number</i>—The autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number.lower 16-bit decimal number. <i>name</i>—Peer template name. The name can be any case-sensitive alphanumeric string up to 64 characters. <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.
clear ip bgp dampening [<i>ip-neighbor</i> <i>ip-prefix</i>] [vrf <i>vrf-name</i>]	Clears route flap dampening in one or more networks. The arguments are as follows: <ul style="list-style-type: none"> <i>ip-neighbor</i>—IPv4 address of neighbor. <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.

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

Command	Purpose
clear ip bgp flap-statistics [<i>ip-neighbor</i> <i>ip-prefix</i>] [vrf <i>vrf-name</i>]	Clears route flap statistics in one or more networks. The arguments are as follows: <ul style="list-style-type: none"> • <i>ip-neighbor</i>—IPv4 address of neighbor. • <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared. • <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.
clear ip mbgp { ip { unicast multicast }} { <i>neighbor</i> * <i>as-number</i> peer-template <i>name</i> <i>prefix</i> } [vrf <i>vrf-name</i>]	Clears one or more neighbors. * clears all neighbors in the address family. The arguments are as follows: <ul style="list-style-type: none"> • <i>neighbor</i>—IPv4 or IPv6 address of neighbor. • <i>as-number</i>—The autonomous system number. The AS number can be a 16-bit integer or a 32-bit integer in the form of higher 16-bit decimal number.lower 16-bit decimal number. • <i>name</i>—Peer template name. The name can be any case-sensitive alphanumeric string up to 64 characters. • <i>prefix</i>—IPv4 or IPv6 prefix. All neighbors within that prefix are cleared. • <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.

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

Command	Purpose
clear ip mbgp dampening [<i>ip-neighbor</i> <i>ip-prefix</i>] [vrf <i>vrf-name</i>]	Clears route flap dampening one or more networks. The arguments are as follows: <ul style="list-style-type: none"> <i>ip-neighbor</i>—IPv4 address of neighbor. <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.
clear ip mbgp flap-statistics [<i>ip-neighbor</i> <i>ip-prefix</i>] [vrf <i>vrf-name</i>]	Clears route flap statistics one or more networks. The arguments are as follows: <ul style="list-style-type: none"> <i>ip-neighbor</i>—IPv4 address of neighbor. <i>ip-prefix</i>—IPv4. All neighbors within that prefix are cleared. <i>vrf-name</i>—VRF name. All neighbors in that VRF are cleared. The name can be any case-sensitive alphanumeric string up to 64 characters.

Verifying Basic BGP Configuration

To verify the BGP configuration, use the following commands:

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

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>] filter-list list-name	Displays the information for the 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 next-hop-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 list-name	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 the 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 map-name	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.

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

Basic BGP Example Configuration

The following example shows a basic BGP configuration:

```
feature bgp
router bgp 40000
  address-family ipv4 multicast
    network 192.0.2.0/24
    network 209.165.201.0/27
  address-family ipv4 unicast
    network 192.0.2.0/24
    network 209.165.201.0/27
  address-family ipv6 multicast
    network 2001::0DB8::/64
    network 2001::0DB8:0:1::/64
  address-family ipv6 unicast
    network 2001:0DB8::/64
    network 2001:0DB8:0:1::/64
  neighbor 2001:0DB8:0:1::55 remote-as 30
    address-family ipv6 multicast
    address-family ipv6 unicast
  neighbor 209.165.201.1 remote-as 45000
    address-family ipv4 multicast
    address-family ipv4 unicast
```

Related Topics

The following topics relate to BGP:

- [Chapter 15, “Configuring Route Policy Manager.”](#)

Where to Go Next

See [Chapter 10, “Configuring Advanced BGP”](#) for details on the following features:

- Peer templates
- Route redistribution
- Route maps

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

Default Settings

Table 9-2 lists the default settings for BGP parameters.

Table 9-2 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 9-21](#)
- [MIBs, page 9-21](#)

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

Feature History for BGP

Table 9-3 lists the release history for this feature.

Table 9-3 Feature History for BGP

Feature Name	Releases	Feature Information
Clearing BGP	4.0(3)	Added support for IPv6 neighbors in the clear [ip] {bgp mbgp} commands.
BGP	4.0(1)	This feature was introduced.

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