

# Chapter 44: BGP



# Table of Contents

<b>Chapter 44 BGP</b> .....	2
44.1 BGP Overview .....	2
44.2 Configure BGP .....	3
44.2.1 BGP Configuration List .....	3
44.2.2 Enable BGP .....	3
44.2.3 Configure BGP Peers .....	4
44.2.4 Configure BGP Parameters .....	6
44.2.5 Monitoring and Maintaining BGP .....	7

# Chapter 44 BGP

## 44.1 BGP Overview

Border Gateway Protocol (BGP) is a dynamic routing protocol deployed between autonomous systems (ASs). It automatically exchanges loop-free routing information between ASs and builds up the topological structure of ASs through exchange of network reachability information with the AS Path attribute.

BGP normative references include RFC1105 (BGP-1), RFC1163 (BGP-2), RFC1267 (BGP-3), RFC1771 (BGP-4), and RFC4271 (BGP-4). RFC1771 has seen the widest application and RFC4271 is the latest issue. BGP is suitable for a distributed network and supports Classless Interdomain Routing (CIDR). With BGP, users can customize policies. BGP-4 is becoming a matter-of-factor standard for Internet exterior routing protocols. BGP is usually deployed between ISPs.

BGP has the following features:

Interior routing protocols such as OSPF and RIP are designed to discover and calculate routes. As an exterior routing protocol, BGP focuses on control of route distribution and selection of the best route.

The AS Path attribute is added to BGP routes to eliminate the routing loop problem.

With TCP as the transport layer protocol, BGP presents better protocol reliability.

Support for CIDR is a significant characteristic of BGP-4 compared with BGP-3. The CIDR technology does not categorized IP addresses into class A, class B, and class C IP addresses. For example, 192.168.0.0 (255.255.0.0) is naturally an invalid class C IP address. This IP address, however, is expressed as 192.168.0.0/16 in CIDR and becomes a valid network address. /16 indicates that the subnet mask is composed of the first 16 bits counted from the left of the IP address. CIDR also simplifies route aggregation, which is a process of consolidating several different routes. With the route aggregation technology, multiple routes are advertised as one route, which reduces the overhead of BGP tables and network bandwidth usage.

In the case of route updates, BGP transmits only incremental routes and substantially reduces the bandwidth used by BGP route transmission. Therefore, BGP is appropriate when a large number of routes need to be transmitted on Internet.

In consideration of management and security, each AS expects to control its incoming and outgoing routes. BGP-4 provides abundant routing policies for flexible route filtering and selection. In addition, BGP-4 is easy to expand and conducive to network development.

BGP runs on a specific router as an upper-layer protocol. Upon startup of BGP, the BGP router sends the entire BGP table to its peer for routing information exchange and then only Update messages are exchanged between them for processing of changed routes. BGP detects the connection between routers by sending and receiving Keepalive messages.

The router sending a BGP message is called the BGP speaker, which constantly receives or generates new routing information and advertises it to other BGP speakers. After receiving a new route advertisement from another AS, the BGP speaker distributes the route advertisement to all the other BGP speakers in the same AS if the route is better than the current one or has not been received ever. If two BGP speakers are exchanging messages, they call each other the peer.

BGP runs on a router in either of the following modes:

**Internal BGP (IBGP)**

**External BGP (EBGP)**

BGP is regarded as IBGP when deployed within an AS and as EBGP when deployed between ASs.

BGP running is driven by messages, which are classified as follows:

**Open message**

**Update message**

**Notification message**

### Keepalive message

An Open message is the first message to be sent after setup of a TCP connection and used to establish a BGP peer relationship. A Notification message is sent when there is an error. A Keepalive message is sent to detect the validity of a connection. As the most important message in BGP, an Update message is transmitted between BGP peers for routing information exchange. It consists of three parts at most: unreachable route, path attributes, and Network Layer Reachability Information (NLRI).

## 44.2 Configure BGP

### 44.2.1 BGP Configuration List

Configuration Task	Description	Detailed Configuration
Enable BGP	Required	44.2.2
Configure BGP peers	Required	44.2.3
Configure BGP Parameters	Required	44.2.4
Monitoring and Maintain BGP	Required	44.2.5

### 44.2.2 Enable BGP

Operation	Command	Remarks
Enter the global configuration mode	<b>system-view</b>	
Runs the command in global configuration mode.	<b>router bgp</b> <i>as-number</i>	
Runs the command in global configuration mode.	<b>undo router bgp</b> <i>as-number</i>	
Configures the local route to be advertised by BGP.	<b>network</b> <i>ip-address</i> [ <b>mask</b> <i>address-mask</i> ]	
Cancels the local route to be advertised by BGP.	<b>undo network</b> <i>ip-address</i> [ <b>mask</b> <i>address-mask</i> ]	
Establishes a neighbor relationship and sets the AS number of the peer.	<b>neighbor</b> <i>neighbor-name</i> <b>peer-group</b>	
Cancels neighbor relationship	<b>undo neighbor</b> <i>neighbor-name</i> <b>peer-group</b>	

## 44.2.3 Configure BGP Peers

Operation	Command	Remarks
Enter the global configuration mode	<b>system-view</b>	
Runs the command in global configuration mode.	<b>router bgp</b> <i>as-number</i>	
Establishes a neighbor relationship and sets the AS number of the peer.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>remote-as</b> <i>as-number</i>	
Deletes the established neighbor relationship.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>remote-as</b>	
Configures peer-group member	<b>neighbor</b> <i>neighbor-address</i> <b>peer-group</b> <i>neighbor-name</i>	
Delete peer-group member	<b>undo neighbor</b> <i>neighbor-address</i> <b>peer-group</b> <i>neighbor-name</i>	
Configures that a connection can be established with an EBGP peer on an indirectly-connected	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>ebgp-multihop</b> [ <i>tth</i> ]	

network.		
Configures that a connection can be established only with an EBGP peer on a directly-connected network.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>ebgp-multihop</b>	
Configures the Keepalive interval and hold timer of a BGP peer.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>timers</b> <i>keepalive-interval</i> <i>hold-time</i>	
Restores the Keepalive interval and hold timer of a BGP peer to the default values.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>timers</b>	

Configures the interval a BGP peer waits before sending a route update message.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>advertisement-interval</b> <i>seconds</i>	
Restores the interval a BGP peer waits before sending a route update message to the default value.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>advertisement-interval</b>	
Configures that its own address is used as the next hop during route advertisement.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>next-hop-self</b>	
Cancels the configuration that its	<b>undo neighbor</b> { <i>neighbor-address</i>	
own address is used as the next hop during route advertisement.	<i>neighbor-name</i> } <b>next-hop-self</b>	
Configures an IP ACL-based route filtering policy for the peer.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>distribute-list</b> { <i>ip-acl-name</i>   <i>ip-acl-number</i> } { <b>in</b>   <b>out</b> }	
Deletes an IP ACL-based route filtering policy of the peer.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>distribute-list</b> { <i>ip-acl-name</i>   <i>ip-acl-number</i> } { <b>in</b>   <b>out</b> }	
Configures an AS Path-based route filtering policy for the peer.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>filter-list</b> <i>aspath-list-number</i> { <b>in</b>   <b>out</b> }	
Deletes an AS Path-based route filtering policy for the peer.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>filter-list</b> <i>aspath-list-number</i> { <b>in</b>   <b>out</b> }	
Configures an IP-Prefix list route filtering policy for the peer.	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>prefix-list</b> <i>list-name</i> { <b>in</b>   <b>out</b> }	

Deletes an IP-Prefix list route filtering policy for the peer.	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>prefix-list</b> <i>list-name</i> { <b>in</b>   <b>out</b> }	
--	--	--

#### 44.2.4 Configure BGP Parameters

Operation	Command	Remarks
Enter the global configuration mode	<b>system-view</b>	
Runs the command in global configuration mode.	<b>router bgp</b> <i>as-number</i>	
Runs the command in BGP configuration mode.	<b>timers bgp</b> <i>keepalive-interval hold-time</i>	
Restores the default value of the timer.	<b>undo timers bgp</b>	
Disable sending connection request packet	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>passive</b>	
Enable sending connection request packet	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>passive</b>	
Shutdown the neighbor connection	<b>neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>shutdown</b>	
Open the neighbor connection	<b>undo neighbor</b> { <i>neighbor-address</i>   <i>neighbor-name</i> } <b>shutdown</b>	
Configures a local priority.	<b>bgp default local-preference</b> <i>value</i>	
Restores the default local priority.	<b>undo bgp default local-preference</b>	
Compares the MED values of neighbors from different ASs.	<b>bgp always-compare-med</b>	
Compares the MED values of neighbors from different ASs.	<b>undo bgp always-compare-med</b>	

Configures local route aggregation.	<b>aggregate-address</b> { <i>ip-address mask</i>   <i>ip-address/mask-length</i> } [ <b>summary-only</b> ] [ <b>as-set</b> ]	
Disables local route aggregation.	<b>undo aggregate-address</b> { <i>ip-address mask</i>   <i>ip-address/mask-length</i> }	
Imports IGP routes into BGP.	<b>redistribute</b> { <b>babel</b>   <b>connected</b>   <b>isis</b>   <b>kernel</b>   <b>ospf</b>   <b>rip</b>   <b>static</b> } [ <b>metric</b> <i>metric</i> ] [ <b>route-map</b> <i>route-map</i> ]]	
Cancels the import of IGP routes into BGP.	<b>undo redistribute</b> { <b>babel</b>   <b>connected</b>   <b>isis</b>   <b>kernel</b>   <b>ospf</b>   <b>rip</b>   <b>static</b> }	

#### 44.2.5 Monitoring and Maintaining BGP

Operation	Command	Remarks
Displays the detailed information of BGP peers.	<b>display ip bgp neighbors</b> <i>neighbor-address</i> [ <b>vpn-instance</b> <i>instance</i> ]	
Displays the brief information of BGP peers.	<b>display ip bgp summary</b> [ <b>vpn-instance</b> <i>instance</i> ]	