BGP Configuration

Contents

Chapter 1 BGP Configuration

1.1 BGP Description

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 (2555.255.0.0) 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 BPG 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).

1.2 BGP Configuration

1.2.1 Configuring BGP

Table 1-1 BGP configuration tasks

Configuration Task		Notes	Details
Basic BGP configuration	Enables/Disables BGP.	Mandatory	1.2.2
	Specifies the route to be advertised by BGP.	Mandatory	1.2.2
Configuring BGP peers	Configures BGP peers.	Mandatory	1.2.3
	Configures a BGP timer.	Mandatory	1.2.3
Configuring BGP parameters	Configures a local priority.	Mandatory	1.2.4
	Configures the MED value of an AS.	Optional	1.2.4
	Compares the MED values of neighbors from different ASs.	Optional	1.2.4
	Configures BGP route aggregation.	Optional	1.2.4
	Imports IGP routes into BGP.	Optional	1.2.4
	Configures BGP route filtering.	Optional	1.2.4
	Defines a distribute list.	Optional	1.2.4
	Defines an AS path list.	Optional	1.2.4
Monitoring and maintaining BGP	Displays the information in a BGP table.	Mandatory	1.2.5

1.2.2 Basic BGP Configuration

Table 1-2 Basic BGP configuration

Operation	Command	Remarks
Runs the command in global configuration mode.	**router bgp *as-number***	-
Runs the command in global configuration mode.	**no router bgp *as-number***	-
Configures the local route to be advertised by BGP.	**network ip-address [mask address-mask]**	
Cancels the local route to be advertised by BGP.	**no network ip-address [mask address-mask]**	
Establishes a neighbor relationship and sets the AS number of the peer.	**neighbor neighbor-address remote-as as-number**	

1.2.3 Configuring BGP Peers

Table 1-3 Configuring BGP peers

Operation	Command	Remarks

Operation	Command	Remarks
Establishes a neighbor relationship and sets the AS number of the peer.	**neighbor neighbor-address remote-as as-number**	-
Deletes the established neighbor relationship.	**no neighbor neighbor-address remote-as as-number**	-
Configures that a connection can be established with an EBGP peer on an indirectly-connected network.	**neighbor *neighbor-address* ebgp-multihop**	
Configures that a connection can be established only with an EBGP peer on a directly-connected network.	**no neighbor *neighbor-address* ebgp-multihop**	
Configures the Keepalive interval and hold timer of a BGP peer.	**neighbor neighbor-address timers keepalive-interval hold-time**	
Restores the Keepalive interval and hold timer of a BGP peer to the default values.	**no neighbor *neighbor-address* timers**	
Configures the interval a BGP peer waits before sending a route update message.	**neighbor neighbor-address advertisement-interval *seconds***	
Restores the interval a BGP peer waits before sending a route update message to the default value.	**no neighbor *neighbor-address* advertisement-interval**	
Configures that its own address is used as the next hop during route advertisement.	**neighbor *neighbor-address* next-hop-self**	
Cancels the configuration that its own address is used as the next hop during route advertisement.	**no neighbor *neighbor-address* next-hop-self**	
Configures an IP ACL-based route filtering policy for the peer.	**neighbor *neighbor-address* distribute-list *access-list-number* { in | out }**	
Deletes an IP ACL-based route filtering policy of the peer.	**no neighbor *neighbor-address* distribute-list *access-list-number* { in | out }**	
Configures an AS Path-based route filtering policy for the peer.	**neighbor neighbor-address filter-list aspath-list-number { in | out }**	
Deletes an AS Path-based route filtering policy for the peer.	**no neighbor *neighbor-address* filter-list *aspath-list-number* { in | out }**	

1.2.4 Configuring BGP Parameters

Table 1-4 Configuring BGP parameters

Operation	Command	Remarks
Runs the command in BGP configuration mode.	**timers bgp keepalive-interval hold-time**	-
Restores the default value of the	**no timers bgp**	By default, the K

Operation	Command	Remarks
timer.		eepalive interval is 30s and the duration of the hold timer is 180s.
Configures a local priority.	**bgp default local-preference *localpref***	
Restores the default local priority.	**no bgp default local-preference**	
Configures the MED value of an AS.	**default-metric *metric***	
Restores the default MED value of an AS.	**no default-metric**	
Compares the MED values of neighbors from different ASs.	**bgp always-compare-med**	
Compares the MED values of neighbors from different ASs.	**no bgp always-compare-med**	
Configures local route aggregation.	**aggregate-address *address mask* [summary-only]**	
Disables local route aggregation.	**no aggregate-address *address mask* [summary-only**	
Imports IGP routes into BGP.	**redistribute { connected | static | rip | ospf } [metric *metric*]**	
Cancels the import of IGP routes into BGP.	**no redistribute { connected | static | rip | ospf }**	
Defines a distribute list.	**ip distribute-list *list-number* { permit | deny } *net-addr wildcard-netmask***	
Deletes a distribute list.	**no ip distribute-list *list-number* { permit | deny } *net-addr wildcard-netmask***	

1.2.5 Defining an AS Path List

The routing information of BGP contains an AS Path field, which records the numbers of the ASs BGP routes have traversed.

Defining an AS path list is to match AS paths with a regular expression. The regular expression is a mode matching tool, allowing users to build a matching mode out of a series of special characters, compare the matching mode with the target, and determine the action to be taken on the target based on the comparison. For example, ^200$ is a regular expression to match the targets containing the string 200.

In the foregoing regular expression, "^" and "$" are characters with particular meanings. They together specify the position of a matching mode in a target and are called locators. In addition to locators, regular expressions may contain other special characters. The characters that appear before such special characters are called preamble characters, and how a preamble character appears in a target depends on special characters. Therefore, users may feel free

in customizing matching modes. The following table describes some special characters that may appear in AS regular expressions and their meanings:

Symbol	Matching Target
.	Any single character, with the space included
*	A preamble character does not appear or appears several times in a target.
- (hyphen)	Any characters beside a hyphen
_ (underline)	Comma, left curly brace, right curly brace, left bracket, and right bracket, start of an input string, end of an input string, or a blank space
[ ]	Any characters in square brackets
[^]	Any characters excluding those in square brackets (^ is located at the beginning)
|	Characters at two sides of a vertical line are in an "or" relationship.
^	The matching mode must appear at the start of a target.
$	The matching mode must appear at the end of a target.

When defining an AS path list, you can use aspath-list-number to form a regular expression to match the AS-Path field of routes, so that the routes that do not meet the predefined conditions will be filtered. Multiple AS path lists may share a list number, that is, a list number may represent a group of AS path lists. Among those, each AS path list is identified by a number.

Table 1-5 Defining an AS path list

Operation	Command	Remarks
Configures a regular expression for an AS.	**ip as-path access-list** ***aspath-list-number*** **{ permit | deny }** ***as-regular-expression***	-
Deletes a defined regular expression.	**no ip as-path access-list** ***aspath-list-number*** **{ permit | deny }** ***as-regular-expression***	-

1.2.6 Monitoring and Maintaining BGP

Table 1-6 Monitoring and maintaining BGP

Operation	Command	Remarks
Displays the information in a BGP table.	show ip bgp [ip-address | A.B.C.D/M]	-
Displays the AS path lists of BGP.	show ip as-path access-list [*aspath-list-number*]	-
Displays the detailed information of BGP peers.	show ip bgp neighbors [*neighbor-address*]	
Displays the brief information of BGP peers.	show ip bgp summary	

1.2.7 Configuration Examples

Switch(config-router-bgp)# aggregate-address 192.168.0.0 255.255.0.0

Switch(config-router-bgp)# bgp always-compare-med

```
Switch(config-router-bgp)# bgp default local-preference 200
Switch(config-router-bgp)# bgp router-id 192.168.3.4
Switch(config-router-bgp)# default-metric 10
Switch(config-router-bgp)# neighbor 192.168.3.3 advertisement-interval 10
Switch(config-router-bgp)# neighbor 192.168.3.3 distribute-list 3 in
Switch(config-router-bgp)# neighbor 192.168.3.7 ebgp-multihop
Switch(config-router-bgp)# neighbor 192.168.3.3 filter-list 3 out
Switch(config-router-bgp)# neighbor 192.168.3.3 next-hop-self
! To configure neighbor 192.168.3.7 as an EBGP neighbor, run the following commands:
Switch(config-router-bgp)# router bgp 400
Switch(config-router-bgp)# neighbor 192.168.3.7 remote-as 700
Switch(config-router-bgp)# neighbor 192.168.3.3 timers 60 180
Switch(config-router-bgp)# network 10.1.0.0 mask 255.255.0.0
Switch(config-router-bgp)# redistribute ospf
Switch(config)# router bgp 400
Switch# show ip as-path 4
ip as path access list 4, 1 rule:
0     permit ^400$
Switch# show ip bgp
Autonomous System number 400, local router ID 192.168.3.3
Status codes: s suppressed, * valid, > best, i internal
Origin codes: i - IGP, e - EGP, ? – incomplete
  Network                NextHop          Metric   LocalPref   Path
*> 192.168.5.0/24        0.0.0.0                            100    i
Switch# show ip bgp neighbors
BGP Neighbor 192.168.3.3 Status ENABLED remote AS 400, internal link
Local host 192.168.3.4 Mask 255.255.255.0 AS 400
Configured Timers: Hold 30 Keepalive 180 Connect Retry 30
Update 30 Update For IntraAS Route 15
Param: Local Preference 100 OutBound Metric 0
Route Reflector Client is DISABLED
BGP State = Established      Socket State = ESTAB Remote Initialized
Remote Router ID = 192.168.3.3 Connection Up Times 0
Running Timers: Hold 180 Keepalive 30 Connect Retry DISABLED
Update 30 Update For IntraAS Route 15
Switch# show ip bgp summary
```

Neighbor	V	AS	MsgRcvd	MsgSent	Up/Down	State/PfxRcd
192.168.3.3	4	400	1	2	04:41:13	Established
192.168.3.7	4	700	2	0	00:44:15	Established
192.168.3.8	4	400	4	1	06:27:29	Established

Switch# show ip distribute-list 3

ip distribute-list 3, 2 rule:

0 deny 192.168.9.0 0.0.0.255

1 permit 0.0.0.0 255.255.255.255

Switch(config-router-bgp)# timers bgp 30 180