

BGP: Frequently Asked Questions

Document ID: 5816

Questions

Introduction

How do I configure BGP?

How do I configure BGP with the use of a loopback address?

What is the order of preference of attributes when some or all are applied to one neighbor in BGP?

What does a next hop of 0.0.0.0 mean in the show ip bgp command output?

What are the well known communities of the BGP community attribute?

What formats can I use to configure the BGP community attribute?

How does BGP behave differently with auto-summary enabled or disabled?

How can I verify if a BGP router announces its BGP networks and propagates them to the global BGP mesh?

When and how should I reset a BGP session?

When I perform MD5 Authentication for BGP through a PIX, is there anything special that needs to be done on the PIX?

What is an autonomous system (AS) number and how do I obtain one?

What is the BGP path selection criteria?

What is the difference between always-compare-med and deterministic-med?

Do internal BGP (iBGP) sessions modify the next hop?

Do external BGP (eBGP) sessions between confederations modify the next hop?

In external BGP (eBGP) sessions, which IP address is sent as the next hop?

Does the route reflector change the next hop attribute of a reflected prefix?

How can I announce a prefix conditionally to one ISP only when I lose the connection to my primary ISP?

How can I configure BGP to provide load sharing and redundancy in my network?

How much memory should I have in my router to receive the complete BGP routing table from my ISP?

What are the benefits of configuring BGP peer groups?

What is synchronization, and how does it influence BGP routes installed in the IP routing table?

How do I know which Cisco IOS software release supports a particular BGP feature?

How can I set the Multi Exit Discriminator (MED) value on prefixes advertised to external BGP (eBGP) neighbors to match the Interior Gateway Protocol (IGP) next hop metric?

What is the default BGP ConnectRetry timer, and is it possible to tune the BGP ConnectRetry timer?

What does r RIB-Failure mean in the show ip bgp command output?

How can I redistribute internal BGP (iBGP) learned default-route (0.0.0.0/0) route into EIGRP/OSPF/IS-IS?

How can I filter all IP routes advertised to a BGP neighbor except the default route 0.0.0.0/0?

How to resolve the error Protocol not in this image?

BGP: timer-wheel running slow by 1 ticks appears in the debug output.

Is it possible to track an interface and change the route availability?

How does IP RIB Update allocate memory?

Related Information

Introduction

This document contains frequently asked questions (FAQs) about Border Gateway Protocol (BGP).

Q. How do I configure BGP?

A. Refer to these documents for information on how to configure BGP and BGP functioning:

- ◆ Configuring BGP
- ◆ BGP Case Studies

Q. How do I configure BGP with the use of a loopback address?

A. The use of a loopback interface ensures that the neighbor stays up and is not affected by malfunctioning hardware.

BGP uses the IP address configured on the physical interface directly connected to the BGP peer as the source address when it establishes the BGP peering session, by default. Issue the **neighbor <ip address> update-source <interface>** command in order to change this behavior and configure the BGP that speaks to the router to establish peering with the use of a loopback address as the source address.

Refer to Sample Configuration for iBGP and eBGP With or Without a Loopback Address for more information.

Q. What is the order of preference of attributes when some or all are applied to one neighbor in BGP?

A. The order of preference varies based on whether the attributes are applied for inbound updates or outbound updates.

For inbound updates the order of preference is:

1. route-map
2. filter-list
3. prefix-list, distribute-list

For outbound updates the order of preference is:

1. prefix-list, distribute-list
2. filter-list
3. route-map

Note: The attributes prefix-list and distribute-list are mutually exclusive, and only one command (**neighbor prefix-list** or **neighbor distribute-list**) can be applied to each inbound or outbound direction for a particular neighbor.

Q. What does a next hop of 0.0.0.0 mean in the show ip bgp command output?

A. A network in the BGP table with a next hop address of 0.0.0.0 means that the network is locally originated via redistribution of Interior Gateway Protocol (IGP) into BGP, or via a **network** or **aggregate** command in the BGP configuration.

Q. What are the well known communities of the BGP community attribute?

A. The community attribute is a transitive, optional attribute designed to group destinations in a certain community and apply certain policies (such as accept, prefer, or redistribute). This table shows the well known BGP communities.

Community	Description
Local-AS	Use in confederation scenarios to prevent sending packets outside the local autonomous system (AS).
no-export	Do not advertise to external BGP (eBGP) peers. Keep this route within an AS.
no-advertise	Do not advertise this route to any peer, internal or external.
none	Apply no community attribute when you want to clear the communities associated with a route.
internet	Advertise this route to the internet community, and any router that belongs to it.

Refer to the Configuring BGP Community Filtering section of Configuring BGP for more information about the configuration of communities.

Q. What formats can I use to configure the BGP community attribute?

A. In Cisco IOS® Software Release 12.0 and later, you can configure communities in three different formats called decimal, hexadecimal, and AA:NN. By default, IOS uses the older decimal format. In order to configure and display in AA:NN, where the first part is the AS number and the second part is a 2-byte number, issue the **ip bgp new-format** global configuration command.

Note: Although the community attribute can be represented in decimal, hexadecimal, or AA:NN, it is still a 32-bit number. For example, any of these three configuration commands specify the community 30:20 (AS 30, number 20):

- ◆ **set community 30:20**
- ◆ **set community 0x1E0014**
- ◆ **set community 1966100**

Regardless of which command you use, the community displayed in the router configuration file and the BGP table is 30:20.

Refer to the Community Attribute section of BGP Case Studies, and Using BGP Community Values to Control Routing Policy in Upstream Provider Network for more information.

Q. How does BGP behave differently with auto-summary enabled or disabled?

A. Auto-summary behavior has changed across Cisco IOS software releases. Initially, auto-summary was enabled by default. However, with Cisco bug ID CSCdu81680 (registered customers only) this behavior has changed. In the latest Cisco IOS, auto-summary is disabled by default. When auto-summary is enabled, it summarizes the locally originated BGP networks to their classfull boundaries. Auto-summary is enabled by default in BGP. When auto-summary is disabled, the routes introduced locally into the BGP table are not summarized to their classfull boundaries. When a subnet exists in the routing table and these three conditions are satisfied, then any subnet of that classfull network in the local routing table will prompt BGP to install the classfull network into the BGP table.

- ◆ Classfull network statement for a network in the routing table
- ◆ Classfull mask on that network statement
- ◆ Auto-summary enabled

For example, if the subnet in the routing table is 75.75.75.0 mask 255.255.255.0, and you configure **network 75.0.0.0** under the **router bgp** command, and auto-summary is enabled, BGP introduces the classfull network 75.0.0.0 mask 255.0.0.0 in the BGP table.

If these three conditions are not all met, then BGP does not install any entry in the BGP table unless there is an exact match in the local routing table.

Note: If the AS that performs BGP does not own the complete classfull network, Cisco recommends that you issue the **no auto-summary** command under **router bgp** in order to disable auto-summary.

Q. How can I verify if a BGP router announces its BGP networks and propagates them to the global BGP mesh?

A. Use these commands to check if the IP blocks are announced to the directly connected ISP:

- ◆ The **show ip bgp neighbors [address] advertised-routes** command shows which messages are being sent.
- ◆ The **show ip bgp neighbors [address] routes** command shows which messages are being received.

Note: The **show ip bgp neighbors [address] advertise-routes** command does not take into account any outbound policies you might have applied. In future Cisco IOS software releases, the command output will be changed to reflect the outbound policies. If there are two alternate paths to a destination, BGP always uses the best route to advertise.

In order to verify how the IP blocks get propagated to the global BGP mesh via the directly connected ISP, log onto a route server on the Internet and look for the BGP entries of the prefix in the route server.

Q. When and how should I reset a BGP session?

A. Clear a BGP session when you change the inbound/outbound policy for this session. Issue the **clear ip bgp x.x.x.x soft out** command to clear a BGP session in order to bring outbound policy changes into effect. Issue the **clear ip bgp x.x.x.x** command in order to clear a BGP session to bring inbound policy changes into effect. If the neighbor has the soft reconfiguration capability, you can issue the **clear ip bgp x.x.x.x soft in** command. The BGP

session can be cleared automatically if you setup the Optimized Edge Routing (OER). OER automatically clears the BGP session for both Inbound and Outbound directions. Refer to Setting Up OER Network Components for more information on OER.

Note: With Cisco IOS Software Release 12.0 and later, a new BGP Soft Reset Enhancement feature is introduced. Refer to BGP Soft Reset Enhancement for more information.

Q. When I perform MD5 Authentication for BGP through a PIX, is there anything special that needs to be done on the PIX?

A. Yes. When a BGP 'neighbor ... password ...' is configured, MD5 authentication is used on the TCP pseudo-IP header, TCP header, and data (refer to RFC 2385). TCP uses this data, which includes the TCP sequence and ACK numbers, and the BGP neighbor password, to create a 128-bit hash number. The hash number is included in the packet in a TCP header option field. By default, the PIX offsets the sequence number by a random value per TCP flow. On the sending BGP peer, TCP uses the original sequence number to create the 128-bit MD5 hash number and includes this hash number in the packet. When the receiving BGP peer gets the packet, TCP uses the PIX modified sequence number to create a 128-bit MD5 hash number and compares it to the hash number included in the packet. Because the TCP sequence value was changed by the PIX, the hash is different TCP on the BGP neighbor drops the packet and logs an MD5 failed message similar to this:

```
%TCP-6-BADAUTH: Invalid MD5 digest from 10.28.0.9:1778 to 10.156.50.10:179
```

Use the **norandomseq** keyword in order to solve this problem and stop the PIX from offsetting the TCP sequence number with the **static (inside,DMZ-ICE) 10.0.0.0 10.0.0.0 netmask 255.0.0.0 norandomseq** command.

Q. What is an autonomous system (AS) number and how do I obtain one?

A. AS numbers are globally unique numbers that are used to identify ASes, and which enable an AS to exchange exterior routing information between neighboring ASes. An AS is a connected group of IP networks that adhere to a single and clearly defined routing policy.

There are a limited number of available AS numbers. Therefore, it is important to determine which sites require unique AS numbers and which do not. Sites that do not require a unique AS number should use one or more of the AS numbers reserved for private use, which are in the range from 64512 to 65535. Access the AS Number Registration Services website to obtain an AS number.

Q. What is the BGP path selection criteria?

A. BGP path selection criteria is documented in BGP Best Path Selection Algorithm.

Q. What is the difference between always-compare-med and deterministic-med?

A. A complete explanation of the differences between these commands is documented in How the **bgp deterministic-med** Command Differs from the **bgp always-compare-med** Command.

Q. Do internal BGP (iBGP) sessions modify the next hop?

A. iBGP sessions preserve the next hop attribute learned from eBGP peers. This is why it is important to have an internal route to the next hop. The BGP route is otherwise unreachable. In order to make sure you can reach the eBGP next hop, include the network that the next hop belongs to in the IGP or issue the **next-hop-self neighbor** command to force the router to advertise itself, rather than the external peer, as the next hop. Refer to the BGP Next Hop Attribute section of BGP Case Studies for a more detailed explanation.

Q. Do external BGP (eBGP) sessions between confederations modify the next hop?

A. No, eBGP sessions between confederation sub-ASes do not modify the next hop attribute. All iBGP rules still apply to have the whole AS behave as a single entity. The metric and local preference values also remain unaltered among confederation eBGP peers. Refer to the BGP Confederation section of BGP Case Studies for more information about confederations.

Q. In external BGP (eBGP) sessions, which IP address is sent as the next hop?

A. In eBGP peering, the next hop is the IP address of the neighbor that announces the route. However, when the route is advertised on a multi-access media (such as Ethernet or Frame Relay), the next hop is usually the IP address of the router interface connected to that media, which originated the route. Refer to the BGP Next Hop Attribute of BGP Case Studies for a more detailed explanation.

Q. Does the route reflector change the next hop attribute of a reflected prefix?

A. By default, the next hop attribute is not changed when a prefix is reflected by route reflector. However, you can issue the **neighbor next-hop-self** command in order to change the attribute of the next hop for prefixes reflected from an eBGP peer to any route reflector client.

Q. How can I announce a prefix conditionally to one ISP only when I lose the connection to my primary ISP?

A. BGP advertises routes from its BGP table to external peers by default. The BGP conditional advertisement feature provides additional control of route advertisement depending on the existence of other prefixes in the BGP table. Normally, routes are propagated regardless of the existence of a different path. The BGP conditional advertisement feature uses the **non-exist-map** and **advertise-map** configuration commands to track routes by the route prefix. If a route prefix is not present in the **non-exist-map** command, the route specified by the **advertise-map** command is announced. Refer to the Configuring BGP Conditional Advertisement section of Configuring BGP for more information.

Q. How can I configure BGP to provide load sharing and redundancy in my network?

A. Use these documents for detailed configuration information:

- ◆ Load Sharing with BGP in Single and Multihomed Environments: Sample Configurations
- ◆ How to Use HSRP to Provide Redundancy in a Multihomed BGP Network

Q. How much memory should I have in my router to receive the complete BGP routing table from my ISP?

A. The amount of memory required to store BGP routes depends on many factors, such as the router, the number of alternate paths available, route dampening, community, the number of maximum paths configured, BGP attributes, and VPN configurations. Without knowledge of these parameters it is difficult to calculate the amount of memory required to store a certain number of BGP routes. Cisco typically recommends a minimum of 512 MB of RAM in the router to store a complete global BGP routing table from one BGP peer. However, it is important to understand ways to reduce memory consumption and achieve optimal routing without the need to receive the complete Internet routing table. Refer to *Achieve Optimal Routing and Reduce BGP Memory Consumption* (registered customers only) for more detailed information.

Q. What are the benefits of configuring BGP peer groups?

A. The major benefit of specifying a BGP peer group is that it reduces the amount of system resources (CPU and memory) used in an update generation. It also simplifies BGP configuration since it allows the routing table to be checked only once, and updates to be replicated to all other in-sync peer group members. This can significantly reduce the load, which depends on the number of peer group members, the number of prefixes in the table, and the number of prefixes advertised. Cisco recommends that you group together peers with identical outbound announcement policies. Refer to *BGP Peer Groups* for more detailed information.

Q. What is synchronization, and how does it influence BGP routes installed in the IP routing table?

A. If your AS passes traffic from another AS to a third AS, BGP should not advertise a route before all routers in your AS learn about the route via IGP. BGP waits until IGP propagates the route within the AS and then advertises it to external peers. A BGP router with synchronization enabled does not install iBGP learned routes into its routing table if it is not able to validate those routes in its IGP. Issue the **no synchronization** command under **router bgp** in order to disable synchronization. This prevents BGP from validating iBGP routes in IGP. Refer to *BGP Case Studies: Synchronization* for a more detailed explanation.

Q. How do I know which Cisco IOS software release supports a particular BGP feature?

A. Use the Cisco IOS Software Advisor (registered customers only) to quickly find which Cisco IOS software release supports your feature.

Q. How can I set the Multi Exit Discriminator (MED) value on prefixes advertised to external BGP (eBGP) neighbors to match the Interior Gateway Protocol (IGP) next hop metric?

A. The **set metric-type internal route-map** configuration command causes BGP to

advertise a MED that corresponds to the IGP metric associated with the next hop of the route. This command is available in Cisco IOS Software Release 10.3 and later. Refer to BGP Commands for more information.

Q. What is the default BGP ConnectRetry timer, and is it possible to tune the BGP ConnectRetry timer?

A. The default BGP ConnectRetry timer is 120 seconds. Only after this time passes does the BGP process check to see if the passive TCP session is established. If the passive TCP session is not established, then the BGP process starts a new active TCP attempt to connect to the remote BGP speaker. During this idle 120 seconds of the ConnectRetry timer, the remote BGP peer can establish a BGP session to it. Presently, the Cisco IOS ConnectRetry timer cannot be changed from its default of 120 seconds.

Q. What does `r RIB-Failure` mean in the `show ip bgp` command output?

```
R1> show ip bgp
BGP table version is 5, local router ID is 200.200.200.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
r> 6.6.6.0/24      10.10.13.3         0     130     0 30 i
*> 7.7.7.0/24     10.10.13.3         0     125     0 30 i
```

A. When BGP tries to install the bestpath prefix into Routing Information Base (RIB) (for example, the IP Routing table), RIB might reject the BGP route due to any of these reasons:

- ◆ Route with better administrative distance already present in IGP. For example, if a static route already exists in IP Routing table.
- ◆ Memory failure.
- ◆ The number of routes in VPN routing/forwarding (VRF) exceeds the route-limit configured under the VRF instance.

In such cases, the prefixes that are rejected for these reasons are identified by `r RIB Failure` in the `show ip bgp` command output and are not advertised to the peers. This feature was first made available in Cisco IOS Software Release 12.2(08.05)T.

Q. How can I redistribute internal BGP (iBGP) learned default-route (0.0.0.0/0) route into EIGRP/OSPF/IS-IS?

A. The redistribution of iBGP routes into Interior Gateway Protocol (IGP) Enhanced Interior Gateway Routing Protocol/Open Shortest Path First/Intermediate System-to-Intermediate System (EIGRP/OSPF/IS-IS) can cause routing loops within the Autonomous System, which is not recommended. By default, iBGP redistribution into IGP is disabled. Issue the `bgp redistribute-internal` command in order to enable redistribution of iBGP routes into IGP. Precautions must be taken to redistribute specific routes using `route-maps` into IGP. A sample configuration for redistributing a iBGP learned default route 0.0.0.0/0 into EIGRP is shown in this output. Configurations for OSPF/IS-IS are similar.

```
router bgp 65345
[...]
bgp redistribute-internal
!
```

```
router eigrp 10
[...]
redistribute bgp 65345 route-map check-def
!
ip prefix-list def-route seq 5 permit 0.0.0.0/0
!
route-map check-def permit 10
match ip address prefix-list def-route
```

Q. How can I filter all IP routes advertised to a BGP neighbor except the default route 0.0.0.0/0?

A. The specific routes can be filtered if you use inbound filter-list, distribute-list, prefix-list and route-map all at the same time for the same bgp neighbor. This is the order of operation:

1. Filter-list
2. Router-map
3. Distribute-list (or) prefix-list

Q. How to resolve the error Protocol not in this image?

A. The reason for getting the error message protocol not in this image is because BGP feature is not supported in the IOS version running on the router. To resolve this error upgrade the IOS to newer IOS versions that supports BGP.

Q. BGP: timer-wheel running slow by 1 ticks appears in the debug output.

A. This message only shows up when a BGP debug is turned on the router. It is just an informational message and not an error message. This informational message relates to BGP internal timers. This message can be ignored by issuing **undebg all** command.

Q. Is it possible to track an interface and change the route availability?

A. Yes, it is possible to track the state change of an interface and route availability with the Enhanced Object tracking. Refer to Enhanced Object Tracking for more information.

Q. How does IP RIB Update allocate memory?

A. IP RIB Update allocates the prefixes, and attributes are held in chunks. It is not possible to free the entire chunk until every element in the chunk is freed. If more routes are learned, then those free elements in the chunks are used.

Related Information

- [BGP Support Page](#)
- [Technical Support & Documentation – Cisco Systems](#)

