



Detecting and Mitigating a BGP Slow Peer

The BGP Slow Peer feature allows a network administrator to detect a BGP slow peer and also to configure a peer as a slow peer statically or to dynamically mark it.

- BGP slow peer detection identifies a BGP peer that is not transmitting update messages within a configured amount of time. It is helpful to know if there is a slow peer, which indicates there is a network issue, such as network congestion or a receiver not processing updates in time, that the network administrator can address.
- BGP slow peer configuration moves or splits the peer from its normal update group to a slow update group, thus allowing the normal update group to function without being slowed down and to converge quickly.
- [Finding Feature Information, on page 1](#)
- [Information About Detecting and Mitigating a BGP Slow Peer, on page 2](#)
- [How to Detect and Mitigate a BGP Slow Peer, on page 4](#)
- [Configuration Examples for Detecting and Mitigating a BGP Slow Peer, on page 18](#)
- [Additional References, on page 20](#)
- [Feature Information for BGP—Support for iBGP Local-AS, on page 21](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest caveats and feature information, see [Bug Search Tool](#) and the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the feature information table at the end of this module.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Information About Detecting and Mitigating a BGP Slow Peer

BGP Slow Peer Problem

BGP update generation uses the concept of update groups to optimize performance. An update group is a collection of peers with the identical outbound policy. When generating updates, the group policy is used to format messages that are then transmitted to the members of the group.

In order to maintain fairness in resource utilization, each update group is allocated a quota of formatted messages that it keeps in its cache. Messages are added to the cache when they are formatted by the group, and they are removed when they are transmitted to all the members of the group.

A slow peer is a peer that cannot keep up with the rate at which the Cisco IOS software is generating update messages, and is not keeping up over a prolonged period (in the order of a few minutes). There are several causes of a peer being slow:

- There is packet loss or high traffic on the link to the peer, and the throughput of the BGP TCP connection is very low.
- The peer has a heavy CPU load and cannot service the TCP connection at the required frequency.

When a slow peer is present in an update group, the number of formatted updates pending transmission builds up. When the cache limit is reached, the group does not have any more quotas to format new messages. In order for a new message to be formatted, some of the existing messages must be transmitted by the slow peer and then removed from the cache. The rest of the members of the group that are faster than the slow peer and have completed transmission of the formatted messages will not have anything new to send, even though there may be newly modified BGP networks waiting to be advertised or withdrawn. This effect of blocking formatting of all the peers in a group when one of the peers is slow in consuming updates is the "slow peer" problem.

Temporary Slowness Does Not Constitute a Slow Peer

Events that cause large churn in the BGP table (such as connection resets) can cause a brief spike in the rate of update generation. A peer that temporarily falls behind during such events, but quickly recovers after the event, is not considered a slow peer. In order for a peer to be marked as slow, it must be incapable of keeping up with the average rate of generated updates over a longer period (in the order of a few minutes).

BGP Slow Peer Feature

The BGP Slow Peer feature provides you, the network administrator, with three options:

- You can configure BGP slow peer detection only, which will simply detect a slow peer and provide you with information about it. Such detection is a key feature, especially in a large network of BGP peers, because you can then address the network problem that is causing the slow peer.
- You can configure a dynamic BGP slow peer. When such slow peer protection is configured, slow peer *detection* is enabled by default. The slow peer is moved or "split" from its normal update group to a slow update group, thus allowing the normal update group to function without being slowed down, and to converge more quickly than it would with the slow peer. You have the choice of whether to keep the slow peer in that slow update group until you clear the slow peer (by specifying the **permanent** keyword), or allow the slow peer to dynamically move back to its regular update group as conditions improve. We

recommend that you use the **permanent** keyword and resolve the network issue before you clear the slow peer status.

- You can configure a static BGP slow peer if you already know which peer is slow, perhaps due to a link issue or slow CPU process power. No detection is necessary, and it is more likely that the slow peer will remain there, hence the static configuration.

BGP Slow Peer Detection

You can choose to detect a BGP slow peer, whether or not you also configure the slow peer to be moved to a slow peer update group. Simply detecting a BGP slow peer provides you with useful information about the slow peer without splitting the update group. You should then address the network problem causing the slow peer.

Timestamp on an Update Message

BGP slow peer detection relies on the timestamp on the update messages in an update group. Update messages are timestamped when they are formatted. When BGP slow peer detection is configured, the timestamp of the oldest message in a peers queue is compared to the current time to determine if the peer is lagging more than the configured slow peer time threshold.

For example, if the oldest message in the peers queue was formatted more than 3 minutes ago, but the BGP slow peer detection threshold is configured at 3 minutes, then the peer that formatted that update message is determined to be a slow peer.

The Cisco IOS software generates a syslog event when a slow peer is detected or recovered (when its update group has converged and it has no messages formatted before the threshold time).

Benefit of BGP Slow Peer Detection

Slow peer detection provides you with information about the slow peer, and you can resolve the root cause without moving the peer to a different update group. Therefore, slow peer detection requires just one command that helps you identify something in your network that could be improved.

Benefits of Configuring a Dynamic or Static BGP Slow Peer

When a slow peer is present in an update group, the number of formatted updates pending transmission builds up. New messages cannot be formatted and transmitted until the backlog is reduced. That scenario delays BGP update packets and therefore delays BGP networks from being advertised. The problem can be resolved or prevented by configuring a dynamic slow peer or a static slow peer. Such configuration causes a slow peer to be put into a new, slow peer update group and thus prevents the slow peer from delaying the BGP peers that are not slow.

Static Slow Peer

If you believe that a peer is slow, you can statically configure the peer to be a slow peer. A static slow peer is recommended for a peer that is known to be slow, perhaps having a slow link or low processing power.

Static slow peer configuration causes the Cisco IOS software to create a separate update group for the peer. If you configure two peers belonging to the same update group as slow, these two peers will be moved into

a single slow peer update group because their policy will match. The slow update group will function at the pace of the slowest of the slow peers.

A static slow peer can be configured in either of two ways:

- At the BGP neighbor (address family) level
- Via a peer policy template

You probably want to determine the root cause of the peer being slow, such as network congestion or a receiver not processing updates in time. A static slow peer is not automatically restored to its original update group. You can restore a static slow peer to its original update group by using the **no neighbor slow-peer split-update-group static** command or the **no slow-peer split-update-group static** command.

Dynamic Slow Peer

An alternative to marking a static slow peer is to configure slow peers dynamically, based on the amount of time that the timestamp of the oldest message in a peers queue lags behind the current time. The default threshold is 300 seconds, and is configurable. We recommend that you specify the optional **permanent** keyword, which causes the peer to remain in the slow peer group while you resolve the root cause of the slow peer. You can then use the **clear bgp slow** command to move the peer back to its original group.

If you do not configure the **permanent** keyword, the peer moves back to its original group if and when it regains its non-slow functioning.

When a dynamic slow peer is configured, detection is enabled automatically.

You can configure dynamic slow peers in three ways:

- At the address family view level
- At the neighbor topology (that is, neighbor address-family) level
- Via a peer policy template

How to Detect and Mitigate a BGP Slow Peer

Detecting a Slow Peer

You might want to just detect a slow peer, but not move the slow peer out of its update group. Such detection notifies you by way of a syslog message that a BGP peer is not transmitting update messages within a configurable amount of time. The peer remains in its update group; the update group is not split. The syslog message level is notice level for both detection and recovery.

If you want to dynamically configure a BGP slow peer, see the [Configuring Dynamic Slow Peer Protection, on page 10](#). You will notice that that task includes and requires the step of detecting a slow peer.

Detect a slow peer by performing one of the following tasks:

Detecting Dynamic Slow Peers at the Address-Family Level

Perform this task to detect all dynamic slow peers at the address-family level. (If you want to detect *specific* slow peers, detect slow peers at the neighbor level or by using a peer policy template).

The last step is optional; use it if you want to disable slow peer detection for a specific peer.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **neighbor** {*ip-address* | *ipv6-address[%]* | *peer-group-name*} **remote-as** *autonomous-system-number*
5. **address-family ipv4**
6. **bgp slow-peer detection** [**threshold** *seconds*]
7. **neighbor** {*neighbor-address* | *peer-group-name*} **slow-peer detection disable**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|---|---|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. • Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 5 | Configures the BGP routing process. |
| Step 4 | neighbor { <i>ip-address</i> <i>ipv6-address[%]</i> <i>peer-group-name</i> } remote-as <i>autonomous-system-number</i> Example: Router(config-router)# neighbor 10.4.4.4 remote-as 5 | (Optional) Adds an entry to the BGP or multiprotocol BGP neighbor table. • This step is required if you intend to disable dynamic slow peer protection for a specific peer as shown in Step 7 below. |
| Step 5 | address-family ipv4 Example: Router(config-router)# address-family ipv4 | Enters address family configuration mode. |
| Step 6 | bgp slow-peer detection [threshold <i>seconds</i>] Example: Router(config-router-af)# bgp slow-peer detection threshold 600 | Configures global slow peer detection and specifies the time in seconds that the timestamp of the oldest update message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer. • The range of the threshold is from 120 to 3600. As long as the command is configured, the default is 300. |

| | Command or Action | Purpose |
|---------------|--|--|
| Step 7 | neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } slow-peer detection disable Example: <pre>Router(config-router-af)# neighbor 10.4.4.4 slow-peer detection disable</pre> | (Optional) Disables slow-peer detection for a specific peer. <ul style="list-style-type: none"> Use this command only if you have configured global slow peer detection in Step 5, and now you want to disable slow peer detection for a specific peer or peer group. |

Detecting Dynamic Slow Peers at the Neighbor Level

Perform this task to detect dynamic slow peers at a specific neighbor address or belonging to a specific peer group.

SUMMARY STEPS

- enable
- configure terminal
- router bgp *autonomous-system-number*
- address-family ipv4
- neighbor {*neighbor-address* | *peer-group-name*} **slow-peer detection**[*threshold seconds*]

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | enable Example: <pre>Router> enable</pre> | Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted. |
| Step 2 | configure terminal Example: <pre>Router# configure terminal</pre> | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: <pre>Router(config)# router bgp 5</pre> | Configures the BGP routing process. |
| Step 4 | address-family ipv4 Example: <pre>Router(config-router)# address-family ipv4</pre> | Enters address family configuration mode. |
| Step 5 | neighbor { <i>neighbor-address</i> <i>peer-group-name</i> } slow-peer detection [<i>threshold seconds</i>] Example: | (Optional) Specifies the time in seconds that the timestamp of the oldest message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer. |

| | Command or Action | Purpose |
|--|---|--|
| | <pre>Router(config-router-af)# neighbor 172.60.2.3 slow-peer detection threshold 1200</pre> | <ul style="list-style-type: none"> The range of the threshold is 120 seconds to 3600 seconds. As long as the command is configured, the default is 300 seconds. |

Detecting Dynamic Slow Peers Using a Peer Policy Template

Perform the following task to detect BGP slow peers using a peer policy template.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **template peer-policy** *policy-template-name*
5. **slow-peer detection** [**threshold** *seconds*]
6. **exit**
7. **address-family ipv4**
8. **neighbor** *ip-address inherit peer-policy policy-template-name*

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | <p>enable</p> <p>Example:</p> <pre>Router> enable</pre> | <p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> Enter your password if prompted. |
| Step 2 | <p>configure terminal</p> <p>Example:</p> <pre>Router# configure terminal</pre> | <p>Enters global configuration mode.</p> |
| Step 3 | <p>router bgp <i>autonomous-system-number</i></p> <p>Example:</p> <pre>Router(config)# router bgp 5</pre> | <p>Configures the BGP routing process.</p> |
| Step 4 | <p>template peer-policy <i>policy-template-name</i></p> <p>Example:</p> <pre>Router(config-router)# template peer-policy global</pre> | <p>Enters policy template configuration mode and creates a peer policy template.</p> |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 5 | slow-peer detection [<i>threshold seconds</i>] Example: <pre>Router(config-router-ptmp)# slow-peer detection threshold 600</pre> | Specifies the time in seconds that the timestamp of the oldest message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer. <ul style="list-style-type: none"> The range of the threshold is from 120 to 3600. As long as the command is configured, the default is 300. |
| Step 6 | exit Example: <pre>Router(config-router-ptmp)# exit</pre> | Exits to higher configuration mode. |
| Step 7 | address-family ipv4 Example: <pre>Router(config-router)# address-family ipv4</pre> | Enters address family configuration mode. |
| Step 8 | neighbor <i>ip-address</i> inherit peer-policy <i>policy-template-name</i> Example: <pre>Router(config-router-af)# neighbor 10.0.0.1 inherit peer-policy global</pre> | Sends a peer policy template to a neighbor so that the neighbor can inherit the configuration. |

Marking a Peer as a Static Slow Peer

There are two ways to statically configure a slow peer. Perform one of the following tasks in this section to statically configure a slow peer:

Marking a Peer as a Static Slow Peer at the Neighbor Level

Perform this task to configure a static slow peer at a specific neighbor address or belonging to a specific peer group.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp *autonomous-system-number***
4. **address-family ipv4**
5. **neighbor {*neighbor-address* | *peer-group-name*} slow-peer split-update-group static**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|-------------------|-------------------------------|
| Step 1 | enable | Enables privileged EXEC mode. |

| | Command or Action | Purpose |
|---------------|--|--|
| | Example: Router> enable | <ul style="list-style-type: none"> Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 5 | Configures the BGP routing process. |
| Step 4 | address-family ipv4 Example: Router(config-router)# address-family ipv4 | Enters address family configuration mode. |
| Step 5 | neighbor {<i>neighbor-address</i> <i>peer-group-name</i>} slow-peer split-update-group static Example: Router(config-router-af)# neighbor 172.16.1.1 slow-peer split-update-group static | Configures the neighbor at the specified address as a slow peer. <ul style="list-style-type: none"> Use the no neighbor {<i>neighbor-address</i> <i>peer-group-name</i>} slow-peer split-update-group static command if you want to restore the peer to its original, non-slow update group. |

Marking a Peer as a Static Slow Peer Using a Peer Policy Template

Perform this task to configure a static slow peer by using a peer policy template.

SUMMARY STEPS

- enable
- configure terminal
- router bgp *autonomous-system-number*
- template peer-policy *policy-template-name*
- slow-peer split-update-group static
- exit
- address-family ipv4
- neighbor *ip-address* inherit peer-policy *policy-template-name*

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---------------------------|--|
| Step 1 | enable Example: | Enables privileged EXEC mode. <ul style="list-style-type: none"> Enter your password if prompted. |

| | Command or Action | Purpose |
|---------------|--|--|
| | Router> enable | |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 5 | Configures the BGP routing process. |
| Step 4 | template peer-policy <i>policy-template-name</i> Example: Router(config-router)# template peer-policy global | Enters policy template configuration mode and creates a peer policy template. |
| Step 5 | slow-peer split-update-group static Example: Router(config-router-ptmp)# slow-peer split-update-group static | Configures the neighbor at the specified address as a slow peer. <ul style="list-style-type: none"> • Use the no slow-peer split-update-group static command if you want to restore the peer to its normal status. |
| Step 6 | exit Example: Router(config-router-ptmp)# exit | Exits to higher configuration mode. |
| Step 7 | address-family ipv4 Example: Router(config-router)# address-family ipv4 | Enters address family configuration mode. |
| Step 8 | neighbor <i>ip-address</i> inherit peer-policy <i>policy-template-name</i> Example: Router(config-router-af)# neighbor 10.0.0.1 inherit peer-policy global | Sends a peer policy template to a neighbor so that the neighbor can inherit the configuration. |

Configuring Dynamic Slow Peer Protection

There are three ways to dynamically configure slow peers, also known as slow peer protection. Perform one or more of the tasks in this section to configure dynamic slow peers:

Configuring Dynamic Slow Peers at the Address-Family Level

Configuring dynamic slow peers at the address-family level applies to all peers in the address family specified. (If you want to configure *specific* slow peers, perform this task at the neighbor level or by using a peer policy template.)

The last step is optional; perform it only if you want to disable slow peer protection for a specific peer.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **neighbor** {*ip-address* | *ipv6-address*[%] | *peer-group-name*} **remote-as** *autonomous-system-number*
5. **address-family ipv4**
6. **bgp slow-peer detection** [*threshold seconds*]
7. **bgp slow-peer split-update-group dynamic** [*permanent*]
8. **neighbor** {*neighbor-address* | *peer-group-name*} **slow-peer split-update-group dynamic disable**

DETAILED STEPS

| | Command or Action | Purpose |
|--------|--|---|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. • Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 5 | Configures the BGP routing process. |
| Step 4 | neighbor { <i>ip-address</i> <i>ipv6-address</i> [%] <i>peer-group-name</i> } remote-as <i>autonomous-system-number</i> Example: Router(config-router)# neighbor 10.4.4.4 remote-as 5 | (Optional) Adds an entry to the BGP or multiprotocol BGP neighbor table. • This step is required if you intend to disable dynamic slow peer protection for a specific peer as shown in Step 8 below. |
| Step 5 | address-family ipv4 Example: Router(config-router)# address-family ipv4 | Enters address family configuration mode. |

| | Command or Action | Purpose |
|---------------|---|---|
| Step 6 | bgp slow-peer detection [threshold seconds] Example: <pre>Router(config-router-af)# bgp slow-peer detection threshold 600</pre> | (Optional) Specifies the time in seconds that the timestamp of the oldest update message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer. <ul style="list-style-type: none"> • When a dynamic slow peer is configured, as in the next step, this detection is enabled automatically. • The range of the threshold is from 120 to 3600. The default is 300. |
| Step 7 | bgp slow-peer split-update-group dynamic [permanent] Example: <pre>Router(config-router-af)# bgp slow-peer split-update-group dynamic permanent</pre> | Moves the dynamically detected slow peer to a slow update group. <ul style="list-style-type: none"> • If a static slow peer update group exists (because of a static slow peer), the dynamic slow peer will be moved to the static slow peer update group. • If no static slow peer update group exists, a new slow peer update group will be created and the peer will be moved to that. • We recommend using the permanent keyword. If the permanent keyword is used, the peer will not be moved to its original update group automatically. After you determine the root cause of the slowness, such as network congestion, for example, you can use a clear bgp slow command to move the peer to its original update group. See the Restoring Dynamic Slow Peers as Normal Peers, on page 16 to move a dynamically slow peer back to its original update group. • If the permanent keyword is not used, the slow peer will be moved back to its regular original update group after it becomes a normal peer (converges). |
| Step 8 | neighbor {neighbor-address peer-group-name} slow-peer split-update-group dynamic disable Example: <pre>Router(config-router-af)# neighbor 10.4.4.4 slow-peer split-update-group dynamic disable</pre> | (Optional) Perform this step only if you want to disable dynamic slow peer protection for a specific peer. |

Configuring Dynamic Slow Peers at the Neighbor Level

Perform this task to configure a dynamic slow peer at a specific neighbor address or belonging to a specific peer group.

SUMMARY STEPS

1. enable

2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **address-family ipv4**
5. **neighbor** {*neighbor-address* | *peer-group-name*} **slow-peer detection** [*threshold seconds*]
6. **neighbor** {*neighbor-address* | *peer-group-name*} **slow-peer split-update-group dynamic** [*permanent*]

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | <p>enable</p> <p>Example:</p> <pre>Router> enable</pre> | <p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | <p>configure terminal</p> <p>Example:</p> <pre>Router# configure terminal</pre> | <p>Enters global configuration mode.</p> |
| Step 3 | <p>router bgp <i>autonomous-system-number</i></p> <p>Example:</p> <pre>Router(config)# router bgp 5</pre> | <p>Configures the BGP routing process.</p> |
| Step 4 | <p>address-family ipv4</p> <p>Example:</p> <pre>Router(config-router)# address-family ipv4</pre> | <p>Enters address family configuration mode.</p> |
| Step 5 | <p>neighbor {<i>neighbor-address</i> <i>peer-group-name</i>} slow-peer detection [<i>threshold seconds</i>]</p> <p>Example:</p> <pre>Router(config-router-af)# neighbor 172.60.2.3 slow-peer detection threshold 1200</pre> | <p>(Optional) Specifies the time in seconds that the timestamp of the oldest update message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer.</p> <ul style="list-style-type: none"> • When a dynamic slow peer is configured, as in the next step, this detection is enabled automatically. • The range of the threshold is from 120 to 3600. The default is 300. |
| Step 6 | <p>neighbor {<i>neighbor-address</i> <i>peer-group-name</i>} slow-peer split-update-group dynamic [<i>permanent</i>]</p> <p>Example:</p> <pre>Router(config-router-af)# neighbor 172.60.2.3 slow-peer split-update-group dynamic permanent</pre> | <p>Moves the dynamically detected slow peer to a slow update group.</p> <ul style="list-style-type: none"> • If a static slow peer update group exists (because of a static slow peer), the dynamic slow peer will be moved to the static slow peer update group. • If no static slow peer update group exists, a new slow peer update group will be created and the peer will be moved to that. |

| | Command or Action | Purpose |
|--|-------------------|--|
| | | <ul style="list-style-type: none"> • We recommend using the permanent keyword. If the permanent keyword is used, the peer will not be moved to its original update group automatically. After you resolve the root cause of the slowness, such as network congestion, for example, you can use a clear bgp slow command to move the peer to its original update group. See the Restoring Dynamic Slow Peers as Normal Peers, on page 16 to move a dynamically slow peer back to its original update group. • If the permanent keyword is not used, the slow peer will be moved back to its regular original update group after it becomes a normal peer (converges). |

Configuring Dynamic Slow Peers Using a Peer Policy Template

Perform this task to configure a BGP slow peer using a peer policy template.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp** *autonomous-system-number*
4. **template peer-policy** *policy-template-name*
5. **slow-peer detection** [*threshold seconds*]
6. **slow-peer split-update-group dynamic** [**permanent**]
7. **exit**
8. **address-family ipv4**
9. **neighbor ip-address inherit peer-policy** *policy-template-name*

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|--|--|
| Step 1 | enable Example: Router> enable | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | configure terminal Example: Router# configure terminal | Enters global configuration mode. |
| Step 3 | router bgp <i>autonomous-system-number</i> Example: Router(config)# router bgp 5 | Configures the BGP routing process. |

| | Command or Action | Purpose |
|--------|--|--|
| Step 4 | <p>template peer-policy <i>policy-template-name</i></p> <p>Example:</p> <pre>Router(config-router)# template peer-policy global</pre> | <p>Enters policy template configuration mode and creates a peer policy template.</p> |
| Step 5 | <p>slow-peer detection [threshold <i>seconds</i>]</p> <p>Example:</p> <pre>Router(config-router-ptmp)# slow-peer detection threshold 600</pre> | <p>(Optional) Specifies the time in seconds that the timestamp of the oldest message in a peers queue can be lagging behind the current time before the peer is determined to be a slow peer.</p> <ul style="list-style-type: none"> • When a dynamic slow peer is configured, as in the next step, this detection is enabled automatically. • The range of the threshold is from 120 to 3600. The default is 300. |
| Step 6 | <p>slow-peer split-update-group dynamic [permanent]</p> <p>Example:</p> <pre>Router(config-router-ptmp)# slow-peer split-update-group dynamic permanent</pre> | <p>Moves the dynamically detected slow peer to a slow update group.</p> <ul style="list-style-type: none"> • If a static slow peer update group exists (because of a static slow peer), the dynamic slow peer will be moved to the static slow peer update group. • If no static slow peer update group exists, a new slow peer update group will be created and the peer will be moved to that. • We recommend using the permanent keyword. If the permanent keyword is used, the peer will not be moved to its original update group automatically. After you determine the root cause of the slowness, such as network congestion, for example, you can use a command to move the peer to its original update group. See the Restoring Dynamic Slow Peers as Normal Peers, on page 16 to move a dynamically slow peer back to its original update group. • If the permanent keyword is not used, the slow peer will be moved back to its regular original update group after it becomes a normal peer (converges). |
| Step 7 | <p>exit</p> <p>Example:</p> <pre>Router(config-router-ptmp)# exit</pre> | <p>Exits to higher configuration mode.</p> |
| Step 8 | <p>address-family ipv4</p> <p>Example:</p> <pre>Router(config-router)# address-family ipv4</pre> | <p>Enters address family configuration mode.</p> |

| | Command or Action | Purpose |
|---------------|---|--|
| Step 9 | neighbor <i>ip-address</i> inherit peer-policy <i>policy-template-name</i> Example: <pre>Router(config-router-af)# neighbor 10.0.0.1 inherit peer-policy global</pre> | Sends a peer policy template to a neighbor so that the neighbor can inherit the configuration. |

Displaying Output About Dynamic Slow Peers

Use one or more of the **show** commands in this task to display output about dynamically configured BGP slow peers.

SUMMARY STEPS

1. **enable**
2. **show ip bgp** [**ipv4** {**multicast** | **unicast**} | **vpn4 all** | **vpn6 unicast all** | **topology**{*|*routing-topology-instance-name*}] [**update-group**] **summary slow**
3. **show ip bgp** [**ipv4** {**multicast** | **unicast**} | **vpn4 all** | **vpn6 unicast all**] **neighbors slow**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|--|
| Step 1 | enable Example: <pre>Router> enable</pre> | Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | show ip bgp [ipv4 { multicast unicast } vpn4 all vpn6 unicast all topology {* <i>routing-topology-instance-name</i> }] [update-group] summary slow Example: <pre>Router# show ip bgp summary slow</pre> | Displays information about dynamic BGP slow peers in summary form. |
| Step 3 | show ip bgp [ipv4 { multicast unicast } vpn4 all vpn6 unicast all] neighbors slow Example: <pre>Router# show ip bgp neighbors slow</pre> | Displays information about dynamic BGP slow peer neighbors. |

Restoring Dynamic Slow Peers as Normal Peers

Once you, the network administrator, resolve the root cause of a slow peer (network congestion, or a receiver not processing updates in time, and so forth), use the **clear** commands in the following task to move the peer back to its original group. Both commands perform the same function.



Note Note that *statically* configured slow peers are not affected by these **clear** commands. To restore a statically configured slow peer to its original update group, use the **no** form of the command shown in one of the tasks in the [Marking a Peer as a Static Slow Peer](#), on page 8.

SUMMARY STEPS

1. **enable**
2. **clear ip bgp** {[af] *} neighbor-address | peer-group group-name} **slow**
3. **clear bgp** af {*| neighbor-address | peer-group group-name} **slow**

DETAILED STEPS

| | Command or Action | Purpose |
|---------------|---|---|
| Step 1 | <p>enable</p> <p>Example:</p> <pre>Router> enable</pre> | <p>Enables privileged EXEC mode.</p> <ul style="list-style-type: none"> • Enter your password if prompted. |
| Step 2 | <p>clear ip bgp {[af] *} neighbor-address peer-group group-name} slow</p> <p>Example:</p> <pre>Router# clear ip bgp * slow</pre> | <p>(Optional) Restores neighbor(s) from a slow update peer group to their original update peer group.</p> <ul style="list-style-type: none"> • <i>af</i> is one of the following address families: ipv4, vpn4, or vpn6. Moves all peers in the IPv4, VPNv4 or VPNv6 address family back to their original update groups. • * moves all peers back to their original update groups. |
| Step 3 | <p>clear bgp af {* neighbor-address peer-group group-name} slow</p> <p>Example:</p> <pre>Router# clear bgp ipv4 * slow</pre> | <p>(Optional) Restores neighbor(s) from slow update peer group to their original update peer group.</p> <ul style="list-style-type: none"> • <i>af</i> is one of the following address families: ipv4, vpn4, or vpn6. Moves peers in the IPv4, VPNv4 or VPNv6 address family back to their original update groups. • * moves all peers in the address family back to their original update groups. |

Configuration Examples for Detecting and Mitigating a BGP Slow Peer

Example: Static Slow Peer

The following example marks the neighbor at 192.168.12.10 as a static slow peer.

```
router bgp 5
address-family ipv4
neighbor 192.168.12.10 slow-peer split-update-group static
```

Example: Static Slow Peer Using Peer Policy Template

The following example configures a static slow peer using a peer policy template named ipv4_ucast_pp2. The neighbor at 10.0.101.4 inherits the policy.

```
router bgp 13
template peer-policy ipv4_ucast_pp2
slow-peer split-update-group static
exit-peer-policy
!
no bgp default route-target filter
no bgp enforce-first-as
bgp log-neighbor-changes
neighbor 10.0.101.4 remote-as 13
address-family ipv4
neighbor 10.0.101.4 inherit peer-policy ipv4_ucast_pp2

RouterA# show ip bgp template peer-policy ipv4_ucast_pp2

Template:ipv4_ucast_pp2, index:2.
Local policies:0x180000000, Inherited polices:0x0
Local disable policies:0x0, Inherited disable policies:0x0
Locally configured policies:
slow-peer split-update-group static
Inherited policies:
```

Example: Dynamic Slow Peer at the Neighbor Level

The following example configures a slow peer at the neighbor level. The neighbor at 10.0.101.3 is configured with dynamic slow peer protection at a default threshold of 300 seconds.

```
router bgp 13
no bgp default route-target filter
no bgp enforce-first-as
bgp log-neighbor-changes
neighbor 10.0.101.3 remote-as 13
address-family ipv4
neighbor 10.0.101.3 slow-peer split-update-group dynamic
```

Example: Dynamic Slow Peers Using Peer Policy Template

In the following example, Router A uses a peer policy template named `ipv4_ucast_pp1` and sets a detection threshold of 120 seconds. The **permanent** keyword causes slow peers to remain in the slow update group until the network administrator uses the **clear ip bgp slow** command to move the peer to its original update group. The neighbor at 10.0.101.2 inherits the peer policy, which means that if that neighbor is determined to be slow, it is moved to a slow update group.

```
router bgp 13
  template peer-policy ipv4_ucast_pp1
  slow-peer detection threshold 120
  slow-peer split-update-group dynamic permanent
  exit-peer-policy
!
no bgp default route-target filter
no bgp enforce-first-as
bgp log-neighbor-changes
  neighbor 10.0.101.2 remote-as 13
!
address-family ipv4
  neighbor 10.0.101.2 activate
  neighbor 10.0.101.2 inherit peer-policy ipv4_ucast_pp1
```

The following output displays the locally configured policies.

```
RouterA# show ip bgp template peer-policy ipv4_ucast_pp1

Template:ipv4_ucast_pp1, index:1.
Local policies:0x300000000, Inherited polices:0x0
Local disable policies:0x0, Inherited disable policies:0x0
Locally configured policies:
  slow-peer detection threshold is 120
  slow-peer split-update-group dynamic permanent
Inherited policies:
```

Example: Dynamic Slow Peers Using Peer Group

The following example configures two peer groups: `ipv4_ucast_pg1` and `ipv4_ucast_pg2`. The neighbor at 10.0.101.1 belongs to `ipv4_ucast_pg1`, where slow peer detection is configured for 120 seconds. The neighbor at 10.0.101.5 belongs to `ipv4_ucast_pg2`, where slow peer detection is configured at 140 seconds.

```
router bgp 13
no bgp default route-target filter
no bgp enforce-first-as
bgp log-neighbor-changes
  neighbor ipv4_ucast_pg1 peer-group
  neighbor ipv4_ucast_pg2 peer-group
  neighbor ipv4_ucast_pg1 remote-as 13
  neighbor ipv4_ucast_pg2 remote-as 13
  neighbor 10.0.101.1 peer-group ipv4_ucast_pg1
  neighbor 10.0.101.5 peer-group ipv4_ucast_pg2
address-family ipv4
  neighbor ipv4_ucast_pg1 slow-peer detection threshold 120
  neighbor ipv4_ucast_pg1 slow-peer split-update-group dynamic
  neighbor ipv4_ucast_pg2 slow-peer detection threshold 140
  neighbor ipv4_ucast_pg2 slow-peer split-update-group dynamic
```

The following output displays information about the peer group `ipv4_ucast_pg1`.

```

RouterA# show ip bgp peer-group ipv4_ucast_pg1

BGP peer-group is ipv4_ucast_pg1, remote AS 13
  BGP version 4
  Neighbor sessions:
    0 active, is multiseession capable
  Default minimum time between advertisement runs is 0 seconds
  For address family: IPv4 Unicast
  BGP neighbor is ipv4_ucast_pg1, peer-group internal, members:
  10.0.101.1
  Index 0
  Slow-peer detection is enabled, threshold value is 120
  Slow-peer split-update-group dynamic is enabled
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0

```

The following output displays information about the peer group `ipv4_ucast_pg2`.

```

RouterA# show ip bgp peer-group ipv4_ucast_pg2

BGP peer-group is ipv4_ucast_pg2, remote AS 13
  BGP version 4
  Neighbor sessions:
    0 active, is multiseession capable
  Default minimum time between advertisement runs is 0 seconds
  For address family: IPv4 Unicast
  BGP neighbor is ipv4_ucast_pg2, peer-group internal, members:
  10.0.101.5
  Index 0
  Slow-peer detection is enabled, threshold value is 140
  Slow-peer split-update-group dynamic is enabled
  Update messages formatted 0, replicated 0
  Number of NLRIs in the update sent: max 0, min 0

```

Additional References

Related Documents

| Related Topic | Document Title |
|--|---|
| Cisco IOS commands | Cisco IOS Master Command List, All Releases |
| BGP commands | Cisco IOS IP Routing: BGP Command Reference |
| MPLS Layer 3 VPN configuration tasks | “Configuring MPLS Layer 3 VPNs” module in the <i>MPLS: Layer 3 VPNs Configuration Guide</i> |
| VRF selection using policy based routing | “MPLS VPN VRF Selection Using Policy-Based Routing” module in the <i>MPLS: Layer 3 VPNs Configuration Guide</i> |

Standards

| Standard | Title |
|---|-------|
| No new or modified standards are supported by this feature, and support for existing standards has not been modified by this feature. | — |

MIBs

| MIB | MIBs Link |
|-----|--|
| — | To locate and download MIBs for selected platforms, Cisco IOS releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs |

RFCs

| RFC | Title |
|---|-------|
| No new or modified RFCs are supported by this feature, and support for existing RFCs has not been modified by this feature. | — |

Technical Assistance

| Description | Link |
|---|---|
| The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password. | http://www.cisco.com/cisco/web/support/index.html |

Feature Information for BGP—Support for iBGP Local-AS

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for BGP—Support for iBGP Local-AS

| Feature Name | Releases | Feature Information |
|-------------------------------|----------|---|
| BGP—Support for iBGP Local-AS | | <p>Prior to the BGP—Support for Local-AS feature, the neighbor local-as command was used on a route reflector to customize AS_PATH attributes for routes received from an eBGP neighbor. The neighbor local-as command can now be used to enable the sending of the iBGP attributes (LOCAL_PREF, ORIGINATOR_ID, CLUSTER_ID, and CLUSTER_LIST) over an iBGP local-AS session. This functionality is useful when merging two autonomous systems, when it is advantageous to keep the iBGP attributes in routes.</p> <p>Prior to the BGP—Support for iBGP Local-AS feature, the RR should not have been configured to change iBGP attributes. With the introduction of this feature, the RR can be configured to change iBGP attributes, providing more flexibility.</p> <p>The following command was introduced:</p> <ul style="list-style-type: none"> • neighbor allow-policy <p>The following commands were modified:</p> <ul style="list-style-type: none"> • neighbor local-as • show ip bgp vpnv4 |