



BGP-Multiple Cluster IDs

The BGP—Multiple Cluster IDs feature allows an iBGP neighbor (usually a route reflector) to have multiple cluster IDs: a global cluster ID and additional cluster IDs that are assigned to clients (neighbors). Prior to the introduction of this feature, a device could have a single, global cluster ID.

When a network administrator configures per-neighbor cluster IDs:

- The loop prevention mechanism based on a `CLUSTER_LIST` is automatically modified to take into account multiple cluster IDs.
- A network administrator can disable client-to-client route reflection based on cluster ID.
- [Information About BGP-Multiple Cluster IDs, on page 1](#)
- [How to Use BGP-Multiple Cluster IDs, on page 4](#)
- [Configuration Examples for BGP-Multiple Cluster IDs, on page 9](#)
- [Additional References, on page 10](#)
- [Feature Information for BGP-Multiple Cluster IDs, on page 11](#)

Information About BGP-Multiple Cluster IDs

Benefit of Multiple Cluster IDs Per Route Reflector

The BGP—Multiple Cluster IDs feature allows a route reflector (RR) to belong to multiple clusters, and therefore have multiple cluster IDs. An RR can have a cluster ID configured on a global basis and a per-neighbor basis. A single cluster ID can be assigned to two or more iBGP neighbors. Prior to this feature, an RR had a single, global cluster ID, which was configured by the `bgp cluster-id` router configuration command.

When a cluster ID is configured per neighbor (by the `neighbor cluster-id` router configuration command), the following two changes occur:

- The loop prevention mechanism based on the `CLUSTER_LIST` attribute is automatically modified to take into account multiple cluster IDs.
- The network administrator can disable client-to-client route reflection based on cluster ID, which allows the network design to change.

The loop prevention mechanism and the CLUSTER_LIST propagation rules are described in the section “How a CLUSTER_LIST Attribute is Used.” Disabling client-to-client reflection is described in the section “Behaviors When Disabling Client-to-Client Route Reflection.”

How a CLUSTER_LIST Attribute is Used

The CLUSTER_LIST propagation rules differ among releases, depending on whether the device is running a Cisco software release generated before or after the BGP—Multiple Cluster IDs feature was implemented. The same is true for loop prevention based on the CLUSTER_LIST.

The CLUSTER_LIST behavior is described below. Classic refers to the behavior of software released before the multiple cluster IDs feature was implemented; MCID refers to the behavior of software released after the feature was implemented.

CLUSTER_LIST Propagation Rules

- Classic—Before reflecting a route, the RR appends the global cluster ID to the CLUSTER_LIST. If the received route had no CLUSTER_LIST attribute, the RR creates a new CLUSTER_LIST attribute with that global cluster ID.
- MCID—Before reflecting a route, the RR appends the cluster ID of the neighbor the route was received from to the CLUSTER_LIST. If the received route had no CLUSTER_LIST attribute, the RR creates a new CLUSTER_LIST attribute with that cluster ID. This behavior includes a neighbor that is not a client of the speaker. If the nonclient neighbor the route was received from does not have an associated cluster ID, the RR uses the global cluster ID.

Loop Prevention Based on CLUSTER_LIST

- Classic—When receiving a route, the RR discards the route if the RR's global cluster ID is contained in the CLUSTER_LIST of the route.
- MCID—When receiving a route, the RR discards the route if the RR's global cluster ID or any of the cluster IDs assigned to any of the iBGP neighbors is contained in the CLUSTER_LIST of the route.

Behaviors When Disabling Client-to-Client Route Reflection

With the introduction of multiple cluster IDs per iBGP neighbor, it is possible to disable route reflection from client to client on the basis of cluster ID. Disabling route reflection allows you to change the network design. A typical (but not required) scenario after disabling route reflection is that clients are fully meshed, so they have to send more updates, and the RR has client-to-client reflection disabled, so that it has to send fewer updates.

You might want to disable route reflection in a scenario similar to the one in the figure below. An RR has several clients [Provider-Edge (PE) routers] with which it has sessions. The iBGP neighbors that should belong to one cluster were assigned the same cluster ID.

Because the PEs belonging to the same cluster are fully meshed (PE1 and PE2 have a session between them; PE3 and PE4 have a session between them), there is no need to reflect the routes between them. That is, routes from PE1 should be forwarded to PE3 and PE4, but not to PE2.

It is important to know that when the software changes reflection state for a given cluster ID, BGP sends an outbound soft refresh to all clients.

Disabling client-to-client route reflection is done differently and has different results, depending on whether the device is running Cisco software generated before or after the multiple cluster IDs feature was implemented. Classic refers to the behavior of software released before the multiple cluster IDs feature was implemented; MCID refers to the behavior of software released after the multiple cluster IDs feature was implemented.

- Classic—When receiving a route from a client, the RR does not reflect it to any other client. Other scenarios for reflection (client-to-nonclient and nonclient-to-client) are maintained. Disabling of route reflection from client to client is usually done when all the clients are fully meshed (the routes are advertised between the clients via that mesh, so there is no need for reflection). The command to disable client-to-client route reflection is entered in router configuration mode (after the **router bgp** command) and it applies globally to all address families: **no bgp client-to-client reflection**
- MCID—When receiving a route from a client, the RR does not reflect it to another client if both clients belong to a cluster for which client-to-client reflection has been disabled. Therefore, route reflection is disabled only intracluster (within the cluster specified). Other cases for reflection (client-to-nonclient, nonclient-to-client, and intercluster) are maintained. This functionality is usually configured when all the clients for a particular cluster are fully meshed among themselves (but not with clients of other clusters). The command to disable client-to-client route reflection for a particular cluster is entered in router configuration mode and it applies globally to all address families:

no bgp client-to-client reflection intra-cluster cluster-id {any | cluster-id1 cluster-id2...}

The **any** keyword is used to disable client-to-client reflection for any cluster.

The Classic, previously released command for disabling all client-to-client reflection is also still available during this post-MCID release timeframe:

no bgp client-to-client reflection [all]

(The optional **all** keyword has no effect in either the positive or negative form of the command, and does not appear in configuration files. It is just to remind the network administrator that both intercluster and intracluster client-to-client reflection are enabled or disabled.)

In summary, after the introduction of the multiple cluster IDs feature, there are three levels of configuration that can disable client-to-client reflection. The software performs them in the following order, from least specific to most specific:

1. Least specific: **no bgp client-to-client reflection [all]** Disables intracluster and intercluster client-to-client reflection.
2. More specific: **no bgp client-to-client reflection intra-cluster cluster-id any** Disables intracluster client-to-client reflection for any cluster-id.
3. Most specific: **no bgp client-to-client reflection intra-cluster cluster-id cluster-id1 cluster-id2 ...** Disables intracluster client-to-client reflection for the specified clusters.

When BGP is advertising updates, the software evaluates each level of configuration in order. Once any level of configuration disables client-to-client reflection, no further evaluation of more specific policies is necessary.

Note the results of the base (positive) and negative (**no**) forms of the three commands listed above:

- A negative configuration (that is, with the **no** keyword) overwrites any less specific configuration.
- A positive configuration (that is, without the **no** keyword) will lose out to (default to) what is configured in a less specific configuration.
- Configurations at any level appear in the configuration file only if they are negative.

All levels can be configured independently and all levels appear in the configuration file independently of the configuration of other levels.

Note that negative configuration makes any more specific configuration unnecessary (because even if the more specific configuration is positive, it is not processed after the negative configuration; if the more specific configuration is negative, it is functionally the same as the earlier negative configuration). The following examples illustrate this behavior.

Example 1

no bgp client-to-client reflection

no bgp client-to-client reflection intra-cluster cluster-id any

Intercluster and intracluster reflection are disabled (based on the first command). The second command disables intracluster reflection, but it is unnecessary because intracluster reflection is already disabled by the first command.

Example 2

no bgp client-to-client reflection intra-cluster cluster-id any

bgp client-to-client reflection intra-cluster cluster-id 1.1.1.1

Cluster ID 1.1.1.1 has intracluster route reflection disabled (even though the second command is positive), because the first command is used to evaluate the update. The first command was negative, and once any level of configuration disables client-to-client reflection, no further evaluation is performed.

Another way to look at this example is that the second command, because it is in a positive form, defaults to the behavior of the first command (which is less specific). Thus, the second command is unnecessary.

Note that the second command would not appear in a configuration file because it is not a negative command.

How to Use BGP-Multiple Cluster IDs

Configuring a Cluster ID per Neighbor

Perform this task on an iBGP peer (usually a route reflector) to configure a cluster ID per neighbor. Configuring a cluster ID per neighbor causes the loop-prevention mechanism based on the CLUSTER_LIST to be automatically modified to take into account multiple cluster IDs. Also, you gain the ability to disable client-to-client route reflection on the basis of cluster ID. The software tags the neighbor so that you can disable route reflection with the use of another command. (See the tasks for disabling client-to-client reflection later in this module.)



Note When you change a cluster ID for a neighbor, BGP automatically does an inbound soft refresh and an outbound soft refresh for all iBGP peers.

SUMMARY STEPS

1. enable

2. **configure terminal**
3. **router bgp *as-number***
4. **neighbor {*ip-address* | *ipv6-address*} remote-as *autonomous-system-number***
5. **neighbor {*ip-address* | *ipv6-address*} cluster-id *cluster-id***
6. **end**
7. **show ip bgp cluster-ids**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Device(config)# router bgp 65000	Configures a BGP routing process and enters router configuration mode. <ul style="list-style-type: none"> • The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along. Valid numbers are from 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535.
Step 4	neighbor {<i>ip-address</i> <i>ipv6-address</i>} remote-as <i>autonomous-system-number</i> Example: Device(config-router)# neighbor 192.168.1.2 remote-as 65000	Adds an entry to the BGP routing table.
Step 5	neighbor {<i>ip-address</i> <i>ipv6-address</i>} cluster-id <i>cluster-id</i> Example: Device(config-router)# neighbor 192.168.1.2 cluster-id 0.0.0.1	Assigns a cluster ID to the specified neighbor. <ul style="list-style-type: none"> • The cluster ID can be in dotted decimal format (such as 192.168.7.4) or decimal format (such as 23), with a maximum of 4 bytes. • A cluster ID that is configured in decimal format (such as 23) is modified to dotted decimal format (such as 0.0.0.23) when it appears in a configuration file. • When you change a cluster ID for a neighbor, BGP automatically does an inbound soft refresh and an outbound soft refresh for all iBGP peers.

	Command or Action	Purpose
Step 6	end Example: <pre>Device(config-router)# end</pre>	(Optional) Exits to privileged EXEC mode.
Step 7	show ip bgp cluster-ids Example: <pre>Device# show ip bgp cluster-ids</pre>	(Optional) Lists: <ul style="list-style-type: none"> • the global cluster ID (whether configured or not) • all cluster IDs that are configured to a neighbor • all cluster IDs for which the network administrator has disabled reflection

Disabling Intracluster and Intercluster Client-to-Client Reflection

Perform the following task on a route reflector if you want to disable both intracluster and intercluster client-to-client reflection. Doing so is the broadest (least specific) way to disable client-to-client reflection. Before advertising updates, the software evaluates each level of configuration in order from least specific to most specific. Once any level of configuration disables client-to-client reflection, no further evaluation of more specific policies is needed.



Note When the software changes reflection state for a given cluster ID, BGP sends an outbound soft refresh to all clients.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp *as-number***
4. **no bgp client-to-client reflection [all]**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Device> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.

	Command or Action	Purpose
Step 3	router bgp <i>as-number</i> Example: <pre>Device(config)# router bgp 65000</pre>	Configures a BGP routing process and enters router configuration mode. <ul style="list-style-type: none"> The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along. Valid numbers are from 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535.
Step 4	no bgp client-to-client reflection [all] Example: <pre>Device(config-router)# no bgp client-to-client reflection all</pre>	Disables intracluster and intercluster client-to-client route reflection. <ul style="list-style-type: none"> The all keyword is just to emphasize that the bgp client-to-client reflection command affects both intracluster and intercluster reflection; the all keyword has no effect in the positive or negative form of the command.

Disabling Intracluster Client-to-Client Reflection for Any Cluster ID

Perform the following task on a route reflector to disable intracluster client-to-client reflection for any cluster ID. Doing so is considered to be the middle of the three levels of commands available to disable client-to-client reflection. That is, it is more specific than disabling intracluster and intercluster client-to-client reflection, but it is not as specific as disabling intracluster client-to-client reflection for certain cluster IDs.

Before advertising updates, the software evaluates each level of configuration in order from least specific to most specific. Once any level of configuration disables client-to-client reflection, no further evaluation of more specific policies is needed.



Note When the software changes reflection state for a given cluster ID, BGP sends an outbound soft refresh to all clients.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **router bgp *as-number***
4. **no bgp client-to-client reflection intra-cluster cluster-id any**

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
	Device> enable	
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Device(config)# router bgp 65000	Configures a BGP routing process and enters router configuration mode. <ul style="list-style-type: none"> The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along. Valid numbers are from 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535.
Step 4	no bgp client-to-client reflection intra-cluster cluster-id any Example: Device(config-router)# no bgp client-to-client reflection intra-cluster cluster-id any	Disables intracluster client-to-client route reflection for any cluster.

Disabling Intracluster Client-to-Client Reflection for Specified Cluster IDs

Perform the following task on a route reflector to disable intracluster client-to-client reflection for specified cluster IDs. Doing so is considered to be the most specific of the three levels of commands available to disable client-to-client reflection. Before advertising updates, the software evaluates each level of configuration in order from least specific to most specific. Once any level of configuration disables client-to-client reflection, no further evaluation of more specific policies is needed.



Note When the software changes reflection state for a given cluster ID, BGP sends an outbound soft refresh to all clients.

SUMMARY STEPS

1. enable
2. configure terminal
3. router bgp *as-number*
4. no bgp client-to-client reflection intra-cluster cluster-id *cluster-id1* [*cluster-id2...*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	router bgp <i>as-number</i> Example: Device(config)# router bgp 65000	Configures a BGP routing process and enters router configuration mode. <ul style="list-style-type: none"> • The <i>as-number</i> argument indicates the number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along. Valid numbers are from 0 to 65535. Private autonomous system numbers that can be used in internal networks range from 64512 to 65535.
Step 4	no bgp client-to-client reflection intra-cluster cluster-id <i>cluster-id1</i> [<i>cluster-id2...</i>] Example: Device(config-router)# no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.1 0.0.0.3 105	Disables intracluster client-to-client route reflection within each of the specified clusters. <ul style="list-style-type: none"> • Note that this example command will appear in the configuration file as “no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.1 0.0.0.3 0.0.0.105” because decimal cluster ID numbers appear in the dotted decimal format.

Configuration Examples for BGP-Multiple Cluster IDs

Example: Per-Nearbor Cluster ID

The following example is configured on a route reflector. The neighbor (client) at IPv6 address 2001:DB8:1::1 is configured to have the cluster ID of 0.0.0.6:

```
router bgp 6500
 neighbor 2001:DB8:1::1 cluster-id 0.0.0.6
```

Example: Disabling Client-to-Client Reflection

The following example disables all intracluster and intercluster client-to-client reflection:

```
router bgp 65000
 no bgp client-to-client reflection all
```

The following example disables intracluster client-to-client reflection for any cluster ID:

```
router bgp 65000
 no bgp client-to-client reflection intra-cluster cluster-id any
```

The following example disables intracluster client-to-client reflection for the specified cluster IDs 0.0.0.1, 14, 15, and 0.0.0.6:

```
router bgp 65000
 no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.1 14 15 0.0.0.6
```

Remember that a cluster ID specified in the **neighbor cluster-id** command in decimal format (such as 23) will appear in a configuration file in dotted decimal format (such as 0.0.0.23). The decimal format does not appear in the configuration file. The running configuration might look like this:

```
router bgp 65000
 no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.1
 no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.6
 no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.14
 no bgp client-to-client reflection intra-cluster cluster-id 0.0.0.15
```

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

MIBs

MIB	MIBs Link
None	To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL: http://www.cisco.com/go/mibs

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-Multiple Cluster IDs

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—Multiple Cluster IDs

Feature Name	Releases	Feature Information
BGP—Multiple Cluster IDs		<p>The BGP—Multiple Cluster IDs feature allows an iBGP neighbor (usually a route reflector) to have multiple cluster IDs: a global cluster ID and additional cluster IDs that are assigned to clients (neighbors). Prior to the introduction of this feature, a device could have a single, global cluster ID.</p> <p>When a network administrator configures per-neighbor cluster IDs:</p> <ul style="list-style-type: none"> • The loop prevention mechanism based on a <code>CLUSTER_LIST</code> is automatically modified to take into account multiple cluster IDs. • A network administrator can disable client-to-client route reflection based on cluster ID. <p>The following commands were introduced:</p> <ul style="list-style-type: none"> • bgp client-to-client reflection intra-cluster • neighbor cluster-id • show ip bgp cluster-ids <p>The following commands were modified:</p> <ul style="list-style-type: none"> • bgp client-to-client reflection • show ip bgp neighbors • show ip bgp template peer-session • show ip bgp update-group