



# Configuring Internal BGP Features

---

This module describes how to configure internal Border Gateway Protocol (BGP) features. Internal BGP (iBGP) refers to running Border Gateway Protocol (BGP) on networking devices within one autonomous system. BGP is an interdomain routing protocol designed to provide loop-free routing between separate routing domains (autonomous systems) that contain independent routing policies. Many companies now have large internal networks and there are many issues involved in scaling the existing internal routing protocols to match the increasing traffic demands while maintaining network efficiency.

## Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see 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 for Configuring Internal BGP Features”](#) section on page 16.

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

## Contents

- [How to Configure Internal BGP Features, page 1](#)
- [Internal BGP Feature Configuration Examples, page 11](#)
- [Additional References, page 13](#)
- [Command Reference, page 14](#)
- [Feature Information for Configuring Internal BGP Features, page 16](#)

## How to Configure Internal BGP Features

The following sections contain optional internal BGP (iBGP) configuration tasks:



---

**Americas Headquarters:**  
**Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA**

© 2005–2009 Cisco Systems, Inc. All rights reserved.

- [Configuring a Routing Domain Confederation](#) (optional)
- [Configuring a Route Reflector](#) (optional)
- [Adjusting BGP Timers](#) (optional)
- [Configuring the Router to Consider a Missing MED as Worst Path](#) (optional)
- [Configuring the Router to Consider the MED to Choose a Path from Subautonomous System Paths](#) (optional)
- [Configuring the Router to Use the MED to Choose a Path in a Confederation](#) (optional)
- [Configuring Route Dampening](#) (optional)

## Configuring a Routing Domain Confederation

One way to reduce the internal BGP (iBGP) mesh is to divide an autonomous system into multiple subautonomous systems and group them into a single confederation. To the outside world, the confederation looks like a single autonomous system. Each autonomous system is fully meshed within itself, and has a few connections to other autonomous systems in the same confederation. Even though the peers in different autonomous systems have external BGP (eBGP) sessions, they exchange routing information as if they were iBGP peers. Specifically, the next hop, Multi\_Exit\_Discriminator (MED) attribute, and local preference information is preserved. This feature allows you to retain a single Interior Gateway Protocol (IGP) for all of the autonomous systems.

To configure a BGP confederation, you must specify a confederation identifier. To the outside world, the group of autonomous systems will look like a single autonomous system with the confederation identifier as the autonomous system number. To configure a BGP confederation identifier, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp confederation identifier</b> <i>as-number</i>	Configures a BGP confederation.

In order to treat the neighbors from other autonomous systems within the confederation as special eBGP peers, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp confederation peers</b> <i>as-number [as-number]</i>	Specifies the autonomous systems that belong to the confederation.

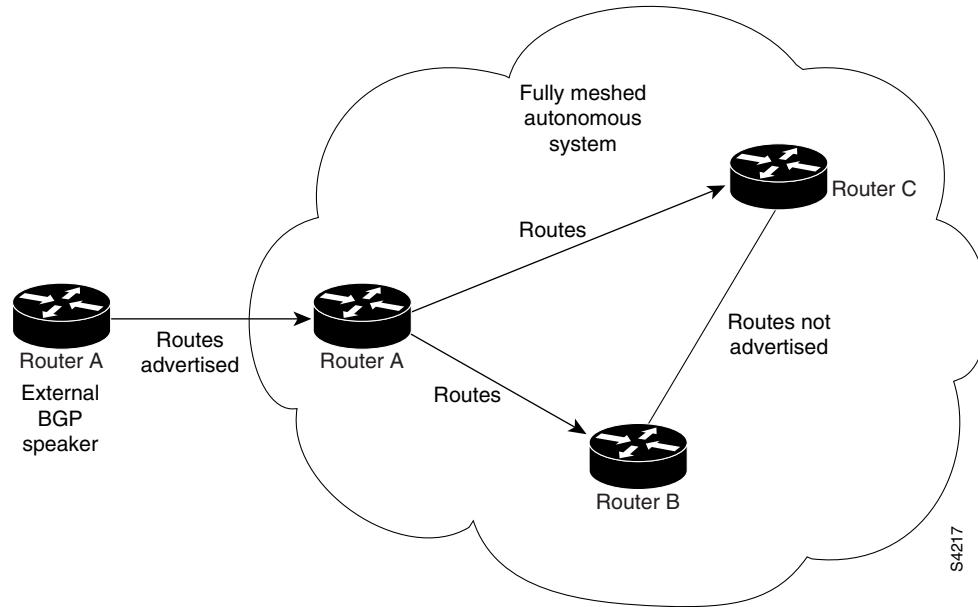
For an alternative way to reduce the iBGP mesh, see “[Configuring a Route Reflector](#).”

## Configuring a Route Reflector

BGP requires that all iBGP speakers be fully meshed. However, this requirement does not scale well when there are many iBGP speakers. Instead of configuring a confederation, another way to reduce the iBGP mesh is to configure a *route reflector*.

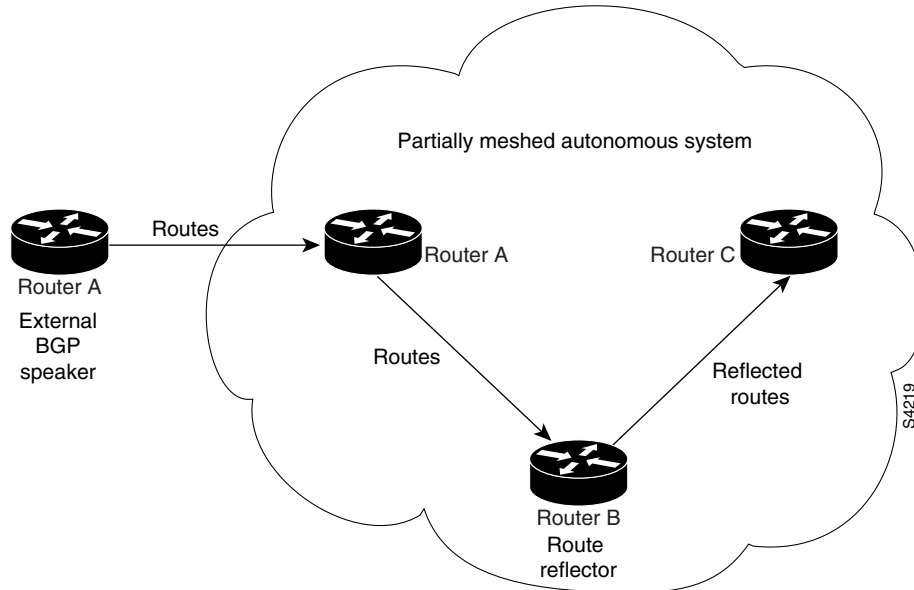
Figure 1 illustrates a simple iBGP configuration with three iBGP speakers (Routers A, B, and C). Without route reflectors, when Router A receives a route from an external neighbor, it must advertise it to both routers B and C. Routers B and C do not readvertise the iBGP learned route to other iBGP speakers because the routers do not pass on routes learned from internal neighbors to other internal neighbors, thus preventing a routing information loop.

**Figure 1** Three Fully Meshed iBGP Speakers



With route reflectors, all iBGP speakers need not be fully meshed because there is a method to pass learned routes to neighbors. In this model, an iBGP peer is configured to be a route reflector responsible for passing iBGP learned routes to a set of iBGP neighbors. In Figure 2, Router B is configured as a route reflector. When the route reflector receives routes advertised from Router A, it advertises them to Router C, and vice versa. This scheme eliminates the need for the iBGP session between Routers A and C.

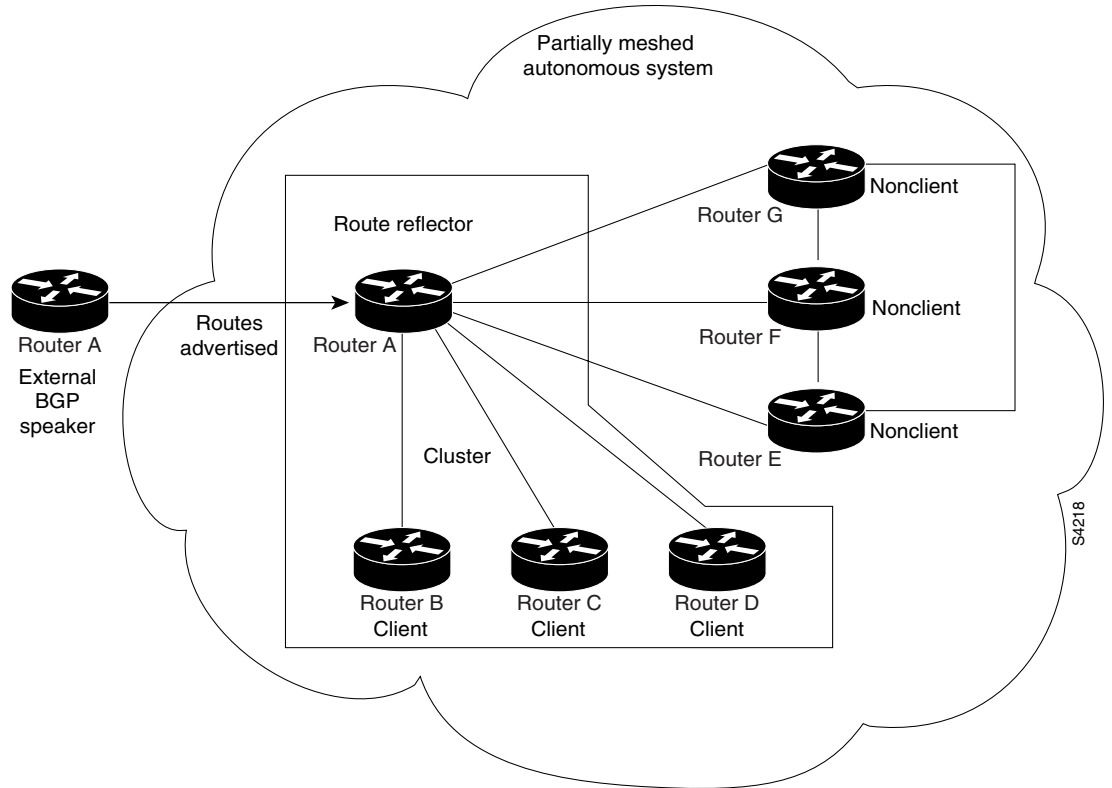
**Figure 2** Simple BGP Model with a Route Reflector



The internal peers of the route reflector are divided into two groups: client peers and all the other routers in the autonomous system (nonclient peers). A route reflector reflects routes between these two groups. The route reflector and its client peers form a *cluster*. The nonclient peers must be fully meshed with each other, but the client peers need not be fully meshed. The clients in the cluster do not communicate with iBGP speakers outside their cluster.

Figure 3 illustrates a more complex route reflector scheme. Router A is the route reflector in a cluster with routers B, C, and D. Routers E, F, and G are fully meshed, nonclient routers.

**Figure 3 More Complex BGP Route Reflector Model**



When the route reflector receives an advertised route, depending on the neighbor, it takes the following actions:

- A route from an external BGP speaker is advertised to all clients and nonclient peers.
- A route from a nonclient peer is advertised to all clients.
- A route from a client is advertised to all clients and nonclient peers. Hence, the clients need not be fully meshed.

To configure a route reflector and its clients, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>neighbor</b> {ip-address   peer-group-name} <b>route-reflector-client</b>	Configures the local router as a BGP route reflector and the specified neighbor as a client.

Along with route reflector-aware BGP speakers, it is possible to have BGP speakers that do not understand the concept of route reflectors. They can be members of either client or nonclient groups allowing an easy and gradual migration from the old BGP model to the route reflector model. Initially, you could create a single cluster with a route reflector and a few clients. All the other iBGP speakers could be nonclient peers to the route reflector and then more clusters could be created gradually.

An autonomous system can have multiple route reflectors. A route reflector treats other route reflectors just like other iBGP speakers. A route reflector can be configured to have other route reflectors in a client group or nonclient group. In a simple configuration, the backbone could be divided into many clusters.

Each route reflector would be configured with other route reflectors as nonclient peers (thus, all the route reflectors will be fully meshed). The clients are configured to maintain iBGP sessions with only the route reflector in their cluster.

Usually a cluster of clients will have a single route reflector. In that case, the cluster is identified by the router ID of the route reflector. To increase redundancy and avoid a single point of failure, a cluster might have more than one route reflector. In this case, all route reflectors in the cluster must be configured with the 4-byte cluster ID so that a route reflector can recognize updates from route reflectors in the same cluster. All the route reflectors serving a cluster should be fully meshed and all of them should have identical sets of client and nonclient peers.

If the cluster has more than one route reflector, configure the cluster ID by using the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp cluster-id</b> <i>cluster-id</i>	Configures the cluster ID.

Use the **show ip bgp** command to display the originator ID and the cluster-list attributes.

By default, the clients of a route reflector are not required to be fully meshed and the routes from a client are reflected to other clients. However, if the clients are fully meshed, the route reflector need not reflect routes to clients.

To disable client-to-client route reflection, use the **no bgp client-to-client reflection** command in router configuration mode:

Command	Purpose
Router(config-router)# <b>no bgp client-to-client reflection</b>	Disables client-to-client route reflection.

As the iBGP learned routes are reflected, routing information may loop. The route reflector model has the following mechanisms to avoid routing loops:

- Originator ID is an optional, nontransitive BGP attribute. It is a 4-byte attribute created by a route reflector. The attribute carries the router ID of the originator of the route in the local autonomous system. Therefore, if a misconfiguration causes routing information to come back to the originator, the information is ignored.
- Cluster-list is an optional, nontransitive BGP attribute. It is a sequence of cluster IDs that the route has passed. When a route reflector reflects a route from its clients to nonclient peers, and vice versa, it appends the local cluster ID to the cluster list. If the cluster list is empty, a new cluster list is created. Using this attribute, a route reflector can identify if routing information is looped back to the same cluster due to misconfiguration. If the local cluster ID is found in the cluster list, the advertisement is ignored.
- The use of **set** clauses in outbound route maps can modify attributes and possibly create routing loops. To avoid this behavior, **set** clauses of outbound route maps are ignored for routes reflected to iBGP peers.

## Adjusting BGP Timers

BGP uses certain timers to control periodic activities such as the sending of keepalive messages and the interval after not receiving a keepalive message after which the Cisco IOS software declares a peer dead. By default, the keepalive timer is 60 seconds, and the hold-time timer is 180 seconds. You can adjust these timers. When a connection is started, BGP will negotiate the hold time with the neighbor. The smaller of the two hold times will be chosen. The keepalive timer is then set based on the negotiated hold time and the configured keepalive time.

To adjust BGP timers for all neighbors, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>timers bgp</b> <i>keepalive holdtime</i>	Adjusts BGP timers for all neighbors.

To adjust BGP keepalive and hold-time timers for a specific neighbor, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>neighbor</b> [ <i>ip-address</i>   <i>peer-group-name</i> ] <b>timers</b> <i>keepalive holdtime</i>	Sets the keepalive and hold-time timers (in seconds) for the specified peer or peer group.



### Note

The timers configured for a specific neighbor or peer group override the timers configured for all BGP neighbors using the **timers bgp** router configuration command.

To clear the timers for a BGP neighbor or peer group, use the **no** form of the **neighbor timers** command.

## Configuring the Router to Consider a Missing MED as Worst Path

To configure the router to consider a path with a missing MED attribute as the worst path, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp bestpath med missing-as-worst</b>	Configures the router to consider a missing MED as having a value of infinity, making the path without a MED value the least desirable path.

## Configuring the Router to Consider the MED to Choose a Path from Subautonomous System Paths

To configure the router to consider the MED value in choosing a path, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp bestpath med confed</b>	Configures the router to consider the MED in choosing a path from among those advertised by different subautonomous systems within a confederation.

The comparison between MEDs is made only if there are no external autonomous systems in the path (an external autonomous system is an autonomous system that is not within the confederation). If there is an external autonomous system in the path, then the external MED is passed transparently through the confederation, and the comparison is not made.

The following example compares route A with these paths:

```
path= 65000 65004, med=2
path= 65001 65004, med=3
path= 65002 65004, med=4
path= 65003 1, med=1
```

In this case, path 1 would be chosen if the **bgp bestpath med confed router configuration** command is enabled. The fourth path has a lower MED, but it is not involved in the MED comparison because there is an external autonomous system in this path.

## Configuring the Router to Use the MED to Choose a Path in a Confederation

To configure the router to use the MED to choose the best path from among paths advertised by a single subautonomous system within a confederation, use the following command in router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp deterministic med</b>	Configures the router to compare the MED variable when choosing among routes advertised by different peers in the same autonomous system.



### Note

If the **bgp always-compare-med** router configuration command is enabled, all paths are fully comparable, including those from other autonomous systems in the confederation, even if the **bgp deterministic med** command is also enabled.

## Configuring Route Dampening

Route dampening is a BGP feature designed to minimize the propagation of flapping routes across an internetwork. A route is considered to be flapping when its availability alternates repeatedly.

For example, consider a network with three BGP autonomous systems: autonomous system 1, autonomous system 2, and autonomous system 3. Suppose the route to network A in autonomous system 1 flaps (it becomes unavailable). Under circumstances without route dampening, the eBGP neighbor of autonomous system 1 to autonomous system 2 sends a withdraw message to autonomous system 2. The border router in autonomous system 2, in turn, propagates the withdraw message to autonomous system 3. When the route to network A reappears, autonomous system 1 sends an advertisement message to

autonomous system 2, which sends it to autonomous system 3. If the route to network A repeatedly becomes unavailable, then available, many withdrawal and advertisement messages are sent. This is a problem in an internetwork connected to the Internet because a route flap in the Internet backbone usually involves many routes.

**Note**

No penalty is applied to a BGP peer reset when route dampening is enabled. Although the reset withdraws the route, no penalty is applied in this instance, even if route flap dampening is enabled.

## Minimizing Flapping

The route dampening feature minimizes the flapping problem as follows. Suppose again that the route to network A flaps. The router in autonomous system 2 (where route dampening is enabled) assigns network A a penalty of 1000 and moves it to history state. The router in autonomous system 2 continues to advertise the status of the route to neighbors. The penalties are cumulative. When the route flaps so often that the penalty exceeds a configurable suppress limit, the router stops advertising the route to network A, regardless of how many times it flaps. Thus, the route is dampened.

The penalty placed on network A is decayed until the reuse limit is reached, upon which the route is once again advertised. At half of the reuse limit, the dampening information for the route to network A is removed.

## Understanding Route Dampening Terms

The following terms are used when describing route dampening:

- **Flap**—A route whose availability alternates repeatedly.
- **History state**—After a route flaps once, it is assigned a penalty and put into history state, meaning the router does not have the best path, based on historical information.
- **Penalty**—Each time a route flaps, the router configured for route dampening in another autonomous system assigns the route a penalty of 1000. Penalties are cumulative. The penalty for the route is stored in the BGP routing table until the penalty exceeds the suppress limit. At that point, the route state changes from history to damp.
- **Damp state**—In this state, the route has flapped so often that the router will not advertise this route to BGP neighbors.
- **Suppress limit**—A route is suppressed when its penalty exceeds this limit. The default value is 2000.
- **Half-life**—Once the route has been assigned a penalty, the penalty is decreased by half after the half-life period (which is 15 minutes by default). The process of reducing the penalty happens every 5 seconds.
- **Reuse limit**—As the penalty for a flapping route decreases and falls below this reuse limit, the route is unsuppressed. That is, the route is added back to the BGP table and once again used for forwarding. The default reuse limit is 750. The process of unsuppressing routes occurs at 10-second increments. Every 10 seconds, the router finds out which routes are now unsuppressed and advertises them to the world.
- **Maximum suppress limit**—This value is the maximum amount of time a route can be suppressed. The default value is four times the half-life.

The routes external to an autonomous system learned via iBGP are not dampened. This policy prevents the iBGP peers from having a higher penalty for routes external to the autonomous system.

## Enabling Route Dampening

To enable BGP route dampening, use the following command in address family or router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp dampening</b>	Enables BGP route dampening.

To change the default values of various dampening factors, use the following command in address family or router configuration mode:

Command	Purpose
Router(config-router)# <b>bgp dampening</b> <i>half-life reuse suppress max-suppress [route-map map-name]</i>	Changes the default values of route dampening factors.

## Monitoring and Maintaining BGP Route Dampening

You can monitor the flaps of all the paths that are flapping. The statistics will be deleted once the route is not suppressed and is stable for at least one half-life. To display flap statistics, use the following commands as needed:

Command	Purpose
Router# <b>show ip bgp flap-statistics</b>	Displays BGP flap statistics for all paths.
Router# <b>show ip bgp flap-statistics regexp</b> <i>regexp</i>	Displays BGP flap statistics for all paths that match the regular expression.
Router# <b>show ip bgp flap-statistics filter-list</b> <i>access-list</i>	Displays BGP flap statistics for all paths that pass the filter.
Router# <b>show ip bgp flap-statistics ip-address</b> <i>mask</i>	Displays BGP flap statistics for a single entry.
Router# <b>show ip bgp flap-statistics ip-address</b> <i>mask longer-prefix</i>	Displays BGP flap statistics for more specific entries.

To clear BGP flap statistics (thus making it less likely that the route will be dampened), use the following commands as needed:

Command	Purpose
Router# <b>clear ip bgp flap-statistics</b>	Clears BGP flap statistics for all routes.
Router# <b>clear ip bgp flap-statistics regexp</b> <i>regexp</i>	Clears BGP flap statistics for all paths that match the regular expression.
Router# <b>clear ip bgp flap-statistics filter-list</b> <i>list</i>	Clears BGP flap statistics for all paths that pass the filter.
Router# <b>clear ip bgp flap-statistics ip-address</b> <i>mask</i>	Clears BGP flap statistics for a single entry.
Router# <b>clear ip bgp ip-address</b> <i>flap-statistics</i>	Clears BGP flap statistics for all paths from a neighbor.

**Note**

The flap statistics for a route are also cleared when a BGP peer is reset. Although the reset withdraws the route, there is no penalty applied in this instance, even if route flap dampening is enabled.

Once a route is dampened, you can display BGP route dampening information, including the time remaining before the dampened routes will be unsuppressed. To display the information, use the following command:

Command	Purpose
Router# <code>show ip bgp dampened-paths</code>	Displays the dampened routes, including the time remaining before they will be unsuppressed.

You can clear BGP route dampening information and unsuppress any suppressed routes by using the following command:

Command	Purpose
Router# <code>clear ip bgp dampening [ip-address network-mask]</code>	Clears route dampening information and unsuppresses the suppressed routes.

## Internal BGP Feature Configuration Examples

The following sections provide internal BGP feature configuration examples:

- [BGP Confederation Configurations with Route Maps Example](#)
- [BGP Confederation Examples](#)

### BGP Confederation Configurations with Route Maps Example

This section contains an example of the use of a BGP confederation configuration that includes BGP communities and route maps. For more examples of how to configure a BGP confederation, see the section “[BGP Confederation Examples](#)” in this chapter.

This example shows how BGP community attributes are used with a BGP confederation configuration to filter routes.

In this example, the route map named *set-community* is applied to the outbound updates to neighbor 172.16.232.50 and the local-as community attribute is used to filter the routes. The routes that pass access list 1 have the special community attribute value local-as. The remaining routes are advertised normally. This special community value automatically prevents the advertisement of those routes by the BGP speakers outside autonomous system 200.

```
router bgp 65000
 network 10.0.1.0 route-map set-community
 bgp confederation identifier 200
 bgp confederation peers 65001
 neighbor 172.16.232.50 remote-as 100
 neighbor 172.16.233.2 remote-as 65001
!
route-map set-community permit 10
 match ip address 1
```

```

set community local-as
!

```

## BGP Confederation Examples

The following is a sample configuration that shows several peers in a confederation. The confederation consists of three internal autonomous systems with autonomous system numbers 6001, 6002, and 6003. To the BGP speakers outside the confederation, the confederation looks like a normal autonomous system with autonomous system number 500 (specified via the **bgp confederation identifier** router configuration command).

In a BGP speaker in autonomous system 6001, the **bgp confederation peers** router configuration command marks the peers from autonomous systems 6002 and 6003 as special eBGP peers. Hence peers 172.16.232.55 and 172.16.232.56 will get the local preference, next hop, and MED unmodified in the updates. The router at 10.16.69.1 is a normal eBGP speaker and the updates received by it from this peer will be just like a normal eBGP update from a peer in autonomous system 6001.

```

router bgp 6001
  bgp confederation identifier 500
  bgp confederation peers 6002 6003
  neighbor 172.16.232.55 remote-as 6002
  neighbor 172.16.232.56 remote-as 6003
  neighbor 10.16.69.1 remote-as 777

```

In a BGP speaker in autonomous system 6002, the peers from autonomous systems 6001 and 6003 are configured as special eBGP peers. 10.70.70.1 is a normal iBGP peer and 10.99.99.2 is a normal eBGP peer from autonomous system 700.

```

router bgp 6002
  bgp confederation identifier 500
  bgp confederation peers 6001 6003
  neighbor 10.70.70.1 remote-as 6002
  neighbor 172.16.232.57 remote-as 6001
  neighbor 172.16.232.56 remote-as 6003
  neighbor 10.99.99.2 remote-as 700

```

In a BGP speaker in autonomous system 6003, the peers from autonomous systems 6001 and 6002 are configured as special eBGP peers. 10.200.200.200 is a normal eBGP peer from autonomous system 701.

```

router bgp 6003
  bgp confederation identifier 500
  bgp confederation peers 6001 6002
  neighbor 172.16.232.57 remote-as 6001
  neighbor 172.16.232.55 remote-as 6002
  neighbor 10.200.200.200 remote-as 701

```

The following is a part of the configuration from the BGP speaker 10.200.200.205 from autonomous system 701 in the same example. Neighbor 172.16.232.56 is configured as a normal eBGP speaker from autonomous system 500. The internal division of the autonomous system into multiple autonomous systems is not known to the peers external to the confederation.

```

router bgp 701
  neighbor 172.16.232.56 remote-as 500
  neighbor 10.200.200.205 remote-as 701

```

# Additional References

The following sections provide references related to configuring internal BGP features.

## Related Documents

Related Topic	Document Title
BGP commands	<a href="#">Cisco IOS IP Routing Protocols Command Reference</a>
BGP overview	<a href="#">Cisco BGP Overview</a>
Basic BGP configuration tasks	<a href="#">Configuring a Basic BGP Network</a>
Connecting to a service provider	<a href="#">Connecting to a Service Provider Using External BGP</a>
Configuring features that apply to multiple IP routing protocols	<a href="#">Configuring IP Routing Protocol-Independent Features</a>

## 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: <a href="http://www.cisco.com/go/mibs">http://www.cisco.com/go/mibs</a>

## RFCs

RFC	Title
RFC 1772	<i>Application of the Border Gateway Protocol in the Internet</i>
RFC 1773	<i>Experience with the BGP Protocol</i>
RFC 1774	<i>BGP-4 Protocol Analysis</i>
RFC 1930	<i>Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)</i>
RFC 2519	<i>A Framework for Inter-Domain Route Aggregation</i>
RFC 2858	<i>Multiprotocol Extensions for BGP-4</i>
RFC 2918	<i>Route Refresh Capability for BGP-4</i>

RFC	Title
RFC 3392	<i>Capabilities Advertisement with BGP-4</i>
RFC 4271	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 4893	<i>BGP Support for Four-octet AS Number Space</i>
RFC 5396	<i>Textual Representation of Autonomous system (AS) Numbers</i>
RFC 5398	<i>Autonomous System (AS) Number Reservation for Documentation Use</i>

## Technical Assistance

Description	Link
<p>The Cisco Support website provides extensive online resources, including documentation and tools for troubleshooting and resolving technical issues with Cisco products and technologies.</p> <p>To receive security and technical information about your products, you can subscribe to various services, such as the Product Alert Tool (accessed from Field Notices), the Cisco Technical Services Newsletter, and Really Simple Syndication (RSS) Feeds.</p> <p>Access to most tools on the Cisco Support website requires a Cisco.com user ID and password.</p>	<a href="http://www.cisco.com/techsupport">http://www.cisco.com/techsupport</a>

## Command Reference

The following commands are introduced or modified in the feature or features documented in this module. For information about these commands, see the *Cisco IOS IP Routing Protocols Command Reference* at [http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp\\_book.html](http://www.cisco.com/en/US/docs/ios/iproute/command/reference/irp_book.html). For information about all Cisco IOS commands, go to the Command Lookup Tool at <http://tools.cisco.com/Support/CLILookup> or to the *Cisco IOS Master Commands List*.

- **bgp always-compare-med**
- **bgp bestpath med confed**
- **bgp bestpath med missing-as-worst**
- **bgp client-to-client reflection**
- **bgp cluster-id**
- **bgp confederation identifier**
- **bgp confederation peers**
- **bgp dampening**
- **bgp deterministic med**
- **clear ip bgp dampening**
- **clear ip bgp flap-statistics**

- **neighbor route-reflector-client**
- **neighbor timers**
- **show ip bgp**
- **show ip bgp dampened-paths**
- **show ip bgp flap-statistics**
- **timers bgp**

# Feature Information for Configuring Internal BGP Features

Table 1 lists the features in this module and provides links to specific configuration information. Only features that were introduced or modified in Cisco IOS Releases 12.2(1), 12.0(3)S, 12.2(27)SBC, 12.2(33)SRB, 12.2(33)SXH, or later release appear in the table.

For information on a feature in this technology that is not documented here, see the “Cisco BGP Features Roadmap.”

Not all commands may be available in your Cisco IOS software release. For release information about a specific command, see the command reference documentation.

Use Cisco Feature Navigator to find information about platform support and software image support. Cisco Feature Navigator enables you to determine which Cisco IOS and Catalyst OS software images support a specific software release, feature set, or platform. To access Cisco Feature Navigator, go to <http://www.cisco.com/go/cfn>. An account on Cisco.com is not required.



## Note

Table 1 lists only the Cisco IOS software release that introduced support for a given feature in a given Cisco IOS software release train. Unless noted otherwise, subsequent releases of that Cisco IOS software release train also support that feature.

**Table 1** Feature Information for Configuring Internal BGP Features

Feature Name	Releases	Feature Configuration Information
Configuring internal BGP features	10.3 12.0(32)S12 12.0(7)T 12.2(33)SRA 12.2(33)SXH	All the features contained in this module are considered to be legacy features and will work in all trains release images.  The following commands were introduced or modified by these features: <b>bgp always-compare-med</b> , <b>bgp bestpath med confed</b> , <b>bgp bestpath med missing-as-worst</b> , <b>bgp client-to-client reflection</b> , <b>bgp cluster-id</b> , <b>bgp confederation identifier</b> , <b>bgp confederation peers</b> , <b>bgp dampening</b> , <b>bgp deterministic med</b> , <b>clear ip bgp dampening</b> , <b>clear ip bgp flap-statistics</b> , <b>neighbor route-reflector-client</b> , <b>neighbor timers</b> , <b>show ip bgp</b> , <b>show ip bgp dampened-paths</b> , <b>show ip bgp flap-statistics</b> , <b>timers bgp</b> .

CCDE, CCENT, CCSI, Cisco Eos, Cisco HealthPresence, Cisco Ironport, the Cisco logo, Cisco Lumin, Cisco Nexus, Cisco Nurse Connect, Cisco Stackpower, Cisco StadiumVision, Cisco TelePresence, Cisco Unified Computing System, Cisco WebEx, DCE, Flip Channels, Flip for Good, Flip Mino, Flip Video, Flip Video (Design), Flipshare (Design), Flip Ultra, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn, Cisco Store, and Flip Gift Card are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0907R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

© 2005–2009 Cisco Systems, Inc. All rights reserved.