



BGP Conditional Route Injection

Feature History

Release	Modification
12.0(14)ST	This feature was introduced.
12.2(4)T	This feature was integrated into Cisco IOS Release 12.2(4)T.
12.2(14)S	This feature was integrated into Cisco IOS Release 12.2(14)S.
12.0(22)S	This feature was integrated into Cisco IOS Release 12.0(22)S.

This feature module describes the BGP Conditional Route Injection feature. It includes the following sections:

- [Feature Overview, page 1](#)
- [Benefits, page 2](#)
- [Related Features and Technologies, page 2](#)
- [Supported Platforms, page 2](#)
- [Supported Standards, MIBs, and RFCs, page 3](#)
- [Configuration Tasks, page 3](#)
- [Monitoring and Maintaining BGP Conditional Route Injection, page 6](#)
- [Configuration Examples, page 7](#)
- [Command Reference, page 7](#)

Feature Overview

Routes that are advertised through the Border Gateway Protocol (BGP) are commonly aggregated to minimize the number of routes that are used and reduce the size of global routing tables. However, common route aggregation can obscure more specific routing information that is more accurate but not necessary to forward packets to their destinations. Routing accuracy is obscured by common route aggregation because a prefix that represents multiple addresses or hosts over a large topological area cannot be accurately reflected in a single route. Cisco IOS software provides several methods in which you can originate a prefix into BGP. The existing methods include redistribution and using the **network** or **aggregate-address** command. These methods assume the existence of more specific routing information (matching the route to be originated) in either the routing table or the BGP table.

The BGP Conditional Route Injection feature allows you to originate a prefix into a BGP routing table without the corresponding match. This feature allows more specific routes to be generated based on administrative policy or traffic engineering information in order to provide more specific control over the forwarding of packets to these more specific routes, which are injected into the BGP routing table only if the configured conditions are met. Enabling this feature will allow you to improve the accuracy of common route aggregation by conditionally injecting or replacing less specific prefixes with more specific prefixes. Only prefixes that are equal to or more specific than the original prefix may be injected. The BGP Conditional Route Injection feature is enabled with the **bgp inject-map exist-map** command. This command uses two route maps (inject-map and exist-map) to install one (or more) more specific prefix into a BGP routing table. The exist-map specifies the prefixes that the BGP speaker will track. The inject-map defines the prefixes that will be created and installed into the local BGP table.

Benefits

The BGP Conditional Route Injection feature allows you to inject more specific prefixes into a BGP routing table over less specific prefixes that were selected through normal route aggregation. These more specific prefixes can be used to provide a finer granularity of traffic engineering or administrative control than is possible with aggregated routes.

Related Features and Technologies

The BGP Conditional Route Injection feature is an extension of the BGP routing protocol. For more information about configuring BGP, conditional advertisement, route summarization and route filtering, refer to the “Configuring BGP” chapter of the Release 12.2 *Cisco IOS IP Configuration Guide* and the *Cisco IOS IP Command Reference, Volume 2 of 3: Routing Protocols*.

Supported Platforms

The BGP Conditional Route Injection feature is supported by all platforms in Cisco IOS Release 12.2(14)S that support BGP:

- Cisco 7200 series
- Cisco 7400 series
- Cisco 7500 series

Determining Platform Support Through Cisco Feature Navigator

Cisco IOS software is packaged in feature sets that support specific platforms. To get updated information regarding platform support for this feature, access Cisco Feature Navigator. Cisco Feature Navigator dynamically updates the list of supported platforms as new platform support is added for the feature.

Cisco Feature Navigator is a web-based tool that enables you to determine which Cisco IOS software images support a specific set of features and which features are supported in a specific Cisco IOS image. You can search by feature or release. Under the release section, you can compare releases side by side to display both the features unique to each software release and the features in common.

To access Cisco Feature Navigator, you must have an account on Cisco.com. If you have forgotten or lost your account information, send a blank e-mail to cco-locksmith@cisco.com. An automatic check will verify that your e-mail address is registered with Cisco.com. If the check is successful, account details with a new random password will be e-mailed to you. Qualified users can establish an account on Cisco.com by following the directions at <http://www.cisco.com/register>

Cisco Feature Navigator is updated regularly when major Cisco IOS software releases and technology releases occur. For the most current information, go to the Cisco Feature Navigator home page at the following URL:

<http://www.cisco.com/go/fn>

Availability of Cisco IOS Software Images

Platform support for particular Cisco IOS software releases is dependent on the availability of the software images for those platforms. Software images for some platforms may be deferred, delayed, or changed without prior notice. For updated information about platform support and availability of software images for each Cisco IOS software release, refer to the online release notes or, if supported, Cisco Feature Navigator.

Supported Standards, MIBs, and RFCs

Standards

No new or modified standards are supported by this feature.

MIBs

No new or modified MIBs are supported by this feature.

To obtain lists of supported MIBs by platform and Cisco IOS release, and to download MIB modules, go to the Cisco MIB website on Cisco.com at the following URL:

<http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml>

RFCs

No new or modified RFCs are supported by this feature.


Configuration Tasks


See the following section for configuration tasks for the BGP Conditional Route Injection feature. Each task in the list is identified as either required or optional.

- [Configuring BGP Conditional Route Injection](#) (required)
- [Verifying BGP Conditional Route Injection](#) (optional)

Configuring BGP Conditional Route Injection

To configure the BGP Conditional Route Injection feature, use the following commands beginning in global configuration mode:

	Command	Purpose
Step 1	Router(config)# router bgp <i>as-number</i>	Places the router in router configuration mode, and configures the router to run a BGP process.
Step 2	Router(config-router)# bgp inject-map <i>ORIGINATE</i> exist-map <i>LEARNED_PATH</i>	Configures the inject-map named <i>ORIGINATE</i> and the exist-map named <i>LEARNED_PATH</i> for conditional route injection.
Step 3	Router(config-router)# exit	Exits router configuration mode, and enters global configuration mode.
Step 4	Router(config)# route-map <i>LEARNED_PATH</i> permit sequence-number	Configures the route map named <i>LEARNED_PATH</i> .
Step 5	Router(config-route-map)# match ip address prefix-list <i>ROUTE</i>	Specifies the aggregate route to which a more specific route will be injected.
Step 6	Router(config-route-map)# match ip route-source prefix-list <i>ROUTE_SOURCE</i>	Configures the prefix list named <i>ROUTE_SOURCE</i> to redistribute the source of the route.  Note The route source is the neighbor address that is configured with the neighbor remote-as command. The tracked prefix must come from this neighbor in order for conditional route injection to occur.
Step 7	Router(config-route-map)# exit	Exits route-map configuration mode, and enters global configuration mode.
Step 8	Router(config)# route-map <i>ORIGINATE</i> permit <i>10</i>	Configures the route map named <i>ORIGINATE</i> .
Step 9	Router(config-route-map)# set ip address prefix-list <i>ORIGINATED_ROUTES</i>	Specifies the routes to be injected.
Step 10	Router(config-route-map)# set community <i>community-attribute</i> additive	Configures the community attribute of the injected routes.
Step 11	Router(config-route-map)# exit	Exits route-map configuration mode, and enters global configuration mode.
Step 12	Router(config)# ip prefix-list <i>ROUTE</i> permit <i>10.1.1.0/24</i>	Configures the prefix list named <i>ROUTE</i> to permit routes from network 10.1.1.0/24.
Step 13	Router(config)# ip prefix-list <i>ORIGINATED_ROUTES</i> permit <i>10.1.1.0/25</i>	Configures the prefix list named <i>ORIGINATED_ROUTES</i> to permit routes from network 10.1.1.0/25.

	Command	Purpose
Step 14	Router(config)# ip prefix-list <i>ORIGINATED_ROUTES</i> permit 10.1.1.128/25	Configures the prefix list named <i>ORIGINATED_ROUTES</i> to permit routes from network 10.1.1.0/25.
Step 15	Router(config)# ip prefix-list <i>ROUTE_SOURCE</i> permit 10.2.1.1/32	Configures the prefix list named <i>ROUTE_SOURCE</i> to permit routes from network 10.2.1.1/32.
		 <p>Note The route source prefix list must be configured with a /32 mask in order for conditional route injection to occur.</p>

**Note**

To enable conditional route injection, the exist-map must contain both the **match ip address prefix-list** and **match ip route-source prefix-list** match clauses in the route map paragraph.

Verifying BGP Conditional Route Injection

To verify that the BGP Conditional Route Injection feature is configured correctly, use the **show ip bgp** or **show ip bgp injected-paths** command.

The following sample output is similar to the output that will be displayed when the **show ip bgp** command is entered:

```
Router# show ip bgp 172.16.0.0
BGP routing table entry for 172.16.0.0/8, version 13
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag:0x200
  Not advertised to any peer
  Local, (injected path from 172.16.0.0/8)
    10.0.0.2 from 10.0.0.2 (2.2.2.2)
      Origin incomplete, localpref 100, valid, external, best
      Community:957874231
  200
    10.0.0.2 from 10.0.0.2 (2.2.2.2)
      Origin incomplete, metric 0, localpref 100, valid, external
```

The following sample output is similar to the output that will be displayed when the **show ip bgp injected-routes** command is entered:

```
Router# show ip bgp injected-paths
BGP table version is 11, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 172.16.0.0       10.0.0.2              0 ?
*> 172.17.0.0/16    10.0.0.2              0 ?
```

Troubleshooting Tips

The BGP Conditional Route Injection feature is based on the injection of a more specific prefix into the BGP routing table when a less specific prefix is present. If conditional route injection is not working properly, check the following:

- If conditional route injection is configured but does not occur, check for the existence of the aggregate prefix in the BGP routing table. The existence (or not) of the tracked prefix in the BGP routing table can be verified with the **show ip bgp** command.
- If the aggregate prefix exists but conditional route injection does not occur, verify that the aggregate prefix is being received from the correct neighbor and the prefix list identifying that neighbor is a /32 match.
- Verify the injection (or not) of the more specific prefix using the **show ip bgp injected-paths** command.
- Verify that the prefix that is being injected is not outside of the scope of the aggregate prefix.
- Ensure that the inject route map is configured with the **set ip address** command and not the **match ip address** command.

Monitoring and Maintaining BGP Conditional Route Injection

To display BGP conditional advertisement information, use the following commands in EXEC mode, as needed:

Command	Purpose
Router# show ip bgp	Displays entries in the BGP routing table.
Router# show ip bgp injected-paths	Displays paths in the BGP routing table that were conditionally injected.
Router# show ip bgp neighbors	Displays information about the TCP and BGP connections to neighbors.

Configuration Examples

This following configuration example configures conditional route injection for the inject-map named ORIGINATE and the exist-map named LEARNED_PATH:

```
router bgp 109

  bgp inject-map ORIGINATE exist-map LEARNED_PATH
  !
  route-map LEARNED_PATH permit 10
    match ip address prefix-list ROUTE
    match ip route-source prefix-list ROUTE_SOURCE
  !
  route-map ORIGINATE permit 10
  set ip address prefix-list ORIGINATED_ROUTES
  set community 14616:555 additive
  !
  ip prefix-list ROUTE permit 10.1.1.0/24
  !
  ip prefix-list ORIGINATED_ROUTES permit 10.1.1.0/25
  ip prefix-list ORIGINATED_ROUTES permit 10.1.1.128/25
  !
  ip prefix-list ROUTE_SOURCE permit 10.2.1.1/32
```

Command Reference

This section documents new and modified commands. All other commands used with this feature are documented in the Cisco IOS Release 12.2 command reference publications.

New Commands

- [bgp inject-map exist-map](#)
- [show ip bgp injected-paths](#)

Modified Commands

- [show ip bgp](#)

bgp inject-map exist-map

To inject a more specific route into a Border Gateway Protocol (BGP) routing table, use the **bgp inject-map exist-map** command in address family or router configuration mode. To disable the conditional injection of a selected route, use the **no** form of this command.

```
bgp inject-map {inject-map-name} exist-map {exist-map-name}[copy-attributes]
```

```
no bgp inject-map {inject-map-name} exist-map {exist-map-name}[copy-attributes]
```

Syntax Description

<i>inject-map-name</i>	Defines the prefixes that will be created and installed to the local BGP table.
<i>exist-map-name</i>	Specifies the prefix that the BGP speaker will track.
copy-attributes	(Optional) Configures the injected route to inherit the attributes of the aggregate route.

Defaults

The BGP Conditional Route Injection feature is not enabled by default.

Command Modes

Address family configuration
Router configuration

Command History

Release	Modification
12.0(14)ST	This command was introduced.
12.2(4)T	This command was integrated into Cisco IOS Release 12.2(4)T.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.

Usage Guidelines

If the **copy-attributes** keyword is not specified when the **bgp inject-map** command is used, the components will use the default attributes for locally originated routes. If the **copy-attribute** keyword is used, the components will inherit the same attributes as the aggregate route.

To enable conditional route injection, the exist-map must contain both the **match ip address prefix-list** and **match ip route-source prefix-list** match clauses in the route map paragraph.

Examples

The following example configures the router for conditional route injection:

```
(config-router)# bgp inject-map map1 exist-map map2 copy-attributes
```

Related Commands

Command	Description
ip prefix-list	Displays information about a prefix list or prefix list entries.
neighbor remote-as	Adds an entry to the BGP or multiprotocol BGP neighbor table.

Command	Description
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.
show ip bgp	Displays entries in the BGP routing table.
show ip bgp injected-paths	Displays injected paths in the BGP routing table.

show ip bgp

To display entries in the Border Gateway Protocol (BGP) routing table, use the **show ip bgp** command in EXEC command.

show ip bgp [*network*] [*network-mask*] [**longer-prefixes**] [**prefix-list** *prefix-list-name* | **route-map** *route-map-name*] [**shorter prefixes** *mask-length*]

Syntax Description

<i>network</i>	(Optional) Network number, entered to display a particular network in the BGP routing table.
<i>network-mask</i>	(Optional) Displays all BGP routes matching the address and mask pair.
longer-prefixes	(Optional) Displays the route and more specific routes.
prefix-list route-map	(Optional) Displays selected routes from a BGP routing table based on the contents of a prefix list or route map.
<i>prefix-list-name</i> <i>route-map-name</i>	(Optional) The name of the route map or prefix list that is specified for the above argument.
shorter prefixes <i>mask-length</i>	(Optional) Displays learned prefixes that are longer than the maximum length but shorter than the specified mask for the prefix.

Command Modes

EXEC

Command History

Release	Modification
10.0	This command was introduced.
12.0	The display of prefix advertisement statistics was added.
12.0(6)T	The display of a message indicating support for route refresh capability was added.
12.0(14)ST	The prefix-list and route-map keywords were added.
12.0(14)ST	The shorter prefixes keyword was added.
12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.

Examples

The following is sample output from the **show ip bgp** command in EXEC mode:

```
Router# show ip bgp

BGP table version is 5, local router ID is 10.0.33.34
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 10.1.0.0         0.0.0.0           0         32768 ?
* 10.2.0.0         10.0.33.35        10         0 35 ?
*>                 0.0.0.0           0         32768 ?
* 10.0.0.0         10.0.33.35        10         0 35 ?
*>                 0.0.0.0           0         32768 ?
*> 192.168.0.0/16  10.0.33.35        10         0 35 ?
```

Table 1 describes the significant fields shown in the display.

Table 1 *show ip bgp Field Descriptions*

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	If shown, the value of the interautonomous system metric.
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

The following is sample output from the **show ip bgp** command in EXEC mode when you specify the **longer-prefixes** keyword:

```
Router# show ip bgp 10.92.0.0 255.255.0.0 longer-prefixes
```

```
BGP table version is 1738, local router ID is 192.168.72.24
```

```
Status codes: s suppressed, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.92.0.0	10.92.72.30	8896		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.1.0	10.92.72.30	8796		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.11.0	10.92.72.30	42482		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.14.0	10.92.72.30	8796		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.15.0	10.92.72.30	8696		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.16.0	10.92.72.30	1400		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.17.0	10.92.72.30	1400		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.18.0	10.92.72.30	8876		32768	?
*	10.92.72.30			0	109 108 ?
*> 10.92.19.0	10.92.72.30	8876		32768	?
*	10.92.72.30			0	109 108 ?

The following is sample output from the **show ip bgp** command in EXEC mode, showing information for prefix 10.0.0.0:

```
Router# show ip bgp 10.0.0.0
```

```
BGP routing table entry for 10.0.0.0/8, version 628
```

```
Paths: (1 available, best #1)
```

```
Advertised to peer-groups:
```

```
ebgp
```

```
Advertised to non-peer-group peers:
```

```
172.30.232.162
```

```
109 65000 297 701 80
```

```
172.30.233.56 from 172.30.233.56 (172.19.185.32)
```

```
Origin incomplete, localpref 100, valid, external, best, ref 2
```



Note

If a prefix has not been advertised to any peer, the display shows “Not advertised to any peer.”

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **prefix-list** keyword:

```
Router# show ip bgp prefix-list ROUTE
```

```
BGP table version is 39, local router ID is 10.0.0.1
```

```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
```

```
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 192.168.1.0	10.0.0.2				0 ?
*	10.0.0.2			0	0 200 ?

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **route-map** keyword:

```
Router# show ip bgp route-map LEARNED_PATH
BGP table version is 40, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 192.168.1.0      10.0.0.2
*                   10.0.0.2           0             0 200 ?
```

The following is sample output from the **show ip bgp** command in privileged EXEC mode when you specify the **shorter-prefixes** keyword:

```
Router# show ip bgp 172.16.0.0/16 shorter-prefixes 8
*> 172.16.0.0      10.0.0.2           0             0 ?
*                   10.0.0.2           0             0 200 ?
```

Related Commands

Command	Description
clear ip bgp	Resets a BGP connection or session.
neighbor soft-reconfiguration	Configures the Cisco IOS software to start storing updates.

show ip bgp injected-paths

To display all the injected paths in the BGP routing table, use the **show ip bgp injected-paths** command in EXEC mode.

show ip bgp injected-paths

Syntax Description This command has no arguments or keywords.

Command Modes EXEC

Command History	Release	Modification
	12.0(14)ST	This command was introduced.
	12.2(4)T	This command was integrated into Cisco IOS Release 12.2(4)T.
	12.2(14)S	This command was integrated into Cisco IOS Release 12.2(14)S.

Examples The following is sample output from the **show ip bgp injected-paths** command in EXEC mode:

```
Router# show ip bgp injected-paths

BGP table version is 11, local router ID is 10.0.0.1
Status codes:s suppressed, d damped, h history, * valid, > best, i -
internal
Origin codes:i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 172.16.0.0       10.0.0.2              0 ?
*> 172.17.0.0/16   10.0.0.2              0 ?
```

[Table 2](#) describes the significant fields shown in the display.

Table 2 *show ip bgp injected-paths Field Descriptions*

Field	Description
BGP table version	Internal version number of the table. This number is incremented whenever the table changes.
local router ID	IP address of the router.

Table 2 *show ip bgp injected-paths Field Descriptions*

Field	Description
Status codes	Status of the table entry. The status is displayed at the beginning of each line in the table. It can be one of the following values: s—The table entry is suppressed. d—The table entry is dampened. h—The table entry history. *—The table entry is valid. >—The table entry is the best entry to use for that network. i—The table entry was learned via an internal BGP (iBGP) session.
Origin codes	Origin of the entry. The origin code is placed at the end of each line in the table. It can be one of the following values: i—Entry originated from Interior Gateway Protocol (IGP) and was advertised with a network router configuration command. e—Entry originated from Exterior Gateway Protocol (EGP). ?—Origin of the path is not clear. Usually, this is a router that is redistributed into BGP from an IGP.
Network	IP address of a network entity.
Next Hop	IP address of the next system that is used when forwarding a packet to the destination network. An entry of 0.0.0.0 indicates that the router has some non-BGP routes to this network.
Metric	The Multi Exit Discriminator (MED) metric for the path. (The name of this metric for BGP versions 2 and 3 is INTER_AS.)
LocPrf	Local preference value as set with the set local-preference route-map configuration command. The default value is 100.
Weight	Weight of the route as set via autonomous system filters.
Path	Autonomous system paths to the destination network. There can be one entry in this field for each autonomous system in the path.

```
■ show ip bgp injected-paths
```