

Send comments to nexus5k-docfeedback@cisco.com



R Commands

This chapter describes the Cisco NX-OS Border Gateway Protocol (BGP) commands that begin with R.

[Send comments to nexus5k-docfeedback@cisco.com](mailto:nexus5k-docfeedback@cisco.com)

redistribute (BGP)

To inject routes from one routing domain into the Border Gateway Protocol (BGP), use the **redistribute** command. To remove the **redistribute** command from the configuration file and restore the system to its default condition in which the software does not redistribute routes, use the **no** form of this command.

```
redistribute { direct | eigrp instance-tag | ospf instance-tag | rip instance-tag | static } [route-map map-name]
```

```
no redistribute { { direct | eigrp instance-tag | ospf instance-tag | rip instance-tag | static } [route-map map-name]
```

Syntax Description		
direct		Distributes routes that are directly connected on an interface.
eigrp <i>instance-tag</i>		Specifies the name of an EIGRP instance. The <i>instance-tag</i> can be any case-sensitive, alphanumeric string up to 20 characters.
ospf <i>instance-tag</i>		Distributes routes from the OSPF protocol. This protocol is supported in the IPv4 address family. The <i>instance-tag</i> can be any case-sensitive, alphanumeric string up to 20 characters.
rip <i>instance-tag</i>		Distributes routes from the RIP protocol. The <i>instance-tag</i> can be any case-sensitive, alphanumeric string up to 20 characters.
static		Redistributes IP static routes.
route-map <i>map-name</i>		(Optional) Specifies the identifier of a configured route map. Use a route map to filter which routes are redistributed into EIGRP.

Command Default Disabled

Command Modes	Release	Modification
	5.0(3)N1(1)	This command was introduced.

Usage Guidelines Use the **redistribute** command to import routes from other routing protocols into BGP. You should always use a route map to filter these routes to ensure that BGP redistributes only the routes that you intend to redistribute.

You must configure a default metric to redistribute routes from another protocol into BGP. You can configure the default metric with the **default-metric** command or with the route map configured with the **redistribute** command.

This command requires the LAN Enterprise Services license.

Examples This example shows how to redistribute BGP routes into an EIGRP autonomous system:

```
switch(config)# router bgp 64496
switch(config-router) address-family ipv4 unicast
switch(config-router-af)# redistribute eigrp 100
```

Send comments to nexus5k-docfeedback@cisco.com

Related Commands

Command	Description
default-metric (BGP)	Sets the default metrics for routes redistributed into BGP.

[Send comments to nexus5k-docfeedback@cisco.com](mailto:nexus5k-docfeedback@cisco.com)

remote-as

To specify the autonomous system (AS) number for a neighbor, use the **remote-as** command. To remove an AS number, use the **no** form of this command.

remote-as *number*

no remote-as *number*

Syntax Description	<i>number</i>
	AS number. The format is x for a two-byte value or x.x for a four-byte value. The range is from 1 to 65535.

Command Default	None
-----------------	------

Command Modes

Release	Modification
5.0(3)N1(1)	This command was introduced.

Usage Guidelines

This command requires the LAN Enterprise Services license.

Examples

This example shows how to configure the neighbor AS number:

```
switch(config)# router bgp 64496
switch(config-router)# neighbor 10.0.0.100
switch(config-router-neighbor)# remote-as 64497
```

Related Commands

Command	Description
feature bgp	Enables BGP on the router.
neighbor	Configures BGP peers.

[Send comments to nexus5k-docfeedback@cisco.com](mailto:nexus5k-docfeedback@cisco.com)

restart (BGP)

To restart a Border Gateway Protocol (BGP) autonomous system and remove all associated neighbors, use the **restart** command.

```
restart bgp as-num[.as-num]
```

Syntax Description

<i>as-num</i>	Number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along; valid values are from 1 to 65535.
<i>.as-num</i>	(Optional) Number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along; valid values are from 0 to 65535.

Command Default

None

Command Modes

Release	Modification
5.0(3)N1(1)	This command was introduced.

Usage Guidelines

This command requires the LAN Enterprise Services license.

Examples

This example shows how to restart the BGP autonomous system:

```
switch# restart bgp 64496
switch#
```

Related Commands

Command	Description
router bgp	Configures a BGP process.

[Send comments to nexus5k-docfeedback@cisco.com](mailto:nexus5k-docfeedback@cisco.com)

route-map

To create a route map, enter the route-map configuration mode, or define the conditions for redistributing routes from one routing protocol into another, use the **route-map** command. To delete an entry, use the **no** form of this command.

```
route-map map-tag [deny | permit] [sequence-number]
```

```
no route-map map-tag [permit | deny] [sequence-number]
```

Syntax Description

<i>map-tag</i>	Route map name.
deny	(Optional) Specifies that the route or packet is not distributed if the match criteria are met for the route map.
permit	(Optional) Specifies that the route or packet is distributed if the match criteria for this route are met.
<i>sequence-number</i>	(Optional) Number that indicates the position a new route map has in the list of route maps already configured with the same name. The no form of this command deletes the position of the route map. Range: 0 to 65535.

Command Default

The **permit** keyword is the default.

Command Modes

Release	Modification
5.0(3)N1(1)	This command was introduced.

Usage Guidelines

If you make changes to a route map that is used by a client, you must exit the route-map configuration submode before the changes take effect in the client. The route-map changes are not propagated to its clients until you exit from the route-map configuration submode or 60 seconds expire since entering the submode.

Once you enter the route-map configuration mode, the following keywords are available:

- **continue** *sequence-number*—Continues on a different entry within the route-map. Range: 0 to 65535
- **description** *description*—Provides a description of the route map. The description can be any alphanumeric string up to 90 characters.
- **exit**—Exits from the current command mode.
- **match**—Matches the values from the specified routing table. The following keywords and arguments are available:
 - **as-path** *name* [*name*]—Specifies the autonomous system (AS) path access list to match. The name can be any alphanumeric string up to 63 characters. See the **match as-path** command for additional information.
 - **community** *name* [*name* | **exact-match**]—Specifies the BGP community list name to match. See the **match community** command for additional information.

Send comments to nexus5k-docfeedback@cisco.com

- **ip**—Configures the IPv4 features. The follow keywords and arguments are available:
 - address** {*access-list-name* [*access-list-name*] | **prefix-list** *ipv4-list-name* [*ipv4-list-name*]}—Specifies the address of the route or packet to match. See the **match ip address** command for additional information.
 - multicast** {**group** *address/length* | **rp** *address/length*}—Specifies the multicast attributes to match. See the **match ip multicast** command for additional information.
 - next-hop**—Matches the next-hop address of the route. See the **match ip next-hop** command for additional information.
 - route-source**—Matches the advertising source address of the route. See the **match ip route-source** command for additional information.
- **no**—Negates a command or set its defaults.
- **set**—Sets the values in the destination routing protocol. The **set** commands specify the routing actions to perform if the criteria enforced by the **match** commands are met. You might want to policy route packets some way other than the obvious shortest path. The following keywords and arguments are available:
 - **as-path**—Prepends a string for a BGP AS-path attribute. See the **set as-path** command for additional information.
 - **comm-list**—Sets the BGP community list (for deletion). See the **set comm-list** command for additional information.
 - **community**—Sets the BGP community attribute. See the **set community** command for additional information.
 - **dampening**—Sets the BGP route flap dampening parameters. See the **set dampening** command for additional information.
 - **forwarding-address**—Sets the forwarding address. See the **set forwarding-address** command for additional information.
 - **level**—Specifies where to import the route. See the **set level** command for additional information.
 - **local-preference**—Specifies the BGP local preference path attribute. See the **set local-preference** command for additional information.
 - **metric**—Sets the metric for the destination routing protocol. See the **set metric** command for additional information.
 - **metric-type**—Sets the type of metric for the destination routing protocol. See the **set metric-type** command for additional information.
 - **origin**—Specifies the BGP origin code. See the **set origin** command for additional information.
 - **tag**—Specifies the tag value for the destination routing protocol. See the **set tag** command for additional information.
 - **weight**—Sets the BGP weight for the routing table. See the **set weight** command for additional information.

Use route maps to redistribute routes.

Send comments to nexus5k-docfeedback@cisco.com

Redistribution

The **redistribute** router configuration command uses the *map-tag* name to reference the route map. Multiple route maps may share the same map tag name.

Use the **route-map** global configuration command and the **match** and **set** route-map configuration commands to define the conditions for redistributing routes from one routing protocol into another. Each **route-map** command has a list of **match** and **set** commands associated with it. The **match** commands specify the match criteria—the conditions under which redistribution is allowed for the current **route-map** command. The **set** commands specify the set actions—the particular redistribution actions to perform if the criteria enforced by the **match** commands are met. The **no route-map** command deletes the route map.

The **match** route-map configuration command has multiple formats. The **match** commands can be given in any order, and all **match** commands must pass to cause the route to be redistributed according to the set actions given with the **set** commands. The **no** forms of the **match** commands remove the specified match criteria.

Use route maps when you want detailed control over how routes are redistributed between routing processes. The destination routing protocol is the one you specify with the **router** global configuration command. The source routing protocol is the one you specify with the **redistribute** router configuration command. See the “Examples” section for an illustration of how route maps are configured.

When you are passing routes through a route map, a route map can have several parts. Any route that does not match at least one **match** clause that relates to a **route-map** command is ignored; that is, the route is not advertised for outbound route maps and is not accepted for inbound route maps. If you want to modify some particular data, you must configure a second route map section with an explicit match specified.

Examples

This example shows how to redistribute Routing Information Protocol (RIP) routes with a hop count equal to 1 into Open Shortest Path First (OSPF). These routes are redistributed into OSPF as external link-state advertisements (LSAs) with a metric type of Type 1, and a tag equal to 1.

```
switch(config)# router ospf 109
switch(config-route-map)# redistribute rip route-map rip-to-ospf
switch(config-route-map)# route-map rip-to-ospf permit
switch(config-route-map)# set metric 5
switch(config-route-map)# set metric-type type1
switch(config-route-map)# set tag 1
```

This example shows how to set the autonomous system path to match BGP autonomous system path access list 20:

```
switch(config)# route-map IGP2BGP
switch(config-route-map)# match as-path 20
```

This example shows how to configure that the routes matching community list 1 have the weight set to 100. Any route that has community 109 has the weight set to 100.

```
switch(config)# ip community-list 1 permit 109
switch(config)# route-map set_weight
switch(config-route-map)# match community 1
switch(config-route-map)# set weight 100
```

This example shows how to configure that the routes matching community list 1 have the weight set to 200. Any route that has community 109 alone has the weight set to 200.

```
switch(config)# ip community-list 1 permit 109
switch(config)# route-map set_weight
switch(config-route-map)# match community 1 exact
```


Send comments to nexus5k-docfeedback@cisco.com

```
switch(config-route-map)# set weight 200
```

This example shows how to configure that the routes match community list LIST_NAME have the weight set to 100. Any route that has community 101 alone has the weight set to 100.

```
switch(config)# ip community-list 1 permit 101
switch(config)# route-map set_weight
switch(config-route-map)# match community LIST_NAME
switch(config-route-map)# set weight 100
```

Related Commands

Command	Description
match as-path	Matches a BGP autonomous system path access list.
match community	Matches a BGP community.
match ip next-hop	Redistributes any routes that have a next-hop router address passed by one of the access lists specified.
match ip route-source	Redistributes routes that have been advertised by routers and access servers at the address specified by the access lists.
match metric	Redistributes routes with the metric specified.
match tag	Redistributes routes in the routing table that match the specified tags.
route-map (IP)	Defines the conditions for redistributing routes from one routing protocol into another.
set as-path	Modifies an autonomous system path for BGP routes.
set community	Sets the BGP communities attribute.
set level	Indicates where to import routes.
set local-preference	Specifies a preference value for the autonomous system path.
set metric	Sets the metric value for a routing protocol.
set metric-type	Sets the metric type for the destination routing protocol.
set next-hop	Specifies the address of the next hop.
set tag	Sets a tag value of the destination routing protocol.
set weight	Specifies the BGP weight for the routing table.

Send comments to nexus5k-docfeedback@cisco.com

route-reflector-client (BGP)

To configure the router as a BGP route reflector and configure the specified neighbor as its client, use the **route-reflector-client** command. To indicate that the neighbor is not a client, use the **no** form of this command.

route-reflector-client

no route-reflector-client

Syntax Description This command has no arguments or keywords.

Command Default There is no route reflector in the autonomous system.

Command Modes BGP Neighbor address-family configuration mode

Command History	Release	Modification
	5.0(3)N1(1)	This command was introduced.

Usage Guidelines Use the **route-reflector-client** command to configure the local router as the route reflector and the specified neighbor as one of its clients. All the neighbors configured with this command will be members of the client group and the remaining BGP peers will be members of the nonclient group for the local route reflector.

Examples This example shows how to configure the local router as a route reflector to the neighbor at 192.168.0.1:

```
switch(config)# router bgp 102
switch(config-router)# neighbor 192.168.0.1 remote-as 201
switch(config-router-neighbor)# address-family ipv4 unicast
switch(config-router-neighbor-af)# route-reflector-client
switch(config-router-neighbor-af)#
```

Related Commands	Command	Description
	address-family (BGP)	Enters the router in address family configuration mode for configuring BGP routing sessions.
	neighbor	Configures a BGP neighbor.
	show ip bgp	Displays entries in the BGP routing table.

[Send comments to nexus5k-docfeedback@cisco.com](mailto:nexus5k-docfeedback@cisco.com)

router bgp

To assign an autonomous system (AS) number to a router and enter the router BGP configuration mode, use the **router bgp** command. To remove an AS number assignment, use the **no** form of this command.

```
router bgp as-num [.as-num]
```

```
no router bgp as-num [.as-num]
```

Syntax Description

<i>as-num</i>	Number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along; valid values are from 1 to 65535.
<i>.as-num</i>	(Optional) Number of an autonomous system that identifies the router to other BGP routers and tags the routing information passed along; valid values are from 0 to 65535.

Command Default

No BGP routing process is enabled by default.

Command Modes

Address-family configuration mode
Neighbor address-family configuration mode
Router BGP configuration mode

Usage Guidelines

Release	Modification
5.0(3)N1(1)	This command was introduced.

The *as-num* is the number for the local BGP speaker and allows you to create a unique identifier for the BGP process on the router.

Once you enter the router BGP configuration mode, the following parameters are available:

- **address-family**—Configures an address-family (router, neighbor, VRF). See the **address-family (BGP)** command for information.
- **bestpath**—Changes the default best path selection algorithm. See the **bestpath** command for information.
- **cluster-id** {*cluster-id* | *cluster-ip-addr*}—Configures the Route Reflector Cluster-ID (router, VRF). Range: 1 to 4294967295. You can enter the cluster identification as a 32-bit quantity or as an IP address. To remove the cluster ID, use the **no** form of this command.
- **confederation** {**identifier** *as-num* [*.as-num*] | **peer** *as-num* [*.as-num*]}—Configures the AS confederation parameters as the routing domain confederation AS or the peer AS in the BGP confederation. To remove the confederation identifier, use the **no** form of this command.

The **confederation** command is used to configure a single autonomous system number to identify a group of smaller autonomous systems as a single confederation. You can use a confederation to divide a large single autonomous system into multiple subautonomous systems and then group them into a single confederation. The subautonomous systems within the confederation exchange routing information. External peers interact with the confederation as if it were a single autonomous system.

Send comments to nexus5k-docfeedback@cisco.com

Each subautonomous system is fully meshed within itself and has a few connections to other autonomous systems within the confederation. Next hop, Multi Exit Discriminator (MED), and local preference information is preserved throughout the confederation, allowing you to retain a single Interior Gateway Protocol (IGP) for all the autonomous systems.

- **enforce-first-as**—Forces BGP to compare an external peer's configured AS number with the first AS in the AS-PATH of the routes received from the peer. In case of a mismatch of AS numbers, the peer is sent an error code update notification message. To disable this feature, use the **no** form of this command.
- **exit**—Exits from the current command mode.
- **fast-external-fallover**—Configures a Border Gateway Protocol (BGP) routing process to immediately reset external BGP peering sessions if the link used to reach these peers goes down. To disable BGP fast external fallover, use the **no** form of this command.

The **fast-external-fallover** command is used to disable or enable fast external fallover for BGP peering sessions with directly connected external peers. The session is immediately reset if the link goes down. Only directly connected peering sessions are supported.

If BGP fast external fallover is disabled, the BGP routing process waits until the default hold timer expires (three keepalives) to reset the peering session.

- **log-neighbor-changes**—Enables logging of the BGP neighbor resets. To disable the logging of changes in BGP neighbor adjacencies, use the **no** form of this command. The **log-neighbor-changes** command enables logging of BGP neighbor status changes (up or down) and resets for troubleshooting network connectivity problems and measuring network stability. Unexpected neighbor resets might indicate high error rates or high packet loss in the network and should be investigated.

Using the **log-neighbor-changes** command to enable status change message logging does not cause a substantial performance impact, unlike, for example, enabling per BGP update debugging. If the UNIX syslog facility is enabled, messages are sent to the UNIX host that is running the syslog daemon so that the messages can be stored and archived. If the UNIX syslog facility is not enabled, the status change messages are retained in the internal buffer of the router and are not stored to the disk. You can set the size of this buffer, which is dependent upon the available RAM, using the **logging buffered** command.

The neighbor status change messages are not tracked if the **bgp log-neighbor-changes** command is not enabled, except for the reset reason, which is always available as output of the **show ip bgp neighbors** command.

The **eigrp log-neighbor-changes** command enables logging of Enhanced Interior Gateway Routing Protocol (EIGRP) neighbor adjacencies, but messages for BGP neighbors are logged only if they are specifically enabled with the **bgp log-neighbor-changes** command.

Use the **show logging command** to display the log for the BGP neighbor changes.

- **neighbor**—Configures a BGP neighbor (router, VRF). See the **neighbor** command for additional information.
- **no**—Negates a command or sets its defaults.
- *router-id*—Specifies the IP address to use as router-id (router, VRF).
- **template**—Enters the template command mode. See the **neighbor** command for additional information.
- **timers**—Configures the BGP-related timers (router, VRF).
 - **bestpath-limit interval**—Configures the timeout for the first best path after a restart, in seconds. Range: 1 to 3600. Default: 300.

Send comments to nexus5k-docfeedback@cisco.com

- **bgp interval**—Configures the different BGP keepalive and holdtimes in seconds. Range: 0 to 3600. Default: 60.
- **prefix-peer-timeout interval**—Configures how long a prefix peer is maintained in seconds. Range: 0 to 1200. Default: 300:
- **vrf**—Configures the virtual router context:
 - *vrf-name*—Specifies the VRF name.
 - **management**—Specifies the configurable VRF name.

This command requires the LAN Enterprise Services license.

Examples

This example shows how to configure a BGP process for autonomous system 120:

```
switch(config)# router bgp 120
switch(config-router)#
```

This example shows how to log neighbor changes for BGP in router configuration mode:

```
switch(config)# bgp router 40000
switch(config-router)# log-neighbor-changes
```

This example shows how to disable the BGP fast external fallover feature. If the link through which this session is carried flaps, the connection is not reset.

```
switch(config)# bgp router 64496
switch(config-router)# no fast-external-fallover
```

This example shows how all incoming updates from eBGP peers are examined to ensure that the first autonomous system number in the AS_PATH is the local AS number of the transmitting peer. The updates from the 10.100.0.1 peer are discarded if the first AS number is not 65001.

```
switch(config)# router bgp 64496
switch(config-router)# bgp enforce-first-as
switch(config-router)# address-family ipv4
switch(config-router-af)# neighbor 10.100.0.1 remote-as 64496
switch(config-router-af)#
```

Related Commands

Command	Description
show ip bgp	Displays entries in the BGP table.

Send comments to nexus5k-docfeedback@cisco.com

Send comments to nexus5k-docfeedback@cisco.com

Send comments to nexus5k-docfeedback@cisco.com