

# Chapter 30

## Summary of BGP Configuration Statements

The following sections explain each of the Border Gateway Protocol (BGP) configuration statements. The statements are organized alphabetically.

### advertise-inactive

<b>Syntax</b>	advertise-inactive;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Have BGP advertise its routes even if the routing table did not select them to be active routes.
<b>Usage Guidelines</b>	See “Have BGP Advertise Inactive Routes” on page 409.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## allow

<b>Syntax</b>	allow [ <i>network/mask-length</i> ];
<b>Hierarchy Level</b>	[edit protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ]
<b>Description</b>	Implicitly configure BGP peers, allowing peer connections from any of the specified networks or hosts. To configure multiple BGP peers, configure one or more networks and hosts within a single allow statement or include multiple allow statements.
<b>Options</b>	<i>network/mask-length</i> —IPv6 or IPv4 network number of a single address or a range of allowable addresses for BGP peers, followed by the number of significant bits in the subnet mask. To allow all addresses, you can specify all, which is equivalent to 0.0.0.0/0 (or ::/0).
<b>Usage Guidelines</b>	See “Minimum BGP Configuration” on page 371 and “Define BGP Groups and Peers” on page 374.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	neighbor on page 435

## as-override

<b>Syntax</b>	as-override;
<b>Hierarchy Level</b>	[edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	<p>Compare the AS path of an incoming advertised route with the AS number of the BGP peer under the group and replace all occurrences of the peer AS number in the AS path with its own AS number before advertising the route to the peer.</p> <p>Enabling the AS override feature allows routes originating from an AS to be accepted by a router residing in the same AS. Without AS override enabled, the router refuses the route advertisement once the AS path shows that the route originated from its own AS. This is done by default to prevent route loops. The as-override statement overrides this default behavior.</p> <p>Note that enabling the AS override feature results in possible occurrence of routing loops. This feature should only be used for specific applications that require this type of behavior, and in situations with strict network control. One application is the IGP protocol between the Provider-Edge router and the Customer-Edge router within a Virtual Private Network. For more information, see the <i>JUNOS Internet Software Configuration Guide: MPLS Applications</i>.</p>
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## authentication-key

<b>Syntax</b>	authentication-key <i>key</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Configure an MD5 authentication key (password). Neighboring routers use the same password to verify the authenticity of BGP packets sent from this system.
<b>Options</b>	<i>key</i> —Authentication password. It can be up to 255 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
<b>Usage Guidelines</b>	See “Configure Authentication” on page 385.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## bgp

<b>Syntax</b>	bgp { ... }
<b>Hierarchy Level</b>	[edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
<b>Description</b>	Enable BGP on the router or for a routing instance.
<b>Default</b>	BGP is disabled.
<b>Usage Guidelines</b>	See “Enable BGP” on page 372.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## cluster

<b>Syntax</b>	cluster <i>cluster-identifier</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.
<b>Options</b>	<i>cluster-identifier</i> —IPv6 or IPv4 address to use as the cluster identifier.
<b>Usage Guidelines</b>	See “Configure Route Reflection” on page 399.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	no-client-reflect on page 437

## damping

<b>Syntax</b>	damping;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Enable route flap damping.
<b>Default</b>	Flap damping is disabled on the router.
<b>Usage Guidelines</b>	See “Enable Route Flap Damping” on page 405 and the <i>JUNOS Internet Software Configuration Guide: Policy Framework</i> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## description

<b>Syntax</b>	<code>description text-description;</code>
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Text description of the global, group, or neighbor configuration.
<b>Options</b>	<i>text-description</i> —Text description of the configuration.
<b>Usage Guidelines</b>	See “Define BGP Global Properties” on page 373, “Define Group Properties” on page 377, and “Define Peer Properties” on page 379.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## disable

<b>Syntax</b>	<code>disable;</code>
<b>Hierarchy Level</b>	[edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
<b>Description</b>	Disable BGP on the system.
<b>Usage Guidelines</b>	See “Define BGP Global Properties” on page 373.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- export

- **Syntax** export [ *policy-names* ];

- **Hierarchy Level** [edit protocols bgp],
- [edit protocols bgp group *group-name*],
- [edit protocols bgp group *group-name* neighbor *address*],
- [edit routing-instances *routing-instance-name* protocols bgp],
- [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],
- [edit routing-instances *routing-instance-name* protocols bgp group *group-name*
- neighbor *address*]

- **Description** Apply one or more policies to routes being exported from the routing table into BGP.

- **Options** *policy-names*—Required Privilege Level

- **Usage Guidelines** See “Configure BGP Routing Policy” on page 408 and the *JUNOS Internet Software Configuration Guide: Policy Framework*.

- **Required Privilege Level** routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

- **See Also** import on page 427, *JUNOS Internet Software Configuration Guide: Policy Framework*

## family

<b>Syntax</b>	<pre> family (inet   inet6   inet-vpn   I2-vpn) {   (any   multicast   unicast) {     prefix-limit {       maximum <i>number</i>;       teardown &lt; <i>percentage</i>&gt; &lt; idle-timeout (forever   <i>time-in-minutes</i>)&gt;;     }     rib-group <i>routing-table-group-name</i>;   }   labeled-unicast {     prefix-limit {       maximum <i>number</i>;       teardown &lt; <i>percentage</i>&gt; &lt; idle-timeout (forever   <i>time-in-minutes</i>)&gt; ;     }     resolve-vpn;     rib-group <i>routing-table-group-name</i>;   } } </pre>
<b>Hierarchy Level</b>	<pre> [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>] </pre>
<b>Description</b>	<p>Enable multiprotocol BGP (MBGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, or to specify MBGP to carry NLRI for the IPv6 address family.</p>
<b>Options</b>	<p><b>any</b>—Configure the family type to be both unicast and multicast.</p> <p><b>multicast</b>—Configure the family type to be multicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by multicast for resolving the multicast routes.</p> <p><b>unicast</b>—Configure the family type to be unicast. This means that the BGP peers only carry the unicast routes that are being used for unicast forwarding purposes.</p> <p><b>Default:</b> unicast</p> <p>The remaining statements are explained separately in this chapter.</p>
<b>Usage Guidelines</b>	<p>See “Enable Multiprotocol BGP” on page 406.</p>
<b>Required Privilege Level</b>	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

- graceful-restart

- **Syntax** graceful-restart {
- disable;
- restart-time *seconds*;
- stale-routes-time *seconds*;
- }

- **Hierarchy Level** [edit protocols bgp],
- [edit protocols bgp group *group-name*],
- [edit protocols bgp group *group-name* neighbor *address*],

- **Description** Configure graceful restart for BGP.

- **Options** disable—Disables graceful restart for BGP.

- *seconds*—Time period when restart is expected to be complete.
- **Range:** 1 through 600 seconds

- *seconds*—Maximum time that stale routes are kept during restart.
- **Range:** 1 through 600 seconds

- **Usage Guidelines** See “Configure Graceful Restart” on page 87 and “Configure Graceful Restart” on page 384.

- **Required Privilege Level** routing—To view this statement in the configuration.
- routing-control—To add this statement to the configuration.

## group

```

Syntax  group group-name {
    advertise-inactive;
    allow [ network/mask-length ];
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family (inet | inet6 | inet-vpn | I2-vpn) {
        (any | multicast | unicast) {
            prefix-limit {
                maximum number;
                teardown < percentage> <idle-timeout (forever | time-in-minutes)>;
            }
            rib-group routing-table-group-name;
        }
        labeled-unicast {
            prefix-limit {
                maximum number;
                teardown < percentage> <idle-timeout (forever | time-in-minutes)>;
            }
            resolve-vpn;
            rib-group routing-table-group-name;
        }
    }
    hold-time seconds;
    import [ policy-names ];
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <private>;
    local-preference local-preference;
    log-updown;
    metric-out metric;
    multihop <ttl-value>;
    multipath;
    no-aggregator-id;
    no-client-reflect;
    out-delay seconds;
    passive;
    peer-as autonomous-system;
    preference preference;
    protocol protocol;
    remove-private;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
    type type;
    neighbor address {
        numerous peer-specific options;
    }
}

```



## hold-time

<b>Syntax</b>	hold-time <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Hold-time value to use when negotiating a connection with the peer. The hold-time value is advertised in open packets and indicates to the peer the length of time that it should consider the sender valid. If the peer does not receive a keepalive, update, or notification message within the specified hold time, the BGP connection to the peer is closed and routers through that peer become unavailable.  The hold time is three times the interval at which keepalive messages are sent.
<b>Options</b>	<i>seconds</i> —Hold time. If you set the hold-time value to 0, the hold timer is never started and the router never sends keepalive messages. <b>Range:</b> 6 through 65,535 seconds <b>Default:</b> 90 seconds
<b>Usage Guidelines</b>	See “Modify the Hold-Time Value” on page 384.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## import

<b>Syntax</b>	import [ <i>policy-names</i> ];
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Apply one or more routing policies to routes being imported into the JUNOS routing table from BGP.
<b>Options</b>	<i>policy-names</i> —Name of one or more policies.
<b>Usage Guidelines</b>	See “Configure BGP Routing Policy” on page 408 and the <i>JUNOS Internet Software Configuration Guide: Policy Framework</i> .
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	export on page 422, <i>JUNOS Internet Software Configuration Guide: Policy Framework</i>

## ipsec-sa

<b>Syntax</b>	<code>ipsec-sa ipsec-sa;</code>
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Apply a security association to BGP peers. You can apply the security association globally for all BGP peers, to a group of peers, or to an individual peer.
<b>Options</b>	<i>ipsec-sa</i> —Security association name.
<b>Usage Guidelines</b>	See “Apply IPsec Security Association” on page 385.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## keep

<b>Syntax</b>	<code>keep (all   none);</code>
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Specify whether routes learned from a BGP peer are retained in the routing table even if they contain an AS number that was exported from the local AS.
<b>Default</b>	If you do not include this statement, most routes are retained in the routing table.
<b>Options</b>	<i>all</i> —Retain all routes.  <i>none</i> —Retain none of the routes. When <i>keep none</i> is configured for the BGP session and the inbound policy changes, the JUNOS software forces readvertisement of the full set of routes advertised by the peer.
<b>Usage Guidelines</b>	See “Configure How Often BGP Exchanges Routes with the Routing Table” on page 410.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## labeled-unicast

<b>Syntax</b>	labeled-unicast { resolve-vpn; }
<b>Hierarchy Level</b>	[edit protocols bgp group family inet]
<b>Description</b>	This statement advertises labeled routes from and places labeled routes into the inet.0 vpn. When you use the labeled-unicast statement, the local router will automatically perform next-hop self on all routes advertised into EBGp from IBRP, and from IBRP to EBGp.
<b>Options</b>	resolve-vpn—Allows labeled routes to be placed in the inet.3 routing table for route resolution of PE router connections when the remote PE is located across another AS. Note that in order for a PE router to install a route in the VRF, the next hop must resolve to a route stored within the inet.3 table.
<b>Usage Guidelines</b>	See “Define BGP Global Properties” on page 373.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## local-address

<b>Syntax</b>	local-address <i>address</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Address of the local end of a BGP session. This address is used to accept incoming connections to the peer and to establish connections to the remote peer. When none of the operational interfaces are configured with the specified local address, a session with a BGP peer is placed in the idle state.
<b>Default</b>	If you do not configure a local address, BGP uses the router’s source address selection rules to set the local address. For more information, see the <i>JUNOS Internet Software Configuration Guide: Interfaces and Class of Service</i> .
<b>Options</b>	<i>address</i> —IPv6 or IPv4 address of the local end of the connection.
<b>Usage Guidelines</b>	See “Assign a BGP Identifier” on page 373.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	router-id on page 123

## local-as

**Syntax** local-as *autonomous-system* <private>;

**Hierarchy Level** [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
 neighbor *address*]

**Description** Set the local AS number.

**Options** *autonomous-system*—AS number.

*private*—(Optional) Hide the local AS in paths learned from this peering.

**Usage Guidelines** See “Configure a Local AS” on page 396.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

## local-interface

**Syntax** local-interface *interface-name*;

**Hierarchy Level** [edit protocols bgp group *group-name* neighbor *ipv6-link-local-address*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
 neighbor *ipv6-link-local-address*]

**Description** Specify interface name of peer for IPv6 peering using link-local addresses. This peer is link-local in scope.

**Options** *interface-name*—Interface name of the EBGp IPv6 peer.

**Usage Guidelines** See “Configure EBGp Peering using IPv6 Link-local Address” on page 411.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

## local-preference

<b>Syntax</b>	local-preference <i>local-preference</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Modify the value of the LOCAL_PREF path attribute, which is a metric used by internal BGP sessions to indicate the degree of preference for an external route. The route with the highest local preference value is preferred.  The LOCAL_PREF path attribute always is advertised to internal BGP peers and to neighboring confederations. It is never advertised to external BGP peers.
<b>Default</b>	If you do omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
<b>Options</b>	<i>local-preference</i> —Preference to assign to routes learned from BGP or from the group or peer. <b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ ) <b>Default:</b> If the LOCAL_PREF path attribute is present, do not modify its value. If a BGP route is received without a LOCAL_PREF attribute, the route is handled locally (it is stored in the routing table and advertised by BGP) as if it were received with a LOCAL_PREF value of 100. A non-BGP route that is advertised by BGP is advertised with a LOCAL_PREF value of 100 by default.
<b>Usage Guidelines</b>	See “Configure the BGP Local Preference” on page 391.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	preference on page 439.

log-updown

**Syntax** log-updown;

**Hierarchy Level** [edit protocols bgp],  
[edit protocols bgp group *group-name*],  
[edit protocols bgp group *group-name* neighbor *address*],  
[edit routing-instances *routing-instance-name* protocols bgp],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
neighbor *address*]

**Description** Log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located under the [edit system syslog] hierarchy.

**Usage Guidelines** See “Configure BGP to Log System Log Messages” on page 412 and the *JUNOS Internet Software Configuration Guide: Getting Started*.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

**See Also** traceoptions on page 443

## metric-out

<b>Syntax</b>	metric-out ( <i>metric</i>   minimum-igp < <i>offset</i> >  igp < <i>offset</i> >);
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	<p>Metric for all routes sent using the multiple exit discriminator (MED, or MULTI_EXIT_DISC) path attribute in update messages. This path attribute is used to discriminate among multiple exit points to a neighboring AS. If all other factors are equal, the exit point with the lowest metric is preferred.</p> <p>You can specify a constant metric value by including the <i>metric</i> option. For configurations in which a BGP peer sends third-party next hops that require the local system to perform next-hop resolution—IBGP configurations, configurations within confederation peers, or EBGP configurations that include the multihop command—you can specify a variable metric by including the minimum-igp or igp option.</p> <p>You can increase or decrease the variable metric calculated from the IGP metric (either from the igp or igp-minimum statement) by specifying a value for &lt;<i>offset</i>&gt;. The metric is increased by specifying a positive value for &lt;<i>offset</i>&gt;, and decreased by specifying a negative value for &lt;<i>offset</i>&gt; .</p>
<b>Options</b>	<p>igp—Set the metric to the most-recent metric value calculated in the IGP to get to the BGP next hop.</p> <p><i>metric</i>—Primary metric on all routes sent to peers. <b>Range:</b> 0 through 4,294,967,295 (<math>2^{32} - 1</math>) <b>Default:</b> No metric is sent.</p> <p>minimum-igp—Set the metric to the minimum metric value calculated in the IGP to get to the BGP next hop. If a newly calculated metric is greater than the minimum metric value, the metric value remains unchanged. If a newly calculated metric is lower, the metric value is lowered to that value.</p> <p><i>offset</i>—(Optional) Increases or decreases the metric by this value. <b>Range:</b> <math>-2^{31}</math> through <math>2^{31} - 1</math> <b>Default:</b> No default.</p>
<b>Usage Guidelines</b>	See “Configure the Multiple Exit Discriminator Metric” on page 387.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## multihop

**Syntax** multihop < *tvl-value* > ;

**Hierarchy Level** [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
 neighbor *address*]

**Description** Configure an EBGp multihop session.

For confederation peerings, you do not need to configure multihop sessions explicitly; multihop behavior is implied.

**Default** If you omit this statement, all EBGp peers are assumed to be directly connected (that is, you are establishing a nonmultihop, or “regular,” BGP session), and the default time-to-live (TTL) value is 1.

**Options** *tvl-value*—Configure the maximum TTL value for the TTL in the IP header of BGP packets.  
**Range:** 1 through 255  
**Default:** 64 (for multihop EBGp sessions, confederations, and internal BGP sessions)

**Usage Guidelines** See “Configure an EBGp Multihop Session” on page 391.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

## multipath

**Syntax** multipath;

**Hierarchy Level** [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
 neighbor *address*]

**Description** Allow load sharing among multiple EBGp paths and multiple IBGP paths.

**Usage Guidelines** See “Configure BGP To Select Multiple BGP Paths” on page 395.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

## neighbor

```

Syntax neighbor address {
    advertise-inactive;
    authentication-key key;
    cluster cluster-identifier;
    damping;
    description text-description;
    export [ policy-names ];
    family (inet | inet6 | inet-vpn | I2-vpn) {
        (any | multicast | unicast) {
            prefix-limit {
                maximum number;
                teardown < percentage> < idle-timeout (forever | time-in-minutes)> ;
            }
            rib-group routing-table-group-name;
        }
        labeled-unicast {
            prefix-limit {
                maximum number;
                teardown < percentage> < idle-timeout (forever | time-in-minutes)> ;
            }
            resolve-vpn;
            rib-group routing-table-group-name;
        }
    }
    hold-time seconds;
    import [ policy-names ];
    ipsec-sa ipsec-sa;
    keep (all | none);
    local-address address;
    local-as autonomous-system <private>;
    local-interface interface-name;
    local-preference preference;
    log-updown;
    metric-out metric;
    multihop < ttl-value>;
    multipath;
    no-aggregator-id;
    no-client-reflect;
    out-delay seconds;
    passive;
    peer-as autonomous-system;
    preference preference;
    traceoptions {
        file name <replace> <size size> <files number> <no-stamp>
            <(world-readable | no-world-readable)>;
        flag flag <flag-modifier> <disable>;
    }
}

```

- **Hierarchy Level** [edit protocols bgp group *group-name*],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*]
- 
- **Description** Explicitly configure a neighbor (peer). To configure multiple BGP peers, include multiple neighbor statements.
- 
- By default, the peer's options are identical to those of the group. You can override these options by including peer-specific option statements within the neighbor statement.
- 
- The neighbor statement is one of the statements you can include in the configuration to define a minimal BGP configuration on the router. (You can include an allow all statement in place of a neighbor statement.)
- 
- **Options** *address*—IPv6 or IPv4 address of a single peer.
- 
- The remaining statements are explained separately in this chapter.
- 
- **Usage Guidelines** See “Minimum BGP Configuration” on page 371 and “Define BGP Groups and Peers” on page 374.
- 
- **Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## no-aggregator-id

- **Syntax** no-aggregator-id;
- 
- **Hierarchy Level** [edit protocols bgp],  
[edit protocols bgp group *group-name*],  
[edit protocols bgp group *group-name* neighbor *address*],  
[edit routing-instances *routing-instance-name* protocols bgp],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*]
- 
- **Description** Set the router ID in the BGP aggregator path attribute to zero. (This is one of the path attributes included in BGP update messages.) Doing this prevents different routers within an AS from creating aggregate routes that contain different AS paths.
- 
- **Default** If you omit this statement, the router ID is included in the BGP aggregator path attribute.
- 
- **Usage Guidelines** See “Update Messages” on page 367 and “Control the Aggregator Path Attribute” on page 390.
- 
- **Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

## no-client-reflect

<b>Syntax</b>	no-client-reflect;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Disable intracluster route redistribution by the system acting as the route reflector. Include this statement when the client cluster is fully meshed to prevent the sending of redundant route advertisements.
<b>Usage Guidelines</b>	See “Configure Route Reflection” on page 399.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	cluster on page 420


## out-delay

<b>Syntax</b>	out-delay <i>seconds</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	How long a route must be present in the JUNOS routing table before it is exported to BGP. Use this time delay to help bundle routing updates.
<b>Default</b>	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
<b>Options</b>	<i>seconds</i> —Output delay time. <b>Range:</b> 0 to 65,535 <b>Default:</b> 0 seconds
<b>Usage Guidelines</b>	See “Configure How Often BGP Exchanges Routes with the Routing Table” on page 410.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## passive

<b>Syntax</b>	passive;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Do not send active open messages to the peer. Rather, wait for the peer to issue an open request.
<b>Default</b>	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
<b>Usage Guidelines</b>	See “Open a Peer Connection Passively” on page 386.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## path-selection

<b>Syntax</b>	path-selection (cisco-non-deterministic   always-compare-med);
<b>Hierarchy Level</b>	[edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
<b>Description</b>	Configures BGP path selection.
<b>Options</b>	cisco-non-deterministic—Configure routing table path selection so that it is performed using the same nondeterministic behavior as the Cisco IOS software. The active path is always first. All nonactive, but eligible, paths follow the active path and are maintained in the order in which they were received, with the most recent path first. Ineligible paths remain at the end of the list.  always-compare-med—Always compare MEDs whether or not the peer ASs of the compared routes are the same.
<div data-bbox="570 1528 1338 1709" style="border: 1px solid black; padding: 10px;">  <p>We recommend that you configure the always-compare-med option.</p> <p><b>Note</b></p> </div>	
<b>Default</b>	If the path-selection statement is not included in the configuration, only the MEDs of routes that have the same peer ASs are compared.
<b>Usage Guidelines</b>	See “Configure Routing Table Path Selection” on page 394.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## peer-as

<b>Syntax</b>	peer-as <i>autonomous-system</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Neighbor (peer) AS number.
<b>Options</b>	<i>autonomous-system</i> —AS number.
<b>Usage Guidelines</b>	See “Define BGP Groups and Peers” on page 374.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## preference

<b>Syntax</b>	preference <i>preference</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	Preference for routes learned from BGP.  At the BGP global level, the preference statement sets the preference for routes learned from BGP. You can override this preference in a BGP group or peer preference statement.  At the group or peer level, the preference statement sets the preference for routes learned from the group or peer. Use this statement to override the preference set in the BGP global preference statement when you want to favor routes from one group or peer over those of another.
<b>Options</b>	<i>preference</i> —Preference to assign to routes learned from BGP or from the group or peer. <b>Range:</b> 0 through 4,294,967,295 ( $2^{32} - 1$ ) <b>Default:</b> 170 for the primary preference
<b>Usage Guidelines</b>	See “Control Route Preference” on page 392.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
<b>See Also</b>	local-preference on page 431

## prefix-limit

**Syntax** prefix-limit {  
     maximum *number*;  
     teardown <*percentage*> <idle-timeout (forever | *timeout-in-minutes*)>;  
 }

**Hierarchy Level** [edit protocols bgp family inet (any | multicast | unicast)],  
 [edit protocols bgp group *group-name* family inet (any | multicast | unicast)],  
 [edit protocols bgp group *group-name* neighbor *address*  
   family inet (any | multicast | unicast)],  
 [edit routing-instances *routing-instance-name* protocols bgp family inet  
   (any | multicast | unicast)],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name* family inet  
   (any | multicast | unicast)],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
   neighbor *address* family inet (any | multicast | unicast)]

**Description** Limit the number of prefixes received on a BGP peering and a rate-limit logging when injected prefixes exceed a set limit.

**Options** maximum *number*—When you set the maximum number of prefixes, a message is logged when that number is reached.

**Range:** 1 through 4,294,967,295

teardown <*percentage*>—If you include the teardown statement, the session is torn down when the maximum number of prefixes is reached. If you specify a percentage, messages are logged when the number of prefixes reaches that percentage. Once the session is torn down, it will reestablish in a short time unless you include the idle-timeout statement. Then the session can be kept down for a specified amount of time, or forever. If you specify forever, the session is reestablished only after you issue a clear bgp neighbor command.

**Range:** 0 through 100

idle-timeout (forever | *timeout-in-minutes*)—If you include the idle-timeout statement, the session is torn down for a specified amount of time, or forever. If you specify a period of time, the session is allowed to reestablish after this timeout period. If you specify forever, the session will be reestablished only after you intervene with a clear bgp neighbor command.

**Range:** 0 through 2400

**Usage Guidelines** See “Enable Multiprotocol BGP” on page 406.

**Required Privilege Level** routing—To view this statement in the configuration.  
 routing-control—To add this statement to the configuration.

## protocol

<b>Syntax</b>	protocol <i>protocol</i> ;
<b>Hierarchy Level</b>	[edit protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ]
<b>Description</b>	Interior gateway protocol (IGP) that BGP should use to resolve the next hop for BGP routes.
<b>Default</b>	If you do not include this statement, BGP uses all active routes when resolving next hops.
<b>Options</b>	<i>protocol</i> —Protocol name. It can be isis or ospf.
<b>Usage Guidelines</b>	See “Choose the Protocol Used to Determine the Next Hop” on page 391.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

## remove-private

<b>Syntax</b>	remove-private;
<b>Hierarchy Level</b>	[edit protocols bgp], [edit protocols bgp group <i>group-name</i> ], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> ], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> ]
<b>Description</b>	When advertising AS paths to remote systems, have the local system strip private AS numbers from the AS path. The numbers are stripped from the AS path starting at the left end of the AS path (the end where AS paths have been most recently added). This operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.  The software recognizes the set of AS numbers that is considered private, a range that is defined in the Internet Assigned Numbers Authority (IANA) assigned numbers document.  The set of reserved AS numbers is in the range 64,512 through 65,534, inclusive.
<b>Usage Guidelines</b>	See “Remove Private AS Numbers from AS Paths” on page 399.
<b>Required Privilege Level</b>	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

- rib-group

- **Syntax** `rib-group group-name;`

- **Hierarchy Level** [edit protocols bgp family inet (any | unicast | multicast)],  
[edit protocols bgp group *group-name* family inet (any | unicast | multicast)],  
[edit protocols bgp group *group-name* neighbor *address*  
family inet (any | unicast | multicast)],  
[edit routing-instances *routing-instance-name* protocols bgp family inet  
(any | unicast | multicast)],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name* family inet  
(any | unicast | multicast)],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
neighbor *address* family inet (any | unicast | multicast)]

- **Description** Add unicast prefixes to unicast and multicast tables.

- **Options** *group-name*—Name of the routing table group. The name must start with a letter and can include letters, numbers, and hyphens. You generally specify only one routing table group.

- **Usage Guidelines** See “Create Routing Table Groups” on page 79, “Configure How Interface Routes Are Imported into Routing Tables” on page 81, and “Configure BGP Routing Table Groups” on page 407.

- **Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

- **See Also** interface-routes on page 108, rib-group on page 120.

## traceoptions

**Syntax** traceoptions {  
     file *name* <replace> <size *size*> <files *number*> <no-stamp>  
         <(world-readable | no-world-readable)>;  
     flag *flag* < *flag-modifier*> < disable> ;  
 }

**Hierarchy Level** [edit protocols bgp],  
 [edit protocols bgp group *group-name*],  
 [edit protocols bgp group *group-name* neighbor *address*],  
 [edit routing-instances *routing-instance-name* protocols bgp],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*],  
 [edit routing-instances *routing-instance-name* protocols bgp group *group-name*  
 neighbor *address*]

**Description** Configure BGP protocol-level tracing options.

To specify more than one tracing operation, include multiple flag statements.

**Default** The default BGP protocol-level tracing options are those inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are those inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are those inherited from the group-level traceoptions statement.

**Options** disable—(Optional) Disable the tracing operation. You can use this option is to disable a single operation when you have defined a broad group of tracing operations, such as all.

file *name*—Name of the file to receive the output of the tracing operation. Enclose the name within quotation marks. All files are placed in the directory /var/log. We recommend that you place BGP tracing output in the file bgp-log.

files *number*—(Optional) Maximum number of trace files. When a trace file named *trace-file* reaches its maximum size, it is renamed *trace-file.0*, then *trace-file.1*, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum number of files, you must also specify a maximum file size with the size option.

**Range:** 2 through 1000 files

**Default:** 2 files

*flag*—Tracing operation to perform. To specify more than one tracing operation, include multiple flag statements.

#### BGP Tracing Flags

as-path—AS path regular expression operations.

damping—Damping operations.

keepalive—BGP keepalive messages.

open—Open packets. These packets are sent between peers when they are establishing a connection.

packets—All BGP protocol packets.

update—Update packets. These packets provide routing updates to BGP systems.

**Global Tracing Flags**

all—All tracing operations.

general—A combination of the normal and route trace operations.

normal—All normal operations.

**Default:** If you do not specify this option, only unusual or abnormal operations are traced.

policy—Policy operations and actions.

route—Routing table changes.

state—State transitions.

task—Interface transactions and processing.

timer—Timer usage.

*flag-modifier*—(Optional) Modifier for the tracing flag. You can specify one or more of these modifiers:

detail—Provide detailed trace information.

filter—Filter trace information. Applies only for route and damping tracing flags.

receive—Packets being received.

send—Packets being transmitted.

no-stamp—(Optional) Do not place timestamp information at the beginning of each line in the trace file.

**Default:** If you omit this option, timestamp information is placed at the beginning of each line of the tracing output.

no-world-readable—(Optional) Disallow any user to read the log file.

replace—(Optional) Replace an existing trace file if there is one.

**Default:** If you do not include this option, tracing output is appended to an existing trace file.

**size** *size*—(Optional) Maximum size of each trace file, in kilobytes (KB), megabytes (MB), or gigabytes (GB). When a trace file named *trace-file* reaches this size, it is renamed *trace-file.0*. When the *trace-file* again reaches its maximum size, *trace-file.0* is renamed *trace-file.1* and *trace-file* is renamed *trace-file.0*. This renaming scheme continues until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.

If you specify a maximum file size, you also must specify a maximum number of trace files with the `files` option.

**Syntax:** *xk* to specify KB, *xm* to specify MB, or *xg* to specify GB

**Range:** 10 KB through the maximum file size supported on your system

**Default:** 1 MB

**world-readable**—(Optional) Allow any user to read the log file.

**Usage Guidelines** See “Trace BGP Protocol Traffic” on page 413.

**Required Privilege Level** routing and trace—To view this statement in the configuration.  
routing-control and trace-control—To add this statement to the configuration.

**See Also** log-updown on page 432

## type

**Syntax** `type type;`

**Hierarchy Level** [edit protocols bgp group *group-name*],  
[edit routing-instances *routing-instance-name* protocols bgp group *group-name*]

**Description** Type of BGP peer group.

**Options** *type*—Type of group:

internal—Internal group

external—External group

**Usage Guidelines** See “Define BGP Groups and Peers” on page 374.

**Required Privilege Level** routing—To view this statement in the configuration.  
routing-control—To add this statement to the configuration.

*type*

.....