

Chapter 31

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

Syntax	advertise-inactive;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Have BGP advertise the best route even if the routing table did not select it to be an active route.
Usage Guidelines	See “Setting BGP to Advertise Inactive Routes” on page 506.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

advertise-peer-as

Syntax	advertise-peer-as;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Disable the default behavior of suppressing AS routes.
Usage Guidelines	See “Suppressing Route Advertisement” on page 507.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

aggregate-label

Syntax	aggregate-label { community <i>community-name</i> ; }
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> protocols bgp family inet-vpn labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp family inet-vpn labeled-unicast]
Description	Enables aggregate labels for VPN traffic.
Options	community <i>community-name</i> —Specify the name of the community to which to apply the aggregate label.
Usage Guidelines	See “Configuring Aggregate Labels for VPNs” on page 480.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

allow

Syntax	<code>allow [network/mask-length];</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
Description	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.
Options	<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 <code>all</code> , which is equivalent to <code>0.0.0.0/0</code> (or <code>::/0</code>).
Usage Guidelines	See “Minimum BGP Configuration” on page 463 and “Defining BGP Groups and Peers” on page 467.
Required Privilege Level	<code>routing</code> —To view this statement in the configuration. <code>routing-control</code> —To add this statement to the configuration.
See Also	<code>neighbor</code> on page 542

as-override

Syntax	as-override;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	<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 may result in routing loops. Use this feature only 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 in a virtual private network. For more information, see the <i>JUNOS MPLS Applications Configuration Guide</i>.</p>
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

authentication-key

Syntax	authentication-key <i>key</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Configure an MD5 authentication key (password). Neighboring routers use the same password to verify the authenticity of BGP packets sent from this system.
Options	<i>key</i> —Authentication password. It can be up to 126 characters. Characters can include any ASCII strings. If you include spaces, enclose all characters in quotation marks (" ").
Usage Guidelines	See “Configuring Authentication” on page 480.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

bgp

Syntax	bgp { ... }
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols], [edit routing-instances <i>routing-instance-name</i> protocols]
Description	Enable BGP on the router or for a routing instance.
Default	BGP is disabled.
Usage Guidelines	See “Enabling BGP” on page 464.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

cluster

Syntax	<code>cluster <i>cluster-identifier</i>;</code>
Hierarchy Level	<p>[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]</p>
Description	Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.
Options	<i>cluster-identifier</i> —IPv6 or IPv4 address to use as the cluster identifier.
Usage Guidelines	See “Configuring Route Reflection” on page 494.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>
See Also	no-client-reflect on page 546

damping

Syntax	damping;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Enable route flap damping.
Default	Flap damping is disabled on the router.
Usage Guidelines	See “Enabling Route Flap Damping” on page 499 and the <i>JUNOS Policy Framework Configuration Guide</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

description

Syntax	description <i>text-description</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Text description of the global, group, or neighbor configuration.
Options	<i>text-description</i> —Text description of the configuration.
Usage Guidelines	See “Defining BGP Global Properties” on page 465, “Defining Group Properties” on page 470, and “Defining Peer Properties” on page 472.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

disable

Syntax	disable;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Description	Disable BGP on the system.
Usage Guidelines	See “Defining BGP Global Properties” on page 465.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

explicit-null

Syntax	explicit-null;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast]
Description	Advertise label 0 to the egress router of an LSP.
Default	If you do not include the explicit-null statement in the configuration, label 3 (implicit null) is advertised.
Usage Guidelines	See “Advertising an Explicit Null Label” on page 479.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

export

Syntax	export [<i>policy-names</i>];
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Apply one or more policies to routes being exported from the routing table into BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Usage Guidelines	See “Configuring BGP Routing Policy” on page 505 and the <i>JUNOS Policy Framework Configuration Guide</i> .
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 530 and the <i>JUNOS Policy Framework Configuration Guide</i> .

family

```

Syntax family {
    (inet | inet6 | inet-vpn | inet6-vpn | l2-vpn) {
        (any | multicast | unicast) {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
            }
            rib-group group-name;
        }
        labeled-unicast {
            aggregate-label {
                community community-name;
            }
            explicit-null {
                connected-only;
            }
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
            }
            resolve-vpn;
            rib inet.3;
            rib-group group-name;
        }
    }
    route-target {
        advertise-default;
        external-paths number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
        }
    }
}

```

Hierarchy Level [edit logical-routers *logical-router-name* protocols bgp],
 [edit logical-routers *logical-router-name* protocols bgp group *group-name*],
 [edit logical-routers *logical-router-name* protocols bgp group *group-name* neighbor *address*],
 [edit logical-routers *logical-router-name* routing-instances *routing-instance-name* protocols bgp],
 [edit logical-routers *logical-router-name* routing-instances *routing-instance-name* protocols bgp group *group-name*],
 [edit logical-routers *logical-router-name* routing-instances *routing-instance-name* protocols bgp group *group-name* neighbor *address*],
 [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 Enable multiprotocol BGP (MBGP) by configuring BGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, to specify MBGP to carry NLRI for the IPv6 address family, or to carry NLRI for VPNs.

Options any—Configure the family type to be both unicast and multicast.

labeled-unicast—Configure the family type to be labeled-unicast. This means that the BGP peers are being used only to carry the unicast routes that are being used by labeled-unicast for resolving the labeled-unicast routes.

multicast—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.

unicast—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.

Default: unicast

The remaining statements are explained separately.

Usage Guidelines See “Enabling Multiprotocol BGP” on page 500.

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

graceful-restart

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

Hierarchy Level [edit logical-routers *logical-router-name* protocols bgp],
[edit logical-routers *logical-router-name* protocols bgp group *group-name*],
[edit logical-routers *logical-router-name* protocols bgp group *group-name*
 neighbor *address*],
[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 the 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 “Configuring Graceful Restart” on page 96 and “Configuring Graceful Restart” on page 478.

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;`
 `[network/mask-length];`
 `authentication-key key;`
 `cluster cluster-identifier;`
 `damping;`
 `description text-description;`
 `export [policy-names];`
 `family {`
 `(inet | inet6 | inet-vpn | inet6-vpn | I2-vpn) {`
 `(any | multicast | unicast) {`
 `prefix-limit {`
 `maximum number;`
 `teardown <percentage> <idle-timeout (forever | time-in-minutes)>;`
 `}`
 `rib-group group-name;`
 `}`
 `labeled-unicast {`
 `explicit-null {`
 `connected-only;`
 `}`
 `prefix-limit {`
 `maximum number;`
 `teardown <percentage> <idle-timeout (forever | time-in-minutes)>;`
 `}`
 `resolve-vpn;`
 `rib inet.3;`
 `rib-group group-name;`
 `}`
 `}`
 `route-target {`
 `advertise-default;`
 `external-paths number;`
 `prefix-limit {`
 `maximum number;`
 `teardown <percentage> <idle-timeout (forever | time-in-minutes)>;`
 `}`
 `}`
 `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 <tth-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;
}
}

```

Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Description	<p>Define a BGP peer group. BGP peer groups share a common type, peer autonomous system (AS) number, and cluster ID, if present. To configure multiple BGP groups, include multiple group statements.</p> <p>By default, the group's options are identical to the global BGP options. To override the global options, include group-specific options within the group statement.</p> <p>The group statement is one of the statements you must include in the configuration to run BGP on the router. See "Minimum BGP Configuration" on page 463.</p>
Options	<p><i>group-name</i>—Name of the BGP group.</p> <p>The remaining statements within the group statement are explained separately.</p>
Usage Guidelines	See "Defining BGP Groups and Peers" on page 467.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

hold-time

Syntax	hold-time <i>seconds</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	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.
Options	<i>seconds</i> —Hold time. Range: 6 through 65,535 seconds Default: 90 seconds
Usage Guidelines	See “Modifying the Hold-Time Value” on page 478.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

import

Syntax	import [<i>policy-names</i>];
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Apply one or more routing policies to routes being imported into the JUNOS routing table from BGP.
Options	<i>policy-names</i> —Name of one or more policies.
Usage Guidelines	See “Configuring BGP Routing Policy” on page 505 and the <i>JUNOS Policy Framework Configuration Guide</i> .
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	export on page 524 and the <i>JUNOS Policy Framework Configuration Guide</i>

include-mp-next-hop

Syntax	include-mp-next-hop;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit protocols bgp]
Description	Enable multiprotocol updates to contain next-hop reachability information.
Usage Guidelines	See “Enabling Next-Hop Reachability Information” on page 511.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

ipsec-sa

Syntax	<code>ipsec-sa ipsec-sa;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	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.
Options	<i>ipsec-sa</i> —Security association name.
Usage Guidelines	See “Applying IPSec Security Association” on page 481.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

keep

Syntax	keep (all none);
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	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.
Default	If you do not include this statement, most routes are retained in the routing table.
Options	all—Retain all routes. none—Retain none of the routes. When keep none 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.
Usage Guidelines	See “Configuring How Often BGP Exchanges Routes with the Routing Table” on page 506.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

labeled-unicast

Syntax	<pre> labeled-unicast { aggregate-label { community community-name; } explicit-null { connected-only; } prefix-limit { maximum <i>number</i>; teardown <percentage> <idle-timeout (forever timeout-in-minutes)>; } resolve-vpn; rib inet.3; rib-group <i>group-name</i>; } </pre>
Hierarchy Level	<pre> [edit logical-routers <i>logical-router-name</i> protocols bgp family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit protocols bgp family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit protocols bgp group <i>group-name</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit routing-instances <i>routing-instance-name</i> protocols bgp family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family (inet inet6 inet-vpn inet6-vpn I2-vpn)] </pre>
Description	Configure the family type to be labeled-unicast.
Options	The statements are explained separately.
Usage Guidelines	See “Enabling Multiprotocol BGP” on page 500.
Required Privilege Level	<pre> routing—To view this statement in the configuration. routing-control—To add this statement to the configuration. </pre>

local-address

Syntax	<code>local-address address;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Specify the 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.
Default	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 Network Interfaces and Class of Service Configuration Guide</i> .
Options	<i>address</i> —IPv6 or IPv4 address of the local end of the connection.
Usage Guidelines	See "Assigning a BGP Identifier" on page 465.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	router-id on page 151

local-as

Syntax	<code>local-as <i>autonomous-system</i> <private>;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Set the local AS number.
Options	<i>autonomous-system</i> —AS number. <i>private</i> —(Optional) Hide the local AS in paths learned from this peering.
Usage Guidelines	See “Configuring a Local AS” on page 490.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

local-interface

Syntax	<code>local-interface <i>interface-name</i>;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>ipv6-link-local-address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>ipv6-link-local-address</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>ipv6-link-local-address</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>ipv6-link-local-address</i>]
Description	Specify the interface name of the peer for IPv6 peering using link-local addresses. This peer is link-local in scope.
Options	<i>interface-name</i> —Interface name of the EBGp IPv6 peer.
Usage Guidelines	See “Configuring EBGp Peering Using IPv6 Link-local Address” on page 508.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

local-preference

Syntax	<code>local-preference local-preference;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	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.
Default	If you omit this statement, the LOCAL_PREF path attribute, if present, is not modified.
Options	<i>local-preference</i> —Preference to assign to routes learned from BGP or from the group or peer. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 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. By default, non-BGP routes that are advertised by BGP are advertised with a LOCAL_PREF value of 100.
Usage Guidelines	See “Configuring the BGP Local Preference” on page 486.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	preference on page 551

log-updown

Syntax	log-updown;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Log a message whenever a BGP peer makes a state transition. Messages are logged using the system logging mechanism located at the [edit system syslog] hierarchy level.
Usage Guidelines	See “Configuring BGP to Log System Log Messages” on page 509 and the <i>JUNOS System Basics Configuration Guide</i> .
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 559

metric-out

Syntax	metric-out (<i>metric</i> minimum-igp < <i>offset</i> > igp < <i>offset</i> >);
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	<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 <<i>offset</i>>. The metric is increased by specifying a positive value for <<i>offset</i>>, and decreased by specifying a negative value for <<i>offset</i>> .</p>
Options	<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. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 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. Range: -2^{31} through $2^{31} - 1$ Default: No default.</p>

Usage Guidelines See “Configuring the Multiple Exit Discriminator Metric” on page 482.

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

mtu-discovery

Syntax mtu-discovery;

Hierarchy Level [edit logical-routers *logical-router-name* protocols bgp],
[edit logical-routers *logical-router-name* protocols bgp group *group-name*],
[edit logical-routers *logical-router-name* protocols bgp group *group-name*
neighbor *address*],
[edit logical-routers *logical-router-name* routing-instances *routing-instance-name*
protocols bgp],
[edit logical-routers *logical-router-name* routing-instances *routing-instance-name*
protocols bgp group *group-name*],
[edit logical-routers *logical-router-name* routing-instances *routing-instance-name*
protocols bgp group *group-name* neighbor *address*],
[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 TCP path MTU discovery. MTU discovery improves convergence times for internal BGP sessions.

Usage Guidelines See “Configuring MTU Discovery” on page 478.

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

multihop

Syntax	<pre>multihop { <ttl-value>; no-nexthop-change; }</pre>
Hierarchy Level	<pre>[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>
Description	<p>Configure an EBGP multihop session.</p> <p>External confederation peering is a special case that allows unconnected third-party next hops. You do not need to configure multihop sessions explicitly in this particular case; multihop behavior is implied.</p> <p>If you have confederation external BGP peer-to-loopback addresses, you still need the multihop configuration.</p>
Default	<p>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.</p>
Options	<p><i>ttl-value</i>—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)</p> <p><i>no-nexthop-change</i>—Specify not to change the TTL value; for next-hop-to-self route advertisements, specify the <i>no-nexthop-self</i> option.</p>
Usage Guidelines	<p>See “Configuring an EBGP Multihop Session” on page 486.</p>
Required Privilege Level	<p><i>routing</i>—To view this statement in the configuration. <i>routing-control</i>—To add this statement to the configuration.</p>

multipath

Syntax	multipath;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Allow load sharing among multiple EBGp paths and multiple IBGP paths.
Usage Guidelines	See “Configuring BGP to Select Multiple BGP Paths” on page 489.
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 | inet6-vpn | l2-vpn) {
            (any | multicast | unicast) {
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
                }
                rib-group routing-table-group-name;
            }
            labeled-unicast {
                explicit-null {
                    connected-only;
                }
                prefix-limit {
                    maximum number;
                    teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
                }
                resolve-vpn;
                rib inet.3;
                rib-group routing-table-group-name;
            }
        }
    }
    route-target {
        advertise-default;
        external-paths number;
        prefix-limit {
            maximum number;
            teardown <percentage> <idle-timeout (forever | time-in-minutes)>;
        }
    }
}
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 logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i>], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>]
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	<i>address</i> —IPv6 or IPv4 address of a single peer. The remaining statements are explained separately.
Usage Guidelines	See “Minimum BGP Configuration” on page 463 and “Defining BGP Groups and Peers” on page 467.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-advertise-peer-as

Syntax	no-advertise-peer-as;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Restore the default behavior of suppressing AS routes.
Usage Guidelines	See “Suppressing Route Advertisement” on page 507.
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 logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
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 460 and “Controlling the Aggregator Path Attribute” on page 485.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

no-client-reflect

Syntax	no-client-reflect;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	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.
Usage Guidelines	See “Configuring Route Reflection” on page 494.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	cluster on page 520


out-delay

Syntax	out-delay <i>seconds</i> ;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Specify 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.
Default	If you omit this statement, routes are exported to BGP immediately after they have been added to the routing table.
Options	<i>seconds</i> —Output delay time. Range: 0 through 65,535 seconds Default: 0 seconds
Usage Guidelines	See “Configuring How Often BGP Exchanges Routes with the Routing Table” on page 506.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

passive

Syntax	passive;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Do not send active open messages to the peer. Rather, wait for the peer to issue an open request.
Default	If you omit this statement, all explicitly configured peers are active, and each peer periodically sends open requests until its peer responds.
Usage Guidelines	See “Opening a Peer Connection Passively” on page 481.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

path-selection

Syntax	path-selection (cisco-non-deterministic always-compare-med);
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit protocols bgp], [edit routing-instances <i>routing-instance-name</i> protocols bgp]
Description	Configure BGP path selection.
Default	If the path-selection statement is not included in the configuration, only the MEDs of routes that have the same peer ASs are compared.
Options	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.
	NOTE: We recommend that you configure the always-compare-med option.
Usage Guidelines	See “Configuring Routing Table Path Selection” on page 488.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

peer-as

Syntax	<code>peer-as <i>autonomous-system</i>;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Specify the neighbor (peer) AS number.
Options	<i>autonomous-system</i> —AS number.
Usage Guidelines	See “Defining BGP Groups and Peers” on page 467.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

preference

Syntax	<code>preference preference;</code>
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>]
Description	Specify the 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.
Options	<i>preference</i> —Preference to assign to routes learned from BGP or from the group or peer. Range: 0 through 4,294,967,295 ($2^{32} - 1$) Default: 170 for the primary preference
Usage Guidelines	See “Controlling Route Preference” on page 487.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.
See Also	local-preference on page 536

prefix-limit

Syntax	<pre>prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>timeout-in-minutes</i>)>; }</pre>
Hierarchy Level	<pre>[edit logical-routers <i>logical-router-name</i> protocols bgp family inet (any labeled-unicast multicast unicast)], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family inet (any labeled-unicast multicast unicast)], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (any labeled-unicast multicast unicast)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet (any labeled-unicast multicast unicast)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (any labeled-unicast multicast unicast)], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (any labeled-unicast multicast unicast)], [edit protocols bgp family inet (any labeled-unicast multicast unicast)], [edit protocols bgp group <i>group-name</i> family inet (any labeled-unicast multicast unicast)], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (any labeled-unicast multicast unicast)], [edit routing-instances <i>routing-instance-name</i> protocols bgp family inet (any labeled-unicast multicast unicast)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet (any labeled-unicast multicast unicast)], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet (any labeled-unicast multicast unicast)]</pre>
Description	Limit the number of prefixes received on a BGP peering and a rate-limit logging when injected prefixes exceed a set limit.
Options	<p>maximum <i>number</i>—When you set the maximum number of prefixes, a message is logged when that number is reached. Range: 1 through 4,294,967,295</p> <p>teardown <<i>percentage</i>>—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</p> <p>idle-timeout (forever <i>timeout-in-minutes</i>)—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 is reestablished only after you intervene with a clear bgp neighbor command. Range: 0 through 2400</p>

Usage Guidelines See “Enabling Multiprotocol BGP” on page 500.

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

protocol

Syntax `protocol protocol;`

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

Description Specify the interior gateway protocol (IGP) that BGP should use to resolve the next hop for BGP routes.

Default If you do not include this statement, BGP uses all active routes when resolving next hops.

Options `protocol`—Protocol name. It can be isis or ospf.

Usage Guidelines See “Choosing the Protocol Used to Determine the Next Hop” on page 486.

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

remove-private

Syntax	remove-private;
Hierarchy Level	<pre>[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>
Description	<p>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). The router stops searching for private ASs when it finds the first non-private AS. This operation takes place after any confederation member ASs have already been removed from the AS path, if applicable.</p> <p>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.</p> <p>The set of reserved AS numbers is in the range from 64,512 through 65,535.</p>
Usage Guidelines	See “Removing Private AS Numbers from AS Paths” on page 493.
Required Privilege Level	<pre>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</pre>

resolve-vpn

Syntax	resolve-vpn;
Hierarchy Level	<p>[edit logical-routers <i>logical-router-name</i> protocols bgp family inet labeled-unicast],</p> <p>[edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast],</p> <p>[edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast],</p> <p>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast],</p> <p>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast],</p> <p>[edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast],</p> <p>[edit protocols bgp family inet labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> family inet labeled-unicast],</p> <p>[edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast],</p> <p>[edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast]</p>
Description	<p>Allow labeled routes to be placed in the inet.3 routing table for route resolution. These routes are then resolved for PE router connections where the remote PE is located across another AS. 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.</p>
Usage Guidelines	See “Enabling Multiprotocol BGP” on page 500.
Required Privilege Level	<p>routing—To view this statement in the configuration.</p> <p>routing-control—To add this statement to the configuration.</p>

rib

Syntax	rib inet.3;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit protocols bgp family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family inet labeled-unicast], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family inet labeled-unicast]
Description	You can allow both labeled and unlabeled routes to be exchanged in a single session. The labeled routes are placed in the inet.3 routing table, and both labeled and unlabeled unicast routes can be sent or received by the router.
Options	inet.3—Name of the routing table.
Usage Guidelines	See “Enabling Multiprotocol BGP” on page 500.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.

rib-group

Syntax	<code>rib-group group-name;</code>
Hierarchy Level	<pre>[edit logical-routers logical-router-name protocols bgp family inet (any labeled-unicast unicast multicast)], [edit logical-routers logical-router-name protocols bgp group group-name family inet (any labeled-unicast unicast multicast)], [edit logical-routers logical-router-name protocols bgp group group-name neighbor address family inet (any labeled-unicast unicast multicast)], [edit logical-routers logical-router-name routing-instances routing-instance-name protocols bgp family inet (any labeled-unicast unicast multicast)], [edit logical-routers logical-router-name routing-instances routing-instance-name protocols bgp group group-name family inet (any labeled-unicast unicast multicast)], [edit logical-routers logical-router-name routing-instances routing-instance-name protocols bgp group group-name neighbor address family inet (any labeled-unicast unicast multicast)], [edit protocols bgp family inet (any labeled-unicast unicast multicast)], [edit protocols bgp group group-name family inet (any labeled-unicast unicast multicast)], [edit protocols bgp group group-name neighbor address family inet (any labeled-unicast unicast multicast)], [edit routing-instances routing-instance-name protocols bgp family inet (any labeled-unicast unicast multicast)], [edit routing-instances routing-instance-name protocols bgp group group-name family inet (any labeled-unicast unicast multicast)], [edit routing-instances routing-instance-name protocols bgp group group-name neighbor address family inet (any labeled-unicast unicast multicast)]</pre>
Description	Add unicast prefixes to unicast and multicast tables.
Options	<code>group-name</code> —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 “Creating Routing Table Groups” on page 88, “Configuring How Interface Routes Are Imported into Routing Tables” on page 89, and “Configuring BGP Routing Table Groups” on page 503.
Required Privilege Level	<pre>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</pre>
See Also	interface-routes on page 132, rib-group on page 149

route-target

Syntax	<pre>route-target { advertise-default; external-paths <i>number</i>; prefix-limit { maximum <i>number</i>; teardown <<i>percentage</i>> <idle-timeout (forever <i>time-in-minutes</i>)>; } }</pre>
Hierarchy Level	<p>[edit logical router <i>logical-router-name</i> protocols bgp family], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> family], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family], [edit protocols bgp family], [edit protocols bgp group <i>group-name</i> family], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i> family], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> family], [edit routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i> family]</p>
Description	<p>Limit the number of prefixes advertised on BGP peerings specifically to the peers that need the updates.</p>
Options	<p>advertise-default—Advertise default routes and suppress more specific routes.</p> <p>external-paths <i>number</i>—Number of external paths accepted for route filtering. Range: 1 through 16 paths Default: 1 path</p> <p>The remaining statements are explained separately.</p>
Usage Guidelines	See “Enabling Route Target Filtering” on page 504.
Required Privilege Level	<p>routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.</p>

traceoptions

Syntax	<pre> traceoptions { file <i>name</i> <replace> <size <i>size</i>> <files <i>number</i>> <no-stamp> <(world-readable no-world-readable)>; flag <i>flag</i> <<i>flag-modifier</i>> <disable>; } </pre>
Hierarchy Level	<pre> [edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> routing-instances <i>routing-instance-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [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>
Description	<p>Configure BGP protocol-level tracing options.</p> <p>To specify more than one tracing operation, include multiple flag statements.</p>
Default	<p>The default BGP protocol-level tracing options are inherited from the routing protocols traceoptions statement included at the [edit routing-options] hierarchy level. The default group-level trace options are inherited from the BGP protocol-level traceoptions statement. The default peer-level trace options are inherited from the group-level traceoptions statement.</p>
Options	<p>disable—(Optional) Disable the tracing operation. You can use this option to disable a single operation when you have defined a broad group of tracing operations, such as all.</p> <p>file <i>name</i>—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.</p> <p>files <i>number</i>—(Optional) Maximum number of trace files. When a trace file named <i>trace-file</i> reaches its maximum size, it is renamed <i>trace-file.0</i>, then <i>trace-file.1</i>, and so on, until the maximum number of trace files is reached. Then, the oldest trace file is overwritten.</p> <p>If you specify a maximum number of files, you must also specify a maximum file size with the size option.</p> <p>Range: 2 through 1000 files Default: 2 files</p>

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 “Tracing BGP Protocol Traffic” on page 511.

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 537

type

Syntax type *type*;

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

Description Specify the type of BGP peer group.

Options *type*—Type of group:

internal—Internal group

external—External group

Usage Guidelines See “Defining BGP Groups and Peers” on page 467.

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

vpn-apply-export

Syntax	vpn-apply-export;
Hierarchy Level	[edit logical-routers <i>logical-router-name</i> protocols bgp], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i>], [edit logical-routers <i>logical-router-name</i> protocols bgp group <i>group-name</i> neighbor <i>address</i>], [edit protocols bgp], [edit protocols bgp group <i>group-name</i>], [edit protocols bgp group <i>group-name</i> neighbor <i>address</i>]
Description	Apply a BGP export policy in addition to a VPN routing and forwarding (VRF) export policy to routes.
Default	The default action is to accept.
Usage Guidelines	See “Applying BGP Export Policy to VRF Routes” on page 510.
Required Privilege Level	routing—To view this statement in the configuration. routing-control—To add this statement to the configuration.