Chapter 21

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

Description Have BGP advertise its routes even if the routing table did not select them to be active routes.

Usage Guidelines See "Have BGP Advertise Inactive Routes" on page 238.

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

allow

Syntax allow [*network*/*mask-length*];

Hierarchy Level [edit protocols bgp group *group-name*]

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 *network*/*mask-length*—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, specify ::/0.

Usage Guidelines See "Define BGP Groups and Peers" on page 220 and "Minimum BGP Configuration" on page 217.

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

See Also neighbor on page 256

as-override

Syntax as-override;

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

Description 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.

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.

Note that enabling the AS override feature might result in routing loops. You should use AS override only for specific applications that require this type of behavior, and in situations with strict network control.

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 *key*;

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

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 *key*—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 (" ").

Usage Guidelines See “Configure Authentication” on page 224.

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 protocols]

Description Enable BGP on the router.

Default BGP is disabled.

Usage Guidelines See “Enable BGP” on page 218.

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

cluster

Syntax cluster *cluster-identifier*;

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

Description Specify the cluster identifier to be used by the route reflector cluster in an internal BGP group.

Options *cluster-identifier*—IPv6 or IPv4 address to use as the cluster identifier.

Usage Guidelines See “Configure Route Reflection” on page 232.

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

See Also no-client-reflect on page 258

damping

Syntax damping;

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

Description Enable route flap damping.

Default Flap damping is disabled on the router.

Usage Guidelines See “Enable Route Flap Damping” on page 233.

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

description

Syntax description *text-description*;

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

Description Text description of the global, group, or neighbor configuration.

Options *text-description*—Text description of the configuration.

Usage Guidelines See "Define BGP Global Properties" on page 219, "Define Group Properties" on page 222, and "Define Peer Properties" on page 223.

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 protocols bgp]

Description Disable BGP on the system.

Usage Guidelines See "Define BGP Global Properties" on page 219.

Required Privilege Level 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*]

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

Options *policy-names*—Name of one or more policies.

Usage Guidelines See "Configure BGP Routing Policy" on page 236.

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 250

family

Syntax

```
family (inet | inet6) {
    (any | multicast | unicast) {
        prefix-limit {
            maximum number;
            teardown < percentage > <idle-timeout (forever | minutes)>;
        }
        rib-group routing-table-group-name;
    }
    labeled-unicast {
        prefix-limit {
            maximum number;
            teardown < percentage > < idle-timeout (forever | minutes) > ;
        }
        resolve-vpn;
        rib-group routing-table-group-name;
    }
}
```

Hierarchy Level [edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit 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, or to specify MBGP to carry NLRI for the IPv6 address family.

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

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 “Enable Multiprotocol BGP” on page 234.

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) {
        (any | multicast | unicast) {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group routing-table-group-name;
        }
        labeled-unicast {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | 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;
    }
}
```

Hierarchy Level [edit protocols bgp]

Description 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.

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.

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 217.

Options *group-name*—Name of the BGP group.

The remaining statements within the group statement are explained separately.

Usage Guidelines See "Define BGP Groups and Peers" on page 220.

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

hold-time

Syntax hold-time *seconds*;

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

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 *seconds*—Hold time. If you set the hold-time value to 0, the hold timer is never started and the router never sends keepalive messages.
Range: 6 through 65,535 seconds
Default: 90 seconds

Usage Guidelines See "Modify the Hold-Time Value" on page 224.

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

import

Syntax import [*policy-names*];

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

Description Apply one or more routing policies to routes being imported into the JUNOS routing table from BGP.

Options *policy-names*—Name of one or more policies.

Usage Guidelines See "Configure BGP Routing Policy" on page 236.

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 246

ipsec-sa

Syntax ipsec-sa *ipsec-sa*;

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 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 *ipsec-sa*—Security association name.

Usage Guidelines See "Apply an IPSec Security Association" on page 225.

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

See Also *JUNOS Internet Software Configuration Guide: Getting Started*

keep

Syntax keep (all | none);

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

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 "Configure How Often BGP Exchanges Routes with the Routing Table" on page 238.

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

local-address

Syntax local-address *address*;

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

Description 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 *JUNOS Internet Software Configuration Guide: Interfaces and Class of Service*.

Options *address*—IPv6 or IPv4 address of the local end of the connection.

Usage Guidelines See "Assign a BGP Identifier" on page 218.

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

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*]

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 231.

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

local-preference

Syntax local-preference *local-preference*;

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

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 do omit this statement, the LOCAL_PREF path attribute, if present, is not modified.

Options *local-preference*—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. A non-BGP route that is advertised by BGP is advertised with a LOCAL_PREF value of 100 by default.

Usage Guidelines See "Configure the BGP Local Preference" on page 229.

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 260

log-updown

Syntax log-updown;

Hierarchy Level [edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit 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 240 and the *JUNOS Internet Software Configur ation Guide: Getting St arted*.

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 263

metric-out

Syntax metric-out (*metric* | minimum-igp <*offset*>| igp <*offset*>);

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

Description 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.

You can specify a constant metric value by including the *metric* 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.

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 <*offset*>. The metric is increased by specifying a positive value for <*offset*>, and decreased by specifying a negative value for <*offset* > .

Options igp—Set the metric to the most-recent metric value calculated in the IGP to get to the BGP next hop.

metric—Primary metric on all routes sent to peers.
Range: 0 through 4,294,967,295 ($2^{32} – 1$)
Default: No metric is sent.

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.

offset—(Optional) Increases or decreases the metric by this value.
Range: $–2^{31}$ through $2^{31} – 1$
Default: No default.

Usage Guidelines See “Configure the Multiple Exit Discriminator Metric” on page 226.

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

multihop

Syntax multihop < *ttl-value* > ;

Hierarchy Level [edit protocols bgp],
[edit protocols bgp group *group-name*],
[edit 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 *ttl-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 228.

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*]

Description Allow loadsharing among multiple EBGP paths and multiple IBGP paths.

Usage Guidelines See "Configure BGP to Select Multiple BGP Paths" on page 230.

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 ) {
        (any | multicast | unicast) {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | minutes)>;
            }
            rib-group routing-table-group-name;
        }
        labeled-unicast {
            prefix-limit {
                maximum number;
                teardown <percentage> <idle-timeout (forever | 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 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*]

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.

Usage Guidelines See "Minimum BGP Configuration" on page 217 and "Define BGP Groups and Peers" on page 220.

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*]

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 213 and "Control the Aggregator Path Attribute" on page 228.

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

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 “Configure Route Reflection” on page 232.

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 245

out-delay

Syntax out-delay *seconds*;

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

Description 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 *seconds*—Output delay time.
Range: 0 to 65,535 seconds
Default: 0 seconds

Usage Guidelines See “Configure How Often BGP Exchanges Routes with the Routing Table” on page 238.

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

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 "Open a Peer Connection Passively" on page 225.

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 protocols bgp]

Description Configures BGP path selection.

Options always-compare-med—Always compare MEDs whether or not the peer ASs of the compared routes are the same.

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.


Note

We recommend that you configure the always-compare-med option.

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.

Usage Guidelines See "Configure Routing Table Path Selection" on page 230.

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

peer-as

Syntax peer-as *autonomous-system*;

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

Description Neighbor (peer) AS number.

Options *autonomous-system*—AS number.

Usage Guidelines See “Define BGP Groups and Peers” on page 220.

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

preference

Syntax preference *preference*;

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

Description 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 *preference*—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 “Control Route Preference” on page 229.

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 252

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)]
```

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

Options

idle-timeout (forever | *minutes*)—If you include the idle-timeout statement, the session is torn down for a specified amount of time or forever. If you specify an amount 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 minutes

maximum *number*—When you set the maximum number of prefixes, log a message 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 of the maximum. Once the session is torn down, it reestablishes within a short period of 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 intervene with a clear bgp neighbor command.
Range: 0 through 100

Usage Guidelines See “Enable Multiprotocol BGP” on page 234.

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 protocols bgp group group-name]
```

Description 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 ripng.

Usage Guidelines See “Choose the Protocol Used to Determine the Next Hop” on page 228.

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

Description 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 Number Authority (IANA) assigned numbers document.

The set of reserved AS numbers ranges from 64,512 through 65,534, inclusive.

Usage Guidelines See "Remove Private AS Numbers from AS Paths" on page 231.

Required Privilege Level 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)]

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 "Configure BGP Routing Table Groups" on page 235.

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

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*]

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

aspath—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.

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: *x*k to specify KB, *x*m to specify MB, or *x*g 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 241.

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 253

type

Syntax type *type*;

Hierarchy Level [edit 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 220.

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

type