Configuration Commands

Generic Commands

shutdown

Syntax [**no**] **shutdown**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command administratively enables an entity.

Unlike other commands and parameters where the default state is not indicated in the configuration file, the **shutdown** and **no shutdown** states are always indicated in system generated configuration files.

Default administrative states for services and service entities are described in Special Cases.

The **no** form of the command places an entity in an administratively enabled state.

Special Cases **BGP Global —** The BGP protocol is created in the **no shutdown** state.

BGP Group — BGP groups are created in the **no shutdown** state.

BGP Neighbor — BGP neighbors/peers are created in the **no shutdown** state.

description

Syntax **description** *description-string*
no description

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command creates a text description stored in the configuration file for a configuration context.

The **no** form of the command removes the description string from the context.

Default No description is associated with the configuration context.

Parameters *string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

BGP Commands

bgp

Syntax [**no**] **bgp**

Context config>router

Description This command creates the BGP protocol instance and BGP configuration context. BGP is administratively enabled upon creation.

The **no** form of the command deletes the BGP protocol instance and removes all configuration parameters for the BGP instance. BGP must be **shutdown** before deleting the BGP instance. An error occurs if BGP is not **shutdown** first.

add-paths

Syntax [**no**] **add-paths**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command allows adds the add-paths node to be the configured for one or more families configuration of the BGP instance, a group or a neighbor. The BGP add-paths capability allows the router to send and/or receive multiple paths per prefix to/from a peer.The add-paths command without additional parameters is equivalent to removing Add-Paths support for all address families, which causes sessions that previously negotiated the add-paths capability for one or more address families to go down and come back up without the add-paths capability.

The no form of the command (no add-paths) removes add-paths from the configuration of BGP, the group or the neighbor, causing sessions established using add-paths to go down and come back up without the add-paths capability.

Default no add-paths

ipv4

Syntax **ipv4 send** *send-limit* **receive** [**none**]
ipv4 send *send-limit*
no ipv4

Context config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths

Description This command is used to configure the add-paths capability for IPv4 routes (including labeled IPv4 routes). By default, add-paths is not enabled for IPv4 routes.

The maximum number of paths per IPv4 prefix to send is the configured send limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default. Entering the command without optional parameters negotiates the ability to both send and receive multiple paths per IPv4 prefix with each peer and configures the router to send the two best paths per prefix to each peer using the default Add-N, N=2 path selection algorithm.

The **no** form of the command disables add-paths support for IPv4 routes, causing sessions established using add-paths for IPv4 to go down and come back up without the add-paths capability.

Default no ipv4

Parameters **send** *send-limit* — The maximum number of paths per IPv4 prefix that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).

Values 1 — 16, none

receive — The router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers

none — The router does not negotiate the Add-Paths receive capability for VPN-IPv64 routes with its peers.

ipv6

Syntax **ipv6 send** *send-limit* **receive** [**none**]
ipv6 send *send-limit*
no ipv6

Context config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths

Description This command is used to configure the add-paths capability for IPv6 routes (including 6PE routes). By default, add-paths is not enabled for IPv6 routes.

The maximum number of paths per IPv6 prefix to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the receive keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of the command disables add-paths support for IPv6 routes, causing sessions established using add-paths for IPv6 to go down and come back up without the add-paths capability.

Default no ipv6

Parameters **send** *send-limit* — The maximum number of paths per IPv6 prefix that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).

Values 1 — 16, none

receive **—** The router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers

none — The router does not negotiate the Add-Paths receive capability for VPN-IPv6 routes with its peers.

vpn-ipv4

Syntax **vpn-ipv4 send** *send-limit* **receive** [**none**]
vpn-ipv4 send *send-limit*
no vpn-ipv4

Context config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths

Description This command is used to configure the add-paths capability for VPN-IPv4 routes. By default, add-paths is not enabled for VPN-IPv4 routes.

The maximum number of paths per VPN-IPv4 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of the command disables add-paths support for VPN-IPv4 routes, causing sessions established using add-paths for VPN-IPv4 to go down and come back up without the add-paths capability.

Default no vpn-ipv4

Parameters *send-limit* — The maximum number of paths per VPN-IPv4 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).

Values 1 — 16, none

receive **—** The router negotiates the add-paths receive capability for VPN-IPv4 routes with its peers

none — The router does not negotiate the Add-Paths receive capability for VPN-IPv64 routes with its peers.

vpn-ipv6

Syntax **vpn-ipv6 send** *send-limit* **receive** [**none**]
vpn-ipv6 send *send-limit*
no vpn-ipv6

Context config>router>bgp>add-paths
config>router>bgp>group>add-paths
config>router>bgp>group>neighbor>add-paths

Description This command is used to configure the add-paths capability for VPN-IPv6 routes. By default, add-paths is not enabled for VPN-IPv6 routes.

The maximum number of paths per VPN-IPv6 NLRI to send is the configured send-limit, which is a mandatory parameter. The capability to receive multiple paths per prefix from a peer is configurable using the **receive** keyword, which is optional. If the **receive** keyword is not included in the command the receive capability is enabled by default.

The **no** form of the command disables add-paths support for VPN-IPv6 routes, causing sessions established using add-paths for VPN-IPv6 to go down and come back up without the add-paths capability.

Default no vpn-ipv6

Parameters *send-limit* — The maximum number of paths per VPN-IPv6 NLRI that are allowed to be advertised to add-paths peers (the actual number of advertised routes may be less depending on the next-hop diversity requirement, other configuration options, route policies and/or route advertisement rules).

Values 1 — 16, none

receive — The router negotiates the add-paths receive capability for VPN-IPv6 routes with its peers

none — The router does not negotiate the add-paths receive capability for VPN-IPv6 routes with its peers.

advertise-external

Syntax [**no**] **advertise-external** [**ipv4**] [**ipv6**]

Context config>router>bgp

Description This command allows BGP to advertise its best external route to a destination even when its best overall route is an internal route. Entering the command (or its no form) with no address family parameters is equivalent to specifying all supported address families.

The no form of the command disables Advertise Best External for the BGP family.

Default no advertise-external

Parameters **ipv4** — Enable/disable best-external advertisement for all IPv4 (unicast and labeled-unicast) routes.

ipv6 — Enable/disable best-external advertisement for all IPv6 (unicast and labeled-unicast) routes.

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables the advertising of inactive BGP routes to other BGP peers. By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

The **no** form of the command disables the advertising of inactive BGP routers to other BGP peers.

Default no advertise-inactive

advertise-label

Syntax **advertise-label** [**ipv4** [**include-ldp-prefix**]] [**ipv6**]
no advertise-label

Context config>router>bgp>group>neighbor

Description This command configures the IPv4 transport peers to exchange IPv6 prefixes using 6PE, LDP FEC prefixes as RFC3107 labeled IPv4, as well as RFC 3107-labeled IPv4 routes.

If IPv4 is enabled all IPv4 routes advertised to the remote BGP peer will be sent with an RFC 3107-formatted label for the destination route. If **include-ldp-fec-prefix** option is also enabled, all activated /32 LDP FEC prefixes will be sent the to remote BGP peer with an RFC 3107 formatted label.

If ipv6 is enabled all IPv6 routes advertised to the remote BGP peer will be sent using the 6PE encapsulation.

The **no** form of the command disables any or all configured options.

The command must include one or more of the options above.

Default no advertise-label

Parameters **ipv4 —** Specifies the advertisement label address family for core IPv4 routes. This keyword can be specified only for an IPv4 peer.

include-ldp-prefix — Specifies the inclusion of LDP FEC prefixes in the advertisement of core IPv4 routes as EFC 3107 labeled routes to the peer.

ipv6 — Specifies the advertisement label address family to support the 6PE feature. This keyword can be specified only for an IPv6 peer.

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. This command is used at the group level to revert to the value defined under the global level, while this command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no aggregator-id-zero —** BGP adds the AS number and router ID to the aggregator path attribute.

aigp

Syntax [**no**] **aigp**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables or disables Accumulated IGP (AIGP) path attribute support with one or more BGP peers. BGP path selection among routes with an associated AIGP metric is based on the end-to-end IGP metrics of the different BGP paths, even when these BGP paths span more than one AS and IGP instance.

The effect of disabling AIGP (using the no form of the command or implicit) is to remove the AIGP attribute from advertised routes, if present, and to ignore the AIGP attribute in received routes.

Default no aigp

always-compare-med

Syntax **always-compare-med** {**zero** | **infinity**}
no always-compare-med strict-as {zero | infinity}
no always-compare-med

Context config>router>bgp>best-path-selection
config>service>vprn>bgp>best-path-selection

Description This command configures the comparison of BGP routes based on the MED attribute. The default behavior of SR-OS (equivalent to the **no** form of the command) is to only compare two routes on the basis of MED if they have the same neighbor AS (the first non-confed AS in the received AS_PATH attribute). Also by default, a route without a MED attribute is handled the same as though it had a MED attribute with the value 0. The **always-compare-med** command without the **strict-as** keyword allows MED to be compared even if the paths have a different neighbor AS; in this case, if neither **zero** or **infinity** is specified, the **zero** option is inferred, meaning a route without a MED is handled the same as though it had a MED attribute with the value 0. When the **strict-as** keyword is present, MED is only compared between paths from the same neighbor AS, and in this case, **zero** or **infinity** is mandatory and tells BGP how to interpret paths without a MED attribute.

Default **no always-compare-med**

Parameters **zero —** Specifies that for routes learned without a MED attribute that a zero (0) value is used in the MED comparison. The routes with the lowest metric are the most preferred.

infinity — Specifies for routes learned without a MED attribute that a value of infinity (2^32-1) is used in the MED comparison. This in effect makes these routes the least desirable.

strict-as — Specifies BGP paths to be compared even with different neighbor AS.

as-path-ignore

Syntax **as-path-ignore** [**ipv4**] [**vpn-ipv4**] [**ipv6**] [**vpn-ipv6**] [**mcast-ipv4**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**l2-vpn**]
no as-path-ignore

Context config>router>bgp>best-path-selection
config>service>vprn>bgp>best-path-selection

Description This command determines whether the AS path is used to determine the best BGP route.

If this option is present, the AS paths of incoming routes are not used in the route selection process.

The **no** form of the command removes the parameter from the configuration.

Default **no as-path-ignore**

Parameters **ipv4 —** Specifies that the AS-path length will be ignored for all IPv4 routes.

vpn-ipv4 — Specifies that the length AS-path will be ignored for all IPv4 VPRN routes.

ipv6 — Specifies that the AS-path length will be ignored for all IPv6 routes.

vpn-ipv6 — Specifies that the AS-path length will be ignored for all IPv6 VPRN routes.

mcast-ipv4 — Specifies that the AS-path length will be ignored for all IPv4 multicast routes.

mvpn-ipv4 — Specifies that the AS-path length will be ignored for all mVPN IPv4 multicast routes.

mvpn-ipv6 — Specifies that the AS-path length will be ignored for all mVPN IPv6 multicast routes.

l2-vpn **—** The AS-path length will be ignored for all L2-VPN NLRIs.

compare-origin-validation-state

Syntax **compare-origin-validation-state**
no compare-origin-validation-state

Context config>router>bgp>best-path-selection

Description When this command is configured, a new step is inserted in the BGP decision process after removal of invalid routes and before the comparison of Local Preference. The new step compares the origin validation state so that a BGP route with a ‘Valid’ state is preferred over a BGP route with a ‘Not-Found’ state, and a BGP route with a ‘Not-Found’ state is preferred over a BGP route with an ‘Invalid’ state assuming that these routes are considered ‘usable’.

The new step is skipped when **no compare-origin-validation-state** is configured.

Default no compare-origin-validation-state

deterministic-med

Syntax [**no**] **deterministic-med**

Context config>router>bgp>best-path-selection

Description This command controls how the BGP decision process compares routes on the basis of MED. When **deterministic-med** is configured, BGP groups paths that are equal up to the MED comparison step based on neighbor AS, and then compares the best path from each group to arrive at the overall best path. This change to the BGP decision process makes best path selection completely deterministic in all cases. Without **deterministic-med**, the overall best path selection is sometimes dependent on the order of the route arrival because of the rule that MED cannot be compared in routes from different neighbor AS.

Default no deterministic-med

auth-keychain

Syntax **auth-keychain** *name*

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures a TCP authentication keychain to use for the session. The keychain allows the rollover of authentication keys during the lifetime of a session.

Default no auth-keychain

Parameters *name* — Specifies the name of the keychain, up to 32 characters, to use for the specified TCP session or sessions.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message based digest.

The authentication *key* can be any combination of ASCII characters up to 255 characters long.

The **no** form of the command reverts to the default value.

Default MD5 Authentication is disabled by default.

Parameters *authentication-key* — The authentication key. The key can be any combination of ASCII characters up to 255 characters in length (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — The hash key. The key can be any combination of ASCII characters up to 342 characters in length (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the **hash2** parameter is not used, the less encrypted **hash** form is assumed.

backup-path

Syntax [**no**] **backup-path** [**ipv4**] [**ipv6**]

Context config>router
config>service>vprn

Description This command enables the computation and use of a backup path for IPv4 and/or IPv6 BGP-learned prefixes belonging to the base router or a particular VPRN. Multiple paths must be received for a prefix in order to take advantage of this feature. When a prefix has a backup path and its primary path(s) fail the affected traffic is rapidly diverted to the backup path without waiting for control plane re-convergence to occur. When many prefixes share the same primary path(s), and in some cases also the same backup path, the time to failover traffic to the backup path is independent of the number of prefixes. In some cases prefix independent convergence may require use of FP2 or later IOMs/IMMs/XMAs.

By default, IPv4 and IPv6 prefixes do not have a backup path installed in the IOM.

Default no backup-path

Parameters **ipv4** — enable the use of a backup path for BGP-learned IPv4 prefixes

ipv6 — enable the use of a backup path for BGP-learned IPv6 prefixes

best-path-selection

Syntax **best-path-selection**

Context config>router>bgp

Description This command enables path selection configuration.

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables the use of bi-directional forwarding (BFD) to control the state of the associated protocol interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command removes BFD from the associated IGP/BGP protocol adjacency.

Default no bfd-enable

cluster

Syntax **cluster** *cluster-id*
no cluster

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of the command deletes the cluster ID and effectively disables the Route Reflection for the given group.

Default **no cluster** — No cluster ID is defined.

Parameters *cluster-id* — The route reflector cluster ID is expressed in dot decimal notation.

Values Any 32 bit number in dot decimal notation. (0.0.0.1 — 255.255.255.255)

confederation

Syntax **confederation** *confed-as-num* **members** *member-as-num*
no confederation *confed-as-num* [**members** *member-as-num*]

Context config>router

Description This command creates confederation autonomous systems within an AS.

This technique is used to reduce the number of IBGP sessions required within an AS. Route reflection is the other technique that is commonly deployed to reduce the number of IBGP sessions.

The **no** form of the command deletes the specified member AS from the confederation.

When members are not specified in the **no** statement, the entire list is removed and confederations is disabled.

When the last member of the list is removed, confederations is disabled.

Default **no confederation** — No confederations are defined.

Parameters *confed-as-num* — The confederation AS number expressed as a decimal integer.

Values 1 — 65535

members *member-as-num* **—** *T*he AS number(s) of members that are part of the confederation expressed as a decimal integer. Configure up to 15 members per *confed-as-num.*

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer. This configuration parameter can be set at three levels: global level (applies to all peers), peer-group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **120**

seconds — The BGP Connect Retry timer value in seconds expressed as a decimal integer.

Values 1 — 65535

damp-peer-oscillations

Syntax **damp-peer-oscillations** [**idle-hold-time** *initial-wait second-wait max-wait*] [**error-interval** *minutes*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command controls how long a BGP peer session remains in the idle-state after some type of error causes the session to reset. In the idle state, BGP does not initiate or respond to attempts to establish a new session. Repeated errors that occur a short while after each session reset cause longer and longer hold times in the idle state. This command supports the DampPeerOscillations FSM behavior described in section 8.1 of RFC 4271, *A Border Gateway Protocol 4 (BGP-4).*

The default behavior, which applies when no damp-peer-oscillations is configured, is to immediately transition out of the idle-state after every reset.

Default *no damp-peer-oscillations*

Parameters *initial-wait* — The amount of time, in minutes, that a session remains in the idle-state after it has been stable for a while.

Values 0 — 2048

Default 0

second-wait — A period of time, in minutes, that is doubled after each repeated session failure that occurs within a relatively short span of time.

Values 1 — 2048

Default 5

max-wait — The maximum amount of time, in minutes, that a session remains in the idle-state after it has experienced repeated instability.

Values 1 — 2048

Default 60

minutes — The interval of time, in minutes after a session reset, during which the session must be error-free in order to reset the penalty counter and return to idle-hold-time to initial-wait.

Values 0 — 2048

Default 30

damping

Syntax [**no**] **damping**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route policy definition.

The **no** form of the command used at the global level reverts route damping.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life:	15 minutes
Max-suppress:	60 minutes
Suppress-threshold:	3000
Reuse-threshold:	750

Default **no damping** — Learned route damping is disabled.

default-route-target

Syntax [**no**] **default-route-target**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command originates the default RTC route (zero prefix length) towards the selected peers.

Default No default RTC routes are advertised by the router.

disable-4byte-asn

Syntax [**no**] **disable-4byte-asn**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer(s).

The **no** form of the command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-capability-negotiation

Syntax [**no**] **disable-capability-negotiation**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the exchange of . When command is enabled and after the peering is flapped, any new are not negotiated and will strictly support IPv4 routing exchanges with that peer.

The **no** form of the command removes this command from the configuration and restores the normal behavior.

Default no disable-capability-negotiation

disable-client-reflect

Syntax [**no**] **disable-client-reflect**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command disables the reflection of routes by the route reflector to the clients in a specific group or neighbor.

This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes.

Default **no disable-client-reflect —** Client routes are reflected to all client peers.

disable-communities

Syntax **disable-communities** [**standard**] [**extended**]
no disable-communities

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures BGP to disable sending communities.

Parameters **standard —** Specifies standard communities that existed before VPRNs or 2547.

extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target.

disable-fast-external-failover

Syntax [**no**] **disable-fast-external-failover**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures BGP fast external failover.

disable-route-table-install

Syntax [**no**] **disable-route-table-install**

Context config>router>bgp

Description This command specifies whether to disable the installation of all (labeled and unlabeled) IPv4 and IPv6 BGP routes into RTM (Routing Table Manager) and the FIB (Forwarding Information Base) on the base router instance.

ebgp-link-bandwidth

Syntax **ebgp-link-bandwidth** *family* [*family* **...** (up to 4 max)]
no ebgp-link-bandwidth

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description When the **egp-link-bandwidth** command is configured, BGP automatically adds a link-bandwidth extended community to every route (of the selected types) received from directly connected (single-hop) EBGP peers within the scope of the command.

The link-bandwidth extended community added by this command encodes the local-AS number of receiving BGP instance and the bandwidth of the interface to the directly connected EBGP peer.

Default no egp-link-bandwidth — No link bandwidth extended community is automatically added to received BGP routes.

Parameters *family* — The BGP address family.

Values		
	ipv4	The command applies to IPv4 and label-IPv4 routes.
	ipv6	The command applies to IPv6 and 6PE routes.
	vpn-ipv4	The command applies to VPN-IPv4 routes.
	vpn-ipv6	The command applies to VPN-IPv6 routes.

enable-origin-validation

Syntax **enable-origin-validation [ipv4] [ipv6]**
no enable-origin-validation

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description When the **enable-origin-validation** command is added to the configuration of a group or neighbor, it causes every inbound IPv4 and/or IPv6 route from that peer to be marked with one of the 3 following origin validation states:

- Valid (0)
- Not-Found (1)
- Invalid (2)

By default (when neither the ipv4 or ipv6 option is present in the command) or when both the ipv4 and ipv6 options are specified, all unicast IPv4 (AFI1/SAFI1), label-IPv4 (AFI1/SAFI4), unicast IPv6 (AFI2/SAFI1), and 6PE (AFI2/SAFI4) routes are evaluated to determine their origin validation states. When only the ipv4 or ipv6 option is present, only the corresponding address family routes (unlabeled and labeled) are evaluated.

The **enable-origin-validation** command applies to all types of BGP peers, but as a general rule, it should only be applied to EBGP peers and groups that contain only EBGP peers.

Default no enable-origin-validation

Parameters **ipv4 —** Enables origin validation processing for IPv4 and label-IPv4 routes.

ipv6 — Enables origin validation processing for IPv6 and 6PE routes.

enable-inter-as-vpn

Syntax [**no**] **enable-inter-as-vpn**

Context config>router>bgp

Description This command specifies whether VPNs can exchange routes across autonomous system boundaries, providing model B connectivity

The **no** form of the command disallows ASBRs to advertise VPRN routes to their peers in other autonomous systems.

Default no enable-inter-as-vpn

enable-peer-tracking

Syntax [**no**] **enable-peer-tracking**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables BGP peer tracking. BGP peer tracking allows a BGP peer to be dropped immediately if the route used to resolve the BGP peer address is removed from the IP routing table and there is no alternative available. The BGP peer will not wait for the holdtimer to expire; therefore, the BGP reconvergance process is accelerated.

The **no** form of the command disables peer tracking.

Default no enable-peer-tracking

enable-rr-vpn-forwarding

Syntax [**no**] **enable-rr-vpn-forwarding**

Context config>router>bgp

Description When this command is configured all received VPN-IP routes, regardless of route target, are imported into the dummy VRF, where the BGP next-hops are resolved. The transport-tunnel command under config>router>bgp determines what types of tunnels are eligible to resolve the next-hops. If a received VPN-IP route from IBGP peer X is resolved and selected as best so that it can be re-advertised to an IBGP peer Y, AND the BGP next-hop is modified towards peer Y (by using the next-hop-self command in Y's group or neighbor context or by using a next-hop action in an export policy applied to Y) then BGP allocates a new VPRN service label value for the route, signals that new label value to Y and programs the IOM to do the corresponding label swap operation. The supported combinations of X and Y are outlined below:

- from X (client) to Y (client)
- from X (client) to Y (non-client)
- from X (non-client) to Y (client)

The no form of the command causes the re-advertisement of a VPN-IP route between one IBGP peer and another IBGP peer does not cause a new VPRN service label value to be signaled and programmed even if the BGP next-hop is changed through group/neighbor configuration or policy.

Note that is it advised to leave this command disabled (for scaling and convergence reasons).

Default no enable-rr-vpn-forwarding

export

Syntax **export** *policy-name* [*policy-name…*]
no export [*policy-name*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the export route policy used to determine which routes are advertised to peers.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of fifteen (15) policy names can be configured. The first policy that matches is applied.

When multiple export commands are issued, the last command entered overrides the previous command.

When no export policies are specified, BGP routes are advertised and non-BGP routes are not advertised by default.

The **no** form of the command removes the policy association with the BGP instance. To remove association of all policies, use the **no export** command without arguments.

Default **no export** — No export policy is specified. BGP routes are advertised and non-BGP routes are not advertised.

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

family

Syntax **family** [**ipv4**] [**vpn-ipv4**] [**ipv6**] [**mcast-ipv4**] [**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**flow-ipv4**] [**flow-ipv6**] [**mdt-safi**] [**ms-pw**] [**route-target**] [**mcast-vpn-ipv4**] [**evpn**]
no family

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the address family or families to be supported over BGP peerings in the base router. This command is additive so issuing the **family** command adds the specified address family to the list.

The **no** form of the command removes the specified address family from the associated BGP peerings. If an address family is not specified, then reset the supported address family back to the default.

Default **ipv4**

Parameters **evpn —** Exchanges Ethernet VPN routes using AFI 25 and SAFI 70.

ipv4 — Provisions support for IPv4 routing information.

vpn-ipv4 — Exchanges IPv4 VPN routing information.

ipv6 — Exchanges IPv6 routing information.

mcast-ipv4 — Exchanges multicast IPv4 routing information.

l2-vpn — Exchanges Layer 2 VPN information.

mvpn-ipv4 — Exchanges Multicast VPN related information.

mvpn-ipv6 — Exchanges Multicast VPN related information.

flow-ipv4 — Exchanges IPv4 flowspec routes belonging to AFI 1 and SAFI 133.

flow-ipv6 — Exchanges IPv6 flowspec routes belonging to AFI 2 and SAFI 133.

mdt-safi — Exchanges Multicast VPN information using MDT-SAFI address family

ms-pw — Exchanges dynamic MS-PW related information.

route-target — Exchanges RT constraint routes for VPN route filtering.

mcast-vpn-ipv4 — –Exchanges Multicast Routes in VPN using SAFI 129.

mcast-ipv6 — –Exchanges multicast IPv6 routing information.

flowspec-validate

Syntax **flowspec-validate**
no flowspec-validate

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables/disables validation of received flowspec routes. A flow route with a destination prefix subcomponent that is received from a particular peer is considered valid if and only if that peer also advertised the best unicast route to the destination prefix and any of its more-specific components. Also, when a flow route is received from an EBGP peer, the left most AS number in the AS_PATH attribute must equal the peer's AS number. If validation is enabled and a flowspec route is not valid, it is not eligible for import into the RIB, it is not used for filtering, a log/trap is generated, and it is not propagated to other flowspec peers.

The **no** form of the command disables the validation procedure.

Default **no flowspec-validate**

route-target-list

Syntax **route-target-list** *comm-id* [*comm-id* ..[up to 15 max]]
no route-target-list [*comm-id*]

Context config>router>bgp

Description This command specifies the route target(s) to be accepted from or advertised to peers. If the **route-target-list** is a non-null list, only routes with one or more of the given route targets are accepted from or advertised to peers.

The **route-target-list** is assigned at the global level and applies to all peers connected to the system.

This command is only applicable if the router is a route-reflector server.

The **no** form of the command with a specified route target community removes the specified community from the **route-target-list**. The **no** form of the command entered without a route target community removes all communities from the list.

Default **no route-target-list**

Parameters *comm-id* — Specifies the route target community in the form <0..65535>:<0..65535>

third-party-nexthop

Syntax third-party-nexthop
no third-party-nexthop

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description Use this command to enable the router to send third-party next-hop to EBGP peers in the same subnet as the source peer, as described in RFC 4271. If enabled when an IPv4 or IPv6 route is received from one EBGP peer and advertised to another EBGP peer in the same IP subnet, the BGP next-hop is left unchanged. Third-party next-hop is not done if the address family of the transport does not match the address family of the route.

The **no** form of the command prevents BGP from performing any third party next-hop processing toward any single-hop EBGP peers within the scope of the command. No third-party next-hop means the next-hop will always carry the IP address of the interface used to establish the TCP connection to the peer.

Default no third-party-nexthop

vpn-apply-export

Syntax [**no**] **vpn-apply-export**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command causes the base instance BGP export route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

The **no** form of the command disables the application of the base instance BGP route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

Default no vpn-apply-export

vpn-apply-import

Syntax [**no**] **vpn-apply-import**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command causes the base instance BGP import route policies to be applied to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

The **no** form of the command disables the application of the base instance BGP import route policies to vpn-ipv4/6, mvpn-ipv4/6, l2-vpn, mdt-safi, mcast-vpn-ipv4, and evpn routes.

Default **no vpn-apply-import**

graceful-restart

Syntax [**no**] **graceful-restart**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description The command enables BGP graceful restart helper procedures (the "receiving router" role defined in the standard) for all received IPv4, IPv6, VPN-IPv4, and VPN-IPv6 routes. In order for helper mode to be available for a particular address family, both peers must signal GR support for the address family during capability negotiation.

When a neighbor covered by GR helper mode restarts its control plane, forwarding can continue uninterrupted while the session is re-established and routes are re-learned.

The **no** form of the command disables graceful restart.

Default no graceful-restart

error-handling

Syntax **error-handling**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies whether updated BGP error handling procedures should be applied.

update-fault-tolerance

Syntax [**no**] **update-fault-tolerance**

Context config>router>bgp>update-error-handling
config>router>bgp>group> update-error-handling
config>router>bgp>group>neighbor> update-error-handling

Description This command enables **treat-as-withdraw** and other similarly non-disruptive approaches for handling a wide range of UPDATE message errors, as long as there are no length errors that prevent all of the NLRI fields from being correctly identified and parsed.

Default no fault-tolerance

enable-notification

Syntax **enable-notification**
no enable-notification

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description When this command is present, the graceful restart capability sent by this router indicates support for NOTIFICATION messages. If the peer also supports this capability then the session can be restarted gracefully (while preserving forwarding) if either peer needs to sends a NOTIFICATION message due to some type of event or error.

Default no enable-notification

restart-time

Syntax **restart-time** *seconds*
no restart-time

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description This command sets the value of the restart-time that is advertised in the router's graceful-restart capability. If this command is not configured.

Default no restart time

Parameters *seconds* — The restart-time that is advertised in the router's graceful-restart capability.

Values 0 — 4095 seconds

Default config>router>bgp>graceful-restart: 120 seconds
config>router>bgp>group>graceful-restart: 300 seconds
config>router>bgp>group>neighbor>graceful-restart: 300 seconds

stale-routes-time

Syntax **stale-routes-time** *time*
no stale-routes-time

Context config>router>bgp>graceful-restart
config>router>bgp>group>graceful-restart
config>router>bgp>group>neighbor>graceful-restart

Description This command configures the maximum amount of time in seconds that stale routes should be maintained after a graceful restart is initiated.

The **no** form of the command resets the stale routes time back to the default of 360 seconds.

Default no restart time

Parameters *time* — Specify the amount of time that stale routes should be maintained after a graceful restart is initiated.

Values 1 — 3600 seconds

group

Syntax [**no**] **group** *name*

Context config>router>bgp

Description This command creates a context to configure a BGP peer group.

The **no** form of the command deletes the specified peer group and all configurations associated with the peer group. The group must be **shutdown** before it can be deleted.

Default No peer groups are defined.

Parameters *name* — The peer group name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

hold-time

Syntax **hold-time** *seconds* [**min** *seconds2*]
no hold-time

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in group) or neighbor level (only applies to specified peer). The most specific value is used.

Even though the implementation allows setting the **keepalive** time separately, the configured **keepalive** timer is overridden by the **hold-time** value under the following circumstances:

1. If the specified hold-time is less than the configured **keepalive** time, then the operational **keepalive** time is set to a third of the **hold-time**; the configured **keepalive** time is not changed.
2. If the **hold-time** is set to zero, then the operational value of the **keepalive** time is set to zero; the configured **keepalive** time is not changed. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **90** seconds

Parameters *seconds* — The hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 — 65535

min *seconds2* **—** The minimum hold-time that will be accepted for the session. If the peer proposes a hold-time lower than this value, the session attempt will be rejected.

ibgp-multipath

Syntax [**no**] **ibgp-multipath**

Context config>router>bgp

Description This command enables IBGP multipath load balancing when adding BGP routes to the route table if the route resolving the BGP nexthop offers multiple nexthops.

The **no** form of the command disables the IBGP multipath load balancing feature.

Default **no ibgp-multipath**

ignore-nh-metric

Syntax **ignore-nh-metric**
no ignore-nh-metric

Context config>router>bgp>best-path-selection
config>service>vprn
config>service>vprn>bgp>best-path-selection

Description This command instructs BGP to disregard the resolved distance to the BGP next-hop in its decision process for selecting the best route to a destination. When configured in the config>router>bgp>best-path-selection context, this command applies to the comparison of two BGP routes with the same NLRI learned from base router BGP peers. When configured in the config>service>vprn context, this command applies to the comparison of two BGP-VPN routes for the same IP prefix imported into the VPRN from the base router BGP instance. When configured in the config>service>vprn>bgp>best-path-selection context, this command applies to the comparison of two BGP routes for the same IP prefix learned from VPRN BGP peers.

The no form of the command (no ignore-nh-metric) restores the default behavior whereby BGP factors distance to the next-hop into its decision process.

Default **no ignore-nh-metric**

ignore-router-id

Syntax **ignore-router-id**
no ignore-router-id

Context config>router>bgp>best-path-selection
config>service>vprn>bgp>best-path-selection

Description When the ignore-router-id command is present and the current best path to a destination was learned from EBGP peer X with BGP identifier x and a new path is received from EBGP peer Y with BGP identifier y the best path remains unchanged if the new path is equivalent to the current best path up to the BGP identifier comparison – even if y is less than x. The no form of the command restores the default behavior of selecting the route with the lowest BGP identifier (y) as best.

Default **no ignore-router-id**

origin-invalid-unusable

Syntax **origin-invalid-unusable**
no origin-invalid-unusable

Context config>router>bgp>best-path-selection

Description When **origin-invalid-unusable** is configured, all routes that have an origin validation state of 'Invalid' are considered unusable by the best path selection algorithm, meaning they are not used for forwarding and not advertised to BGP peers.

With the default of **no origin-invalid-unusable**, routes with an origin validation state of 'Invalid' are compared to other 'usable' routes for the same prefix according to the BGP decision process.

Default **no origin-invalid-unusable**

import

Syntax **import** *policy-name* [*policy-name…*]
no import [*policy-name*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command specifies the import route policy to be used to determine which routes are accepted from peers. Route policies are configured in the **config>router>policy-options** context.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific level is used.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of fifteen (15) policy names can be specified. The first policy that matches is applied.

When multiple **import** commands are issued, the last command entered will override the previous command.

When an import policy is not specified, BGP routes are accepted by default.

The **no** form of the command removes the policy association with the BGP instance. To remove association of all policies, use **no import** without arguments.

Default **no import** — No import policy specified (BGP routes are accepted).

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The **keepalive** parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **keepalive** value is generally one-third of the **hold-time** interval. Even though the implementation allows the **keepalive** value and the **hold-time** interval to be independently set, under the following circumstances, the configured **keepalive** value is overridden by the **hold-time** value:

1. If the specified **keepalive** value is greater than the configured **hold-time**, then the specified value is ignored, and the **keepalive** is set to one third of the current **hold-time** value.
2. If the specified **hold-time** interval is less than the configured **keepalive** value, then the **keepalive** value is reset to one third of the specified **hold-time** interval.
3. If the **hold-time** interval is set to zero, then the configured value of the **keepalive** value is ignored. This means that the connection with the peer is up permanently and no **keepalive** packets are sent to the peer.

The **no** form of the command used at the global level reverts to the default value
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **30** seconds

Parameters *seconds* — The keepalive timer in seconds expressed as a decimal integer.

Values 0 — 21845

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description Configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the router uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of the command removes the configured local-address for BGP.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-address** - The router ID is used when communicating with IBGP peers and the interface address is used for directly connected EBGP peers.

ip-address — The local address expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address.

Values	ipv4-address:	a.b.c.d (host bits must be 0)
	ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 — FFFF]H
		d: [0 — 255]D

local-as

Syntax **local-as** *as-number* [**private**] [**no-prepend-global-as**]
no local-as

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures a BGP local autonomous system (AS) number. In addition to the global AS number configured for BGP using the autonomous-system command, a local AS number can be configured to support various AS number migration scenarios.

When the **local-as** command is applied to a BGP neighbor and the local-as is different from the peer-as, the session comes up as EBGP and by default the global-AS number and then (in that order) the local-as number are prepended to the AS_PATH attribute in outbound routes sent to the peer. In received routes from the EBGP peer, the local AS is prepended to the AS path by default, but this can be disabled with the **private** option.

When the **local-as** command is applied to a BGP neighbor and the local-as is the same as the peer-as, the session comes up as IBGP, and by default, the global-AS number is prepended to the AS_PATH attribute in outbound routes sent to the peer.

This configuration parameter can be set at three levels: global level (applies to all BGP peers), group level (applies to all BGP peers in group) or neighbor level (only applies to one specific BGP neighbor). Thus by specifying this at the neighbor level, it is possible to have a separate **local-as** for each BGP session.

When the optional **no-prepend-global-as** command is configured, the global-as number is not added in outbound routes sent to an IBGP or EBGP peer.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The private option can be added or removed dynamically by reissuing the command. Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-as**

Parameters *as-number* — The virtual autonomous system number expressed as a decimal integer.

Values 1 — 4294967295

private — Specifies the local-as is hidden in paths learned from the peering.

no-prepend-global-as — Specifies that the global-as is hidden in paths announced to the BGP peer.

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute.

This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy. This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no local-preference —** Does not override the local-preference value set in arriving routes and analyze routes without local preference with value of 100.

Parameters *local-preference* — The local preference value to be used as the override value expressed as a decimal integer.

Values 0 — 4294967295

loop-detect

Syntax **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop** | **off**}
no loop-detect

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures how the BGP peer session handles loop detection in the AS path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

Note that dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of the command used at the global level reverts to default, which is **loop-detect ignore-loop**.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **loop-detect ignore-loop**

Parameters **drop-peer —** Sends a notification to the remote peer and drops the session.

discard-route — Discards routes received from a peer with the same AS number as the router itself. This option prevents routes looped back to the router from being added to the routing information base and consuming memory. When this option is changed, the change will not be active for an established peer until the connection is re-established for the peer.

ignore-loop — Ignores routes with loops in the AS path but maintains peering.

off — Disables loop detection.

mdt-safi

Syntax[**[no] mdt-safi**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables peer capability to exchange MDT-SAFI address family advertisements.

med-out

Syntax **med-out** {*number* **| igp-cost**}
no med-out

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set via a route policy.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default where the MED is not advertised.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no med-out**

Parameters *number* — The MED path attribute value expressed as a decimal integer.

Values 0 — 4294967295

igp-cost — The MED is set to the IGP cost of the given IP prefix.

min-route-advertisement

Syntax **min-route-advertisement** *seconds*
no min-route-advertisement

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The **no** form of the command used at the global level reverts to default.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **30** seconds

Parameters *seconds* — The minimum route advertising interval, in seconds, expressed as a decimal integer.

Values 1— 255

mp-bgp-keep

Syntax **[no] mp-bgp-keep**

Context config>router>bgp

Description As a result of enabling this command, route refresh messages are no longer needed, or issued when VPN route policy changes are made; RIB-IN will retain all MP-BGP routes.

The **no** form of the command is used to disable this feature.

multihop

Syntax **multihop** *ttl-value*
no multihop

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

The **no** form of the command is used to convey to the BGP instance that the EBGP peers are directly connected.
The **no** form of the command used at the global level reverts to default.
The **no** form of the command used at the group level reverts to the value defined at the global level.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **1** — EBGP peers are directly connected.

64 — IBGP

Parameters *ttl-value* — The TTL value expressed as a decimal integer.

Values 1 — 255

multipath

Syntax **multipath** *max-paths*
no multipath

Context config>router>bgp

Description This command enables BGP multipath.

When multipath is enabled BGP load shares traffic across multiple links. Multipath can be configured to load share traffic across a maximum of 32 routes. If the equal cost routes available are more than the configured value, then routes with the lowest next-hop IP address value are chosen.

This configuration parameter is set at the global level (applies to all peers).

Multipath is effectively disabled if the value is set to one. When multipath is disabled, and multiple equal cost routes are available, the route with the lowest next-hop IP address will be used.

The **no** form of the command used at the global level reverts to default where **multipath** is disabled.

Default **no multipath**

Parameters *max-paths* — The number of equal cost routes to use for multipath routing. If more equal cost routes exist than the configured value, routes with the lowest next-hop value are chosen. Setting this value to 1 disables multipath.

Values 1 — 16

mvpn-vrf-import-subtype-new

Syntax [**no**] **mvpn-vrf-import-subtype-new**

Context config>router>bgp

Description When enabled, the type/subtype in advertised routes is encoded as 0x010b.

The **no** form of the command (the default) encodes the type/subtype as 0x010a (to preserve backwards compatibility).

next-hop-resolution

Syntax **next-hop-resolution**

Context config>router>bgp

Description This command enables the context to configure next-hop resolution parameters.

label-route-transport-tunnel

Syntax **label-route-transport-tunnel**

Context config>router>bgp>next-hop-res

Description This command enables the context to configure the resolution of RFC 3107 BGP label route prefixes using tunnels to BGP next-hops in TTM.

The **label-route-transport-tunnel** and **family** nodes are simply contexts to configure the binding of IPv4 or IPv6 BGP labeled routes to tunnels.

This command provides a separate control for the different families of RFC 3107 BGP label routes: core IPv4 routes, core IPv6 (6PE), and inter-AS option B vpn-ipv4 and vpn-ipv6 routes at ASBR.

By default, core IPv4 routes and inter-AS option B VPN label routes resolve to LDP without the user needing to enter this command. IPv6 BGP labeled routes routes are currently resolving to IPv4 LDP tunnel only with the 6PE feature and do not require this command.

If the **resolution** option is explicitly set to **disabled**, the default binding to LDP tunnel resumes. If **resolution** is set to **any**, any supported tunnel type in BGP label route context will be selected following TTM preference.

The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

When the **sr-isis** (**sr-ospf**) value is enabled, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF).

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

family

Syntax **family ipv4**

Context config>router>bgp>next-hop-resolution>label-route-transport-tunnel

Description This command configures the address family for configuring the resolution of RFC 3107 BGP label routes using tunnels to BGP peers.

Parameters **ipv4 —** selects the IPv4 address family for configuring the resolution of BGP label routes using tunnels to BGP peers.

resolution

Syntax **resolution** {**any** | **filter** | **disabled**}

Context config>router>bgp>next-hop-resolution>label-route-transport-tunnel>family

Description This command configures the resolution mode in the resolution of BGP label routes using tunnels to BGP peers.

Parameters **any —** enables the binding to any supported tunnel type in BGP label route context following TTM preference.

filter — enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled — disables the resolution of BGP label routes using tunnels to BGP peers.

resolution-filter

Syntax **resolution-filter** [**ldp**] [**rsvp**] [**sr-isis**] [**sr-ospf**]

Context config>router>bgp>next-hop-resolution>label-route-transport-tunnel>family

Description This command configures the subset of tunnel types which can be used in the resolution of BGP label routes using tunnels to BGP peers.

The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).

Parameters **ldp** — selects the LDP tunnel type.

rsvp — selects the RSVP-TE tunnel type.

sr-isis — selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM.

sr-ospf — selects the Segment Routing (SR) tunnel type programmed by an OSPF instance in TTM.

policy

Syntax **policy** *policy-name*
no policy

Context config>router>bgp>next-hop-res

Description This command specifies the name of a policy statement to use with the BGP next-hop resolution process. The policy controls which IP routes in RTM are eligible to resolve the BGP next-hop addresses of IPv4 and IPv6 routes. The policy has no effect on the resolution of BGP next-hops to MPLS tunnels. If a BGP next-hop of an IPv4 or IPv6 route R is resolved in RTM and the longest matching route for the next-hop address is an IP route N that is rejected by the policy then route R is unresolved; if the route N is accepted by the policy then it becomes the resolving route for R.

The default next-hop resolution policy (when the **no policy** command is configured) is to use the longest matching active route in RTM that is not a BGP route (unless **use-bgp-routes** is configured), an aggregate route or a subscriber management route.

Default no policy

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>policy-options** context.

shortcut-tunnel

Syntax **shortcut-tunnel**

Context config>router>bgp>next-hop-res

Description This command enables the context to configure the resolution of BGP prefixes using tunnels to BGP next-hops in TTM.

The **shortcut-tunnel** and **family** nodes are simply contexts to configure the binding of BGP unlabelled routes to tunnels.

The default resolution of a BGP unlabelled route is performed in RTM. The user must configure the **resolution** option to enable resolution to tunnels in TTM. If the **resolution** option is explicitly set to **disabled**, the binding to tunnel is removed and resolution resumes in RTM to IP next-hops.

If **resolution** is set to **any**, any supported tunnel type in BGP shortcut context will be selected following TTM preference. If one or more explicit tunnel types are specified using the **resolution-filter** option, then only these tunnel types will be selected again following the TTM preference.

The following tunnel types are supported in a BGP shortcut context and in order of preference: RSVP, LDP, Segment Routing (SR), and BGP.

The **ldp** value instructs BGP to search for an LDP LSP with a FEC prefix corresponding to the address of the BGP next-hop.

The **bgp** value instructs BGP to search for a BGP LSP with a RFC 107 label route prefix matching the address of the BGP next-hop.

The **rsvp** value instructs BGP to search for the best metric RSVP LSP to the address of the BGP next-hop. This address can correspond to the system interface or to another loopback used by the BGP instance on the remote node. The LSP metric is provided by MPLS in the tunnel table. In the case of multiple RSVP LSPs with the same lowest metric, BGP selects the LSP with the lowest tunnel-id.

When the **sr-isis** (**sr-ospf**) value is enabled, a tunnel to the BGP next-hop is selected in the TTM from the lowest numbered ISIS (OSPF) instance.

The user must set **resolution** to **filter** to activate the list of tunnel-types configured under **resolution-filter**.

If **disallow-igp** is enabled, the BGP route will not be activated using IP next-hops in RTM if no tunnel next-hops are found in TTM.

family

Syntax **family ipv4**

Context config>router>bgp>next-hop-res>shortcut-tunnel

Description This command configures the address family for configuring the resolution of BGP prefixes using tunnels to BGP peers.

Parameters **ipv4 —** selects the IPv4 address family for configuring the resolution of BGP prefixes using tunnels to BGP peers.

resolution

Syntax **resolution** {**any | filter | disabled**}

Context config>router>bgp>next-hop-res>shortcut-tunnel>family

Description This command configures the resolution mode in the resolution of BGP prefixes using tunnels to BGP peers.

Parameters **any —** enables the binding to any supported tunnel type in BGP shortcut context following TTM preference.

filter — enables the binding to the subset of tunnel types configured under **resolution-filter**.

disabled — disables the resolution of BGP prefixes using tunnels to BGP peers.

resolution-filter

Syntax **resolution-filter** [**bgp**] [**ldp**] [**rsvp**] [**sr-isis**] [**sr-ospf**]

Context config>router>bgp>next-hop-res>shortcut-tunnel>family

Description This command configures the subset of tunnel types which can be used in the resolution of BGP label routes using tunnels to BGP peers.

The following tunnel types are supported in a BGP label route context and in order of preference: RSVP, LDP, and Segment Routing (SR).

Parameters **bgp —** selects the BGP label route tunnel type.

ldp — selects the LDP tunnel type.

rsvp — selects the RSVP-TE tunnel type.

sr-isis — selects the Segment Routing (SR) tunnel type programmed by an IS-IS instance in TTM.

sr-ospf — selects the Segment Routing (SR) tunnel type programmed by an OSPF instance in TTM.

peer-tracking-policy

Syntax **peer-tracking-policy** *policy-name*
no peer-tracking-policy

Context config>router>bgp
config>service>vprn>bgp

Description This command specifies the name of a policy statement to use with the BGP peer-tracking function on the BGP sessions where this is enabled. The policy controls which IP routes in RTM are eligible to indicate reachability of IPv4 and IPv6 BGP neighbor addresses. If the longest matching route in RTM for a BGP neighbor address is an IP route that is rejected by the policy, or it is a BGP route accepted by the policy, or if there is no matching route, the neighbor is considered unreachable and BGP tears down the peering session and holds it in the idle state until a valid route is once again available and accepted by the policy.

The default peer-tracking policy (when the no peer-tracking-policy command is configured) is to use the longest matching active route in RTM that is not an LDP shortcut route or an aggregate route.

Default no peer-tracking-policy

Parameters *policy-name* — The route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes. Route policies are configured in the **config>router>pol¬icy-options** context.

use-bgp-routes

Syntax [**no**] **use-bgp-routes**

Context config>router>bgp>next-hop-res

Description This command specifies whether to use BGP routes to resolve BGP nexthop for IPv4 and IPv6 families on this router instance.

Default no use-bgp-routes

outbound-route-filtering

Syntax [**no**] **outbound-route-filtering**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command opens the configuration tree for sending or accepting BGP filter lists from peers (outbound route filtering).

Default no outbound-route-filtering

extended-community

Syntax [**no**] **extended-community**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description The extended-community command opens the configuration tree for sending or accepting extended-community based BGP filters.

In order for the **no** version of the command to work, all sub-commands (**send-orf**, **accept-orf**) must be removed first.

Default Community filtering is not enabled by default.

accept-orf

Syntax [**no**] accept-orf

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command instructs the router to negotiate the receive capability in the BGP ORF negotiation with a peer, and to accept filters that the peer wishes to send.

The **no** form of the command causes the router to remove the accept capability in the BGP ORF negotiation with a peer, and to clear any existing ORF filters that are currently in place.

Default Accepting ORFs is not enabled by default.

send-orf

Syntax **send-orf** [*comm-id*...(up to 32 max)]
no send-orf [*comm-id*]

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command instructs the router to negotiate the send capability in the BGP outbound route filtering (ORF) negotiation with a peer.

This command also causes the router to send a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer as an ORF Action ADD.

The **no** form of this command causes the router to remove the send capability in the BGP ORF negotiation with a peer.

The **no** form also causes the router to send an ORF remove action for a community filter, prefix filter, or AS path filter configured as an inbound filter on the BGP session to its peer.

If the *comm-id* parameter(s) are not exclusively route target communities then the router will extract appropriate route targets and use those. If, for some reason, the *comm-id* parameter(s) specified contain no route targets, then the router will not send an ORF.

Default no send-orf — Sending ORF is not enabled by default.

Parameters *comm-id* — Any community policy which consists exclusively of route target extended communities. If it is not specified, then the ORF policy is automatically generated from configured route target lists, accepted client route target ORFs and locally configured route targets.

neighbor

Syntax [**no**] **neighbor** *ip-address*

Context config>router>bgp>group

Description This command creates a BGP peer/neighbor instance within the context of the BGP group.

This command can be issued repeatedly to create multiple peers and their associated configuration.

The **no** form of the command is used to remove the specified neighbor and the entire configuration associated with the neighbor. The neighbor must be administratively **shutdown** before attempting to delete it. If the neighbor is not shutdown, the command will not result in any action except a warning message on the console indicating that neighbor is still administratively up.

Default No neighbors are defined.

Parameters *ip-address* — The IP address of the BGP peer router in dotted decimal notation.

Values	ipv4-address:	a.b.c.d (host bits must be 0)
	ipv6-address:	x:x:x:x:x:x:x:x[-interface] x:x:x:x:x:x:d.d.d.d[-interface] x: [0 — FFFF]H d: [0 — 255]D interface: 32 characters maximum, mandatory for link local addresses

next-hop-self

Syntax [**no**] **next-hop-self** {[**ipv4**] [**vpn-ipv4**] [**ipv6**] [**mcast-ipv4**] [**l2-vpn**]} [**multihoming** *primary-anycast secondary-anycast*]

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the group or neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.

This is primarily used to avoid third-party route advertisements when connected to a multi-access network.

In addition, this command can be used to enable and configure the multi-homing resiliency mechanism replacing the usual BGP nexthop with a configured anycast address.

The **no** form of the command used at the group level allows third-party route advertisements in a multi-access network.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no next-hop-self** — Third-party route advertisements are allowed.

Parameters **ipv4 —** Provisions support for IPv4 routing information.

vpn-ipv4 — Exchanges IPv4 VPN routing information.

ipv6 — Exchanges IPv6 routing information.

mcast-ipv4 — Exchanges multicast IPv4 routing information.

l2-vpn — Exchanges Layer 2 VPN information.

primary-anycast — Specifies the anycast address that the local node will use to replace the BGP nexthop address in route updates associated peers.

secondary-address — Specifies the anycast address that the local node is to track.

passive

Syntax [**no**] **passive**

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description Enables/disables passive mode for the BGP group or neighbor.

When in passive mode, BGP will not attempt to actively connect to the configured BGP peers but responds only when it receives a connect open request from the peer.

The **no** form of the command used at the group level disables passive mode where BGP actively attempts to connect to its peers.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no passive** — BGP will actively try to connect to all the configured peers.

peer-as

Syntax **peer-as** *as-number*

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

For EBGP peers, the peer AS number configured must be different from the autonomous system number configured for this router under the global level since the peer will be in a different autonomous system than this router

For IBGP peers, the peer AS number must be the same as the autonomous system number of this router configured under the global level.

This is required command for each configured peer. This may be configured under the group level for all neighbors in a particular group.

Default No AS numbers are defined.

Parameters *as-number* — The autonomous system number expressed as a decimal integer.

Values 1 — 4294967295

path-mtu-discovery

Syntax [**no**] **path-mtu-discovery**

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command enables path MTU discovery for the associated TCP connections. In doing so, the MTU for the associated TCP session will be initially set to the egress interface MTU. The DF bit will also be set so that if a router along the path of the TCP connection cannot handle a packet of a particular size without fragmenting, it will send back and ICMP message to set the path MTU for the given session to a lower value that can be forwarded without fragmenting.

The **no** form of the command disables path MTU discovery.

Default no path-mtu-discovery

preference

Syntax [**no**] **preference** *preference*

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the route preference for routes learned from the configured peer(s).

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). The most specific value is used.

The lower the preference the higher the chance of the route being the active route. The router assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of the command used at the global level reverts to default value.

The **no** form of the command used at the group level reverts to the value defined at the global level.

The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **170**

Parameters *preference* — The route preference expressed as a decimal integer.

Values 1 — 255

purge-timer

Syntax **purge-timer** *minutes*
no purge-timer

Context config>router>bgp

Description When the system sends a VPN-IP Route-Refresh to a peer it sets all the VPN-IP routes received from that peer (in the RIB-IN) to stale and starts the purge-timer. If the routes are not updated (refreshed) before the purge-timer has expired then the routes are removed.

The BGP purge timer configures the time before stale routes are purged.

The **no** form of the command reverts to the default.

Default 10

Parameters *minutes* — Specifies the maximum time before stale routes are purged.

Values 1 — 60

rapid-update

Syntax **rapid-update** {[**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**evpn**]}
no rapid-update { [**l2-vpn**] [**mvpn-ipv4**] [**mvpn-ipv6**] [**mdt-safi**] [**evpn**]}

Context config>router>bgp

Description This command enables and disables BGP rapid update for specified address-families. When no parameter is given for the no rapid-update statement, rapid update is disabled for all address-families.

Default no rapid-update

Parameters **l2-vpn —** Specifies the BGP rapid update for the 12-byte Virtual Switch Instance identifier (VSI-ID) value consisting of the 8-byte route distinguisher (RD) followed by a 4-byte value.

mvpn-ipv4 — Specifies BGP rapid update for the mvpn-ipv4 address family. The mvpn-pv4 address is a variable size value consisting of the 1-byte route type, 1-byte length and variable size that is route type specific. Route type defines encoding for the route type specific field. Length indicates the length in octets of the route type specific field.

mdt-safi — Specifies BGP rapid update for the mdt-safi address family. The address is a 16-byte value consisting of 12-byte route distinguisher (RD) followed by a 4-byte group address.

mvpn-ipv6 — Specifies BGP rapid update for the mvpn-ipv6 address family.

evpn — Specifies BGP rapid update for the evpn address family.

rapid-withdrawal

Syntax [**no**] **rapid-withdrawal**

Context config>router>bgp

Description This command disables the delay (Minimum Route Advertisement) on sending BGP withdrawals. Normal route withdrawals may be delayed up to the minimum route advertisement to allow for efficient packing of BGP updates.

The **no** form of the command removes this command from the configuration and returns withdrawal processing to the normal behavior.

Default no rapid-withdrawal

prefix-limit

Syntax **prefix-limit** *family limit* [**log-only**] [**threshold** *percentage*] [**idle-timeout** {*minutes* **| forever**}] [**post-import**]
no prefix-limit *family*

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures the maximum number of BGP routes that can be received from a peer before some administrative action is taken. The administrative action can be the generation of a log event or taking down the session. If a session is taken down, then it can be brought back up automatically after an idle-timeout period, or else it can be configured to stay down ('forever') until the operator performs a reset.

The **prefix-limit** command allows each address family to have its own limit; a set of address family limits can be applied to one neighbor or to all neighbors in a group.

The **no** form of the command removes the **prefix-limit**.

Default No prefix limits for any address family.

Parameters **log-only —** Enables the warning message to be sent at the specified threshold percentage, and also when the limit is reached. However, the BGP session is not taken down.

percentage — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 1 — 100

family — The address family to which the limit applies.

Values ipv4|vpn-ipv4|ipv6|vpn-ipv6|mcast-ipv4|l2-vpn|mvpn-ipv4|mdt-safi|ms-pw|flow-ipv4|route-target|mcast-vpn-ipv4|mvpn-ipv6|flow-ipv6|evpn|mcast-ipv6

limit — The number of routes that can be learned from a peer expressed as a decimal integer.

Values 1 — 4294967295

minutes — Specifies duration in minutes before automatically re-establishing a session.

Values 1 — 1024

forever — Specifies that the session is reestablished only after **clear router bgp** command is executed.

post-import — Specifies that the limit should be applied only to the number of routes that are accepted by import policies.

remove-private

Syntax **remove-private** [**limited**] [**skip-peer-as**]
no remove-private

Context config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

When the **remove-private** parameter is set at the global level, it applies to all peers regardless of group or neighbor configuration. When the parameter is set at the group level, it applies to all peers in the group regardless of the neighbor configuration.

The router software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of the command used at the global level reverts to default value. The **no** form of the command used at the group level reverts to the value defined at the global level. The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Parameters **limited** — This optional keyword removes private ASNs up to the first public ASN encountered. It then stops removing private ASNs.

skip-peer-as — This optional keyword causes this command to not remove a private ASN from the AS-Path if that ASN is the same as the BGP peer AS number.

router-id

Syntax **router-id** *ip-address*
no router-id

Context config>router>bgp

Description This command specifies the router ID to be used with this BGP instance.

Changing the BGP router ID on an active BGP instance causes the BGP instance to restart with the new router ID. The router ID must be set to a valid host address.

It is possible to configure an SR OS node to operate with an IPv6 only BOF and no IPv4 system interface address. When configured in this manner, the operator must explicitly define IPv4 router IDs for protocols such as OSPF and BGP as there is no mechanism to derive the router ID from an IPv6 system interface address.

Default — No router-id is configured for BGP by default. The system interface IP address is used.

Parameters — *ip-address* — The router ID expressed in dotted decimal notation. Allowed value is a valid routable IP address on the router, either an interface or system IP address. It is highly recommended that this address be the system IP address.

split-horizon

Syntax — [**no**] **split-horizon**

Context — config>router>bgp
config>router>bgp>group
config>router>bgp>group>neighbor

Description — This command enables the use of split-horizon. Split-horizon prevents routes from being reflected back to a peer that sends the best route. It applies to routes of all address families and to any type of sending peer; confed-EBGP, EBGP and IBGP.

The configuration default is **no split-horizon**, meaning that no effort is taken to prevent a best route from being reflected back to the sending peer.

Default — **no split-horizon**

transport-tunnel

Syntax — **transport-tunnel** *ldp* | *rsvp-te* | *mpls*

Context — config>router>bgp

Description — This command selects the transport LSP option to provide model B or C connectivity.

The **no** form of the command defaults to LDP as transport LSP method for model B or C connectivity.

Default — transport-tunnel ldp

Parameters — *ldp* — Allows LDP-based LSPs to be used as transport from the ASBR to local PE routers.

rsvp-te — Allows RSVP-TE based LSPs to be used as transport from the ASBR to local PE routers.

mpls — Specifies that both LDP and RSVP-TE can be used to resolve the BGP next-hop for VPRN routes in an associated VPRN instance.

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command configures TTL security parameters for incoming packets. When the feature is enabled, BGP/LDP will accept incoming IP packets from a peer only if the TTL value in the packet is greater than or equal to the minimum TTL value configured for that peer.

The **no** form of the command disables TTL security.

Parameters *min-ttl-value* — Specify the minimum TTL value for an incoming packet.

Values 1 — 255

Default 1

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>router>bgp>group
config>router>bgp>group>neighbor

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the router derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of the command used at the group level reverts to the default value.
The **no** form of the command used at the neighbor level reverts to the value defined at the group level.

Default **no type** — Type of neighbor is derived on the local AS specified.

Parameters **internal —** Configures the peer as internal.

external — Configures the peer as external.

Other BGP-Related Commands

autonomous-system

Syntax **autonomous-system** *autonomous-system-number*
no autonomous-system

Context config>router

Description This command configures the autonomous system (AS) number for the router. A router can only belong to one AS. An AS number is a globally unique number with an AS. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself.

If the AS number is changed on a router with an active BGP instance, the new AS number is not used until the BGP instance is restarted either by administratively disabling/enabling (**shutdown**/**no shutdown**) the BGP instance or rebooting the system with the new configuration.

Default No autonomous system number is defined.

Parameters *autonomous-system-number* — The autonomous system number expressed as a decimal integer.

Values 1 — 4294967295

mh-primary-interface

Syntax **mh-primary-interface** *interface-name*
no mh-primary-interface

Context config>router

Description This command creates a loopback interface for the use in multihoming resilency. Once active this interface can be used to advertise reachability information to the rest of the network using the primary address which is backed up by the secondary

This reachability for this address is advertised via IGPs and LDP protocols to allow the resolution of BGP routes advertised with this address.

The no form of the command disables this setting.

Default **no mh-primary-interface**

Parameters *interface-name* — The name of the primary loopback interface.

mh-secondary-interface

Syntax **mh-secondary-interface** *interface-name*
no mh-secondary-interface

Context config>router

Description This command creates a loopback interface for the use in multihoming resiliency. This address is considered the secondary multihoming address and is only used to resolve routes advertised by the primary router in the event that router becomes unavailable. For this purpose, the Reachability for this address is advertised via IGPs and LDP protocols to allow the resolution of BGP routes advertised with this address by the primary multihoming router.

The no form of the command disables this setting.

Default **no mh-secondary-interface**

Parameters *interface-name* — The name of the secondary loopback interface.

address

Syntax **address** {*ip-address/mask* **|** *ip-address netmask*}
no address

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description This command assigns an IP address, IP subnet, and broadcast address format to an IP interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IP interface for the interface to be active. An IP address and a mask combine to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The local subnet that the address command defines must not be part of the services address space within the routing context by use of the config router service-prefix command. Once a portion of the address space is allocated as a service prefix, that portion is not available to IP interfaces for network core connectivity.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. Show commands display CIDR notation and are stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

The no form of the command removes the IP address assignment from the IP interface. Interface specific configurations for IGP protocols like OSPF are also removed. The no form of this command can only be performed when the IP interface is administratively shut down. Shutting down the IP interface will operationally stop any protocol interfaces or MPLS LSPs that explicitly reference that IP address. When a new IP address is defined, the IP interface can be administratively enabled (no

shutdown), which reinitializes the protocol interfaces and MPLS LSPs associated with that IP interface.

If a new address is entered while another address is still active, the new address will be rejected.

Default **no address**

Parameters *ip-address* — The IP address of the IP interface. The ip-addr portion of the address command specifies the IP host address that will be used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation.

Values 1.0.0.0 — 223.255.255.255

/ — The forward slash is a parameter delimiter that separates the ip-addr portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the ipaddr, the "/" and the mask-length parameter. If a forward slash does not ediately follow the ipaddr, a dotted decimal mask must follow the prefix.

mask-length — The subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the ip-addr from the masklength parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address. Allowed values are integers in the range 1— 32. Note that a mask length of 32 is reserved for system IP addresses.

Values 1— 3

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the ip-addr from a traditional dotted decimal mask. The mask parameter indicates the complete mask that will be used in a logical 'AND' function to derive the local subnet of the IP address. Note that a mask of 255.255.255.255 is reserved for system IP addresses.

Values 128.0.0.0 — 255.255.255.255

net-mask — he subnet mask in dotted decimal notation.

Values 0.0.0.0 — 223.255.255.255 (network bits all 1 and host bits all 0)

description

Syntax **description** *description-string*
no description

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description This command creates a text description stored in the configuration file for a configuration context.

The no form of the command removes the description string from the context.

Default **no description**

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax **shutdown**
no shutdown

Context config>router>mh-primary-interface
config>router>mh-secondary-interface

Description The shutdown command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, shutdown and no shutdown are always indicated in system generated configuration files.

The no form of the command puts an entity into the administratively enabled state.

Default **no shutdown**

hold-time

Syntax **hold-time** *holdover-time*
no hold-time

Context config>router>mh-secondary-interface

Description The optional hold-time parameter is only applicable for the secondary context and specifies how long label information learned about the secondary anycast address should be kept after that peer is declared down. This timer should be set to a value large enough for the remainder of the network to detect the failure and complete the reconvergence process.

The no form of the command resets the hold-time back to the default value.

Default **no hold-time**

Parameters *holdover-time* — (seconds) specifies the number of seconds the router should hold label information learned from the alternate router in it's secondary label table. This is to allow the reset of the network to reconverge after a router failure before the anycast based label assignments are flushed from the forwarding plane.

Values 0—65535

Default 90

router-id

Syntax **router-id** *router-id*
no router-id

Context config>router

Description This command configures the router ID for the router instance.

The router ID is used by both OSPF and BGP routing protocols in this instance of the routing table manager. IS-IS uses the router ID value as its system ID.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is initialized, the new router ID is used. This can result in an interim period of time when different protocols use different router IDs.

To force the new router ID to be used, issue the **shutdown** and **no shutdown** commands for each protocol that uses the router ID, or restart the entire router.

The **no** form of the command to reverts to the default value.

Default The system uses the system interface address (which is also the loopback address).
If a system interface address is not configured, use the last 32 bits of the chassis MAC address.

Parameters *router-id* — The 32 bit router ID expressed in dotted decimal notation or as a decimal value.