Configuring BGP with CLI

This section provides information to configure BGP using the command line interface.

Topics in this section include:

- BGP Configuration Overview on page 700
  - → Preconfiguration Requirements on page 700
  - → BGP Hierarchy on page 700
  - → Internal and External BGP Configurations on page 700
  - → BGP Confederations on page 708
  - → BGP Route Reflectors on page 710
- Basic BGP Configuration on page 702
- Common Configuration Tasks on page 704
  - → Creating an Autonomous System on page 705
  - → Configuring a Router ID on page 706
  - → BGP Components on page 712
  - → Configuring Group Attributes on page 712
  - → Configuring Neighbor Attributes on page 713
  - → Configuring Route Reflection on page 714
  - → Configuring a Confederation on page 715
- BGP Configuration Management Tasks on page 716
  - → Modifying an AS Number on page 716
  - → Modifying the BGP Router ID on page 717
  - → Deleting a Neighbor on page 719
  - → Deleting Groups on page 720

BGP Configuration Overview

Preconfiguration Requirements

Before BGP can be implemented, the following entities must be configured:

- The autonomous system (AS) number for the router.

  An AS number is a globally unique value which associates a router to a specific autonomous system. This number is used to exchange exterior routing information with neighboring ASs and as an identifier of the AS itself. Each router participating in BGP must have an AS number specified.

  In order to implement BGP, the AS number must be specified in the `config>router` context.

- Router ID — The router ID is the IP address of the local router. The router ID identifies a packet's origin. The router ID must be a valid host address.

BGP Hierarchy

BGP is configured in the `config>router>bgp` context. Three hierarchical levels are included in BGP configurations:

- Global level
- Group level
- Neighbor level

Commands and parameters configured on the global level are inherited to the group and neighbor levels although parameters configured on the group and neighbor levels take precedence over global configurations.

Internal and External BGP Configurations

A BGP system is comprised of ASs which share network reachability information. Network reachability information is shared with adjacent BGP peers. BGP supports two types of routing information exchanges:

- External BGP (EBGP) is used between ASs.

EBGP speakers peer to different ASs and typically share a subnet. In an external group, the next hop is dependent upon the interface shared between the external peer and the specific neighbor. The `multihop` command must be specified if an EBGP peer is more than one hop away from the local router.

- Internal BGP (IBGP) is used within an AS.

  IBGP peers belong to the same AS and typically does not share a subnet. Neighbors do not have to be directly connected to each other. Since IBGP peers are not required to be directly connected, IBGP uses the IGP path (the IP next-hop learned from the IGP) to reach an IBGP peer for its peering connection.

Basic BGP Configuration

This section provides information to configure BGP and configuration examples of common configuration tasks. The minimal BGP parameters that need to be configured are:

- An autonomous system number for the router.
- A router ID - Note that if a new or different router ID value is entered in the BGP context, then the new value takes precedence and overwrites the router-level router ID.
- A BGP peer group.
- A BGP neighbor with which to peer.
- A BGP peer-AS that is associated with the above peer.

The BGP configuration commands have three primary configuration levels: **bgp** for global configurations, group **name** for BGP group configuration, and neighbor **ip-address** for BGP neighbor configuration. Within the different levels, many of the configuration commands are repeated. For the repeated commands, the command that is most specific to the neighboring router is in effect, that is, neighbor settings have precedence over group settings which have precedence over BGP global settings.

Following is a sample configuration that includes the above parameters. The other parameters shown below are optional:

```
info
#--------------------------------------------------
echo "IP Configuration"
#--------------------------------------------------
...
        autonomous-system 200
        confederation 300 members 200 400 500 600
        router-id 10.10.10.103
#--------------------------------------------------
...
#--------------------------------------------------
echo "BGP Configuration"
#--------------------------------------------------
        bgp
            graceful-restart
            exit
            cluster 0.0.0.100
            export "direct2bgp"
            router-id 10.0.0.12
            group "To_AS_10000"
                connect-retry 20
                hold-time 90
                keepalive 30
                local-preference 100
                remove-private
                peer-as 10000
                neighbor 10.0.0.8
                    description "To_Router B - EBGP Peer"
```

```
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    local-address 10.0.0.12
                    passive
                    preference 99
                    peer-as 10000
                exit
            exit
            group "To_AS_30000"
                connect-retry 20
                hold-time 90
                keepalive 30
                local-preference 100
                remove-private
                peer-as 30000
                neighbor 10.0.3.10
                    description "To_Router C - EBGP Peer"
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    peer-as 30000
                exit
            exit
            group "To_AS_40000"
                connect-retry 20
                hold-time 30
                keepalive 30
                local-preference 100
                peer-as 65206
                neighbor 10.0.0.15
                    description "To_Router E - Sub Confederation AS 65205"
                    connect-retry 20
                    hold-time 90
                    keepalive 30
                    local-address 10.0.0.12
                    peer-as 65205
                exit
            exit
        exit
#--------------------------------------------------
....
A:ALA-48>config>router#
```

Common Configuration Tasks

This section provides a brief overview of the tasks that must be performed to configure BGP and provides the CLI commands. In order to enable BGP, one AS must be configured and at least one group must be configured which includes neighbor (system or IP address) and peering information (AS number).

All BGP instances must be explicitly created on each router. Once created, BGP is administratively enabled.

Configuration planning is essential to organize ASs and the SRs within the ASs, and determine the internal and external BGP peering.

To configure a basic autonomous system, perform the following tasks:

1. Prepare a plan detailing the autonomous sytem(s), the router belonging to each group, group names, and peering connections.
2. Associate each router with an autonomous system number.
3. Configure each router with a router ID.
4. Associate each router with a peer group name.
5. Specify the local IP address that will be used by the group or neighbor when communicating with BGP peers.
6. Specify neighbors.
7. Specify the autonomous system number associated with each neighbor.

Creating an Autonomous System

Before BGP can be configured, the autonomous system must be configured first. In BGP, routing reachability information is exchanged between autonomous systems (ASs). An AS is a group of networks that share routing information. The **autonomous-system** command associates an autonomous system number to the router being configured. The `autonomous-system` command is configured in the **config>router** context.

Use the following CLI syntax to associate a router to an autonomous system:

CLI Syntax: `config>router# autonomous-system` *`autonomous-system`*

The router series supports 4 bytes AS numbers by default. This means autonomous-system can have any value from 1 to 4294967295. The following example displays autonomous system configuration command usage:

Example: `config>router# autonomous-system 100`

The following example displays the autonomous system configuration:

```
ALA-B>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.104/24
            port 1/1/1
        exit
        autonomous-system 100

#------------------------------------------
ALA-B>config>router#
```

Configuring a Router ID

In BGP, routing information is exchanged between autonomous systems. The BGP router ID, expressed like an IPv4 address, uniquely identifies the router. It can be set to be the same as the system interface address.

Note that 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.

Note that if a new or different router ID value is entered in the BGP context, then the new router ID value is used instead of the router ID configured on the router level, system interface level, or inherited from the MAC address. The router-level router ID value remains intact. The router ID used by BGP is selected in the following order:

- The routed-id configured under **config>router>bgp**
- The router-id configured under **config>router**
- The system interface IPv4 address
- The last 4 bytes of the system MAC address

When configuring a new router ID outside of the **config>router>bgp** context, BGP is not automatically restarted with the new router ID; the next time BGP is (re) initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the `shutdown` and `no shutdown` commands for BGP or restart the entire router. Use the following CLI syntax to configure the router ID for multiple protocols:

CLI Syntax: `config>router# router-id` *`router-id`*

The following example displays router ID configuration command usage:

Example: `config>router# router-id` **`10.10.10.104`**

The following example displays the router ID configuration:

```
ALA-B>config>router# info
----------------------------------------------
# IP Configuration
#------------------------------------------
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.104/24
            port 1/1/1
        exit
        autonomous-system 100
        router-id 10.10.10.104
```

```
#--------------------------------------------
...
ALA-B>config>router#
```

BGP Confederations

Follow these steps to configure a confederation:

1. Configure the autonomous system number of the confederation using the confederation command in the **config>router** context..
2. Configure the BGP confederation members using the `confederation` command in the `config>router` context.
3. Configure IBGP peering within the (local) sub-confederation.
4. Configure one or more confed-EBGP peerings to peers in other neighboring sub-confederations.

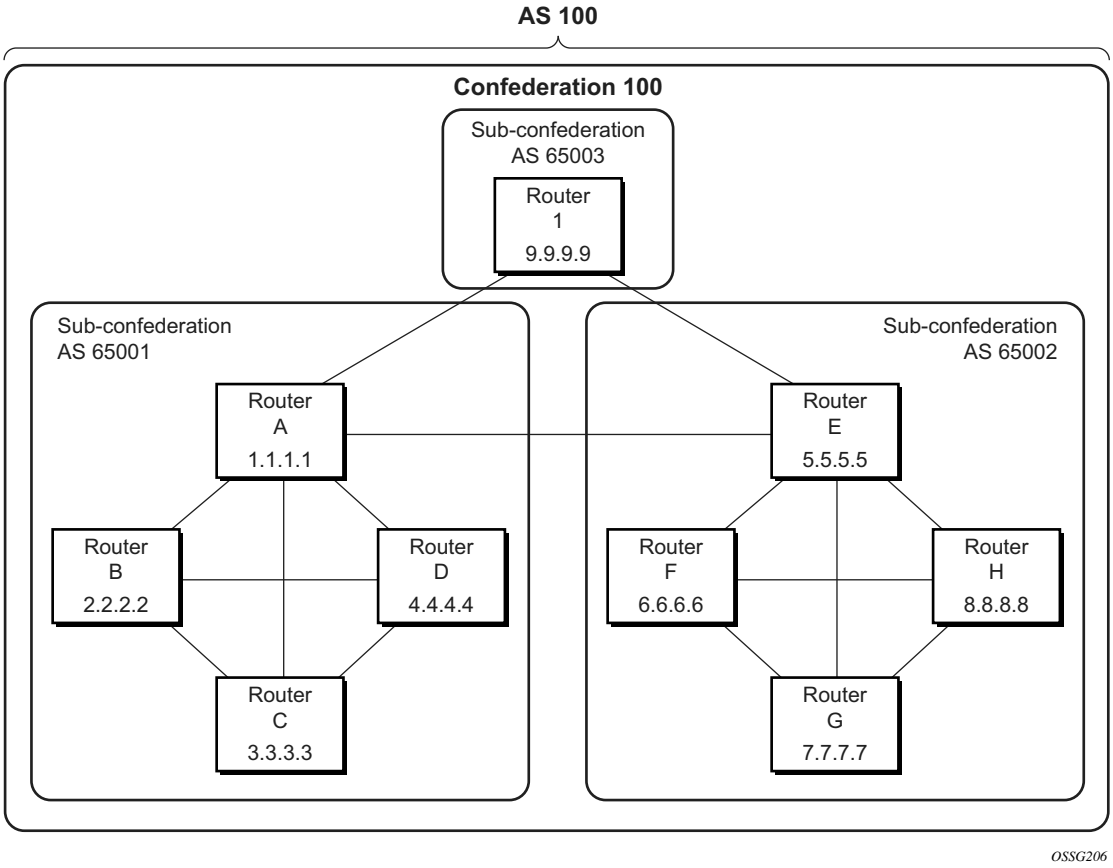

Figure 30: Confederation Network Diagram Example

The following configuration displays the minimum BGP configuration for routers in sub-confederation AS 65001 outlined in Figure 31.

```
ALA-A
     config router
          autonomous-system 65001
          confederation 100 members 65001 65002 65003
          bgp
               group confed1
                    peer-as 65001
                    neighbor 2.2.2.2
                    exit
                    neighbor 3.3.3.3
                    exit
                    neighbor 4.4.4.4
                    exit
               exit
               group external_confed
                    neighbor 5.5.5.5
                         peer-as 65002
                    exit
                    neighbor 9.9.9.9
                         peer-as 65003
                    exit
               exit
          exit
     exit

ALA-D
     config router
          autonomous-system 65001
          confederation 100 members 65001 65002 65003
          bgp
               group confed1
                    peer-as 65001
                    neighbor 1.1.1.1
                    exit
                    neighbor 2.2.2.2
                    exit
                    neighbor 3.3.3.3
                    exit
               exit
          exit
     exit

ROUTER 1
     config router
          autonomous-system 65003
          confederation 100 members 65001 65002 65003
          bgp
               group confed1
                    peer-as 65001
                    neighbor 1.1.1.1
                    exit
                    neighbor 5.5.5.5
                         peer-as 65002
                    exit
               exit
          exit
     exit
```

BGP Route Reflectors

In a standard BGP configuration, all BGP speakers within an AS must have a full BGP mesh to ensure that all externally learned routes are redistributed through the entire AS. IBGP speakers do not re-advertise routes learned from one IBGP peer to another IBGP peer. If a network grows, scaling issues could emerge because of the full mesh configuration requirement. Route reflection circumvents the full mesh requirement but still maintains the full distribution of external routing information within an AS.

Autonomous systems using route reflection arrange BGP routers into groups called *clusters*. Each cluster contains at least one route reflector which is responsible for redistributing route updates to all clients. Route reflector clients do not need to maintain a full peering mesh between each other. They only require a peering to the route reflector(s) in their cluster. The route reflectors must maintain a full peering mesh between all non-clients within the AS.

Each route reflector must be assigned a cluster ID and specify which neighbors are clients and which are non-clients to determine which neighbors should receive reflected routes and which should be treated as a standard IBGP peer. Additional configuration is not required for the route reflector besides the typical BGP neighbor parameters.

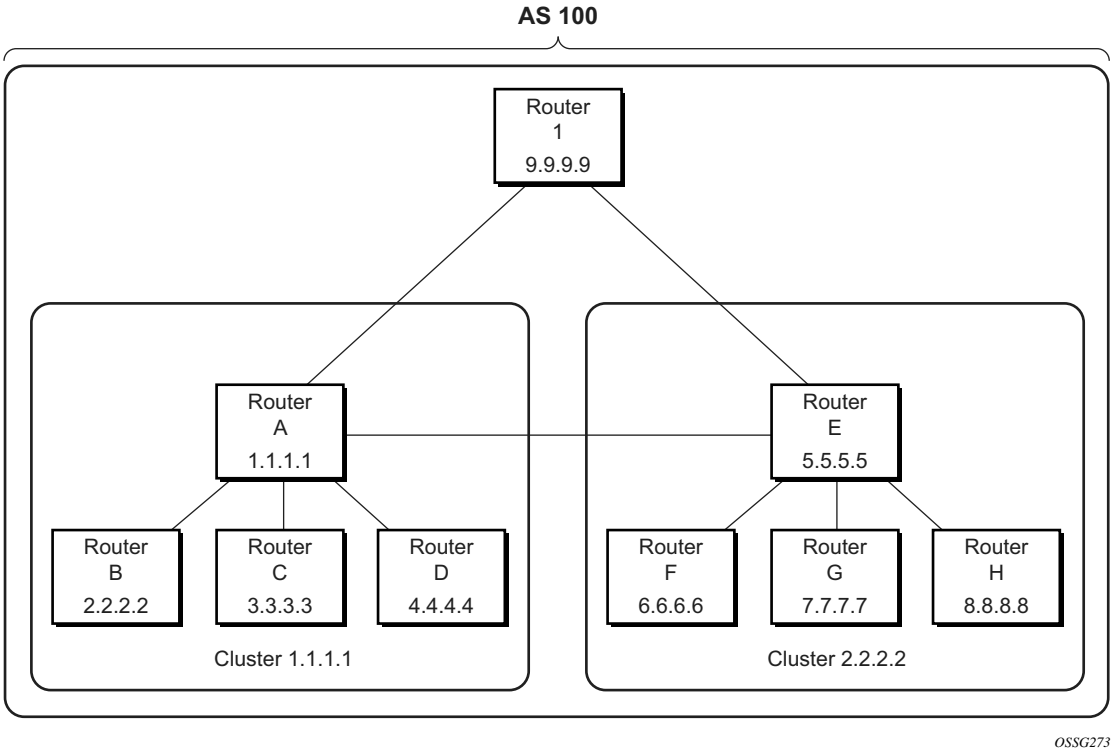

Figure 31: Route Reflection Network Diagram Example

The following configuration displays the minimum BGP configuration for routers in Cluster 1.1.1.1 outlined in Figure 31.

```
ALA-A
     config router bgp
          group cluster1
               peer-as 100
               cluster 1.1.1.1
               neighbor 2.2.2.2
               exit
               neighbor 3.3.3.3
               exit
               neighbor 4.4.4.4
               exit
          exit
          group RRs
               peer-as 100
               neighbor 5.5.5.5
               exit
               neighbor 9.9.9.9
               exit
          exit
     exit

ALA-B
     config router bgp
          group cluster1
               peer-as 100
               neighbor 1.1.1.1
               exit
          exit
     exit

ALA-C
     config router bgp
          group cluster1
               peer-as 100
               neighbor 1.1.1.1
               exit
          exit
     exit

ALA-D
     config router bgp
          group cluster1
               peer-as 100
               neighbor 1.1.1.1
               exit
          exit
     exit
```

BGP Components

Use the CLI syntax displayed below to configure the following BGP attributes:

- BGP Components on page 712
- Configuring Group Attributes on page 712
- Configuring Neighbor Attributes on page 713
- Configuring Route Reflection on page 714
- Configuring a Confederation on page 715

Configuring Group Attributes

A group is a collection of related BGP peers. The group name should be a descriptive name for the group. Follow your group, name, and ID naming conventions for consistency and to help when troubleshooting faults.

All parameters configured for a peer group are applied to the group and are inherited by each peer (neighbor), but a group parameter can be overridden on a specific neighbor-level basis.

The following example displays the BGP group configuration:

```
ALA-B>config>router>bgp# info
----------------------------------------------
...
            group "headquarters1"
                description "HQ execs"
                local-address 10.0.0.104
                disable-communities standard extended
                ttl-security 255
                exit
            exit
...
----------------------------------------------
ALA-B>config>router>bgp#
```

Configuring Neighbor Attributes

After you create a group name and assign options, add neighbors within the same autonomous system to create IBGP connections and/or neighbors in different autonomous systems to create EBGP peers. All parameters configured for the peer group level are applied to each neighbor, but a group parameter can be overridden on a specific neighbor basis.

The following example displays neighbors configured in group “headquarters1”.

```
ALA-B>config>router>bgp# info
----------------------------------------------
...
            group "headquarters1"
                description "HQ execs"
                local-address 10.0.0.104
                disable-communities standard extended
                ttl-security 255
                neighbor 10.0.0.5
                    passive
                    peer-as 300
                exit
                neighbor 10.0.0.106
                    peer-as 100
                exit
                neighbor 17.5.0.2
                    hold-time 90
                    keepalive 30
                    min-as-origination 15
                    local-preference 170
                    peer-as 10701
                exit
                neighbor 17.5.1.2
                    hold-time 90
                    keepalive 30
                    min-as-origination 15
                    local-preference 100
                    min-route-advertisement 30
                    preference 170
                    peer-as 10702
                exit
            exit
...
----------------------------------------------
ALA-B>config>router>bgp#
```

Configuring Route Reflection

Route reflection can be implemented in autonomous systems with a large internal BGP mesh to reduce the number of IBGP sessions required. One or more routers can be selected to act as focal points for internal BGP sessions. Several BGP speaking routers can peer with a route reflector. A route reflector forms peer connections to other route reflectors. A router assumes the role as a route reflector by configuring the `cluster cluster-id` command. No other command is required unless you want to disable reflection to specific peers.

If you configure the `cluster` command at the global level, then all subordinate groups and neighbors are members of the cluster. The route reflector cluster ID is expressed in dotted decimal notation. The ID should be a significant topology-specific value. No other command is required unless you want to disable reflection to specific peers.

If a route reflector client is fully meshed, the `disable-client-reflect` command can be enabled to stop the route reflector from reflecting redundant route updates to a client.

The following example displays a route reflection configuration:

```
ALA-B>config>router>bgp# info
----------------------------------------------
            cluster 0.0.0.100
            group "Santa Clara"
                local-address 10.0.0.103
                neighbor 10.0.0.91
                    peer-as 100
                exit
                neighbor 10.0.0.92
                    peer-as 100
                exit
                neighbor 10.0.0.93
                    disable-client-reflect
                    peer-as 100
                exit
            exit
----------------------------------------------
ALA-B>config>router>bgp#
```

Configuring a Confederation

Reducing a complicated IBGP mesh can be accomplished by dividing a large autonomous system into smaller autonomous systems. The smaller ASs can be grouped into a confederation. A confederation looks like a single AS to routers outside the confederation. Each confederation is identified by its own (confederation) AS number.

To configure a BGP confederation, you must specify a confederation identifier, an AS number expressed as a decimal integer. The collection of autonomous systems appears as a single autonomous system with the confederation number acting as the "all-inclusive" autonomous system number. Up to 15 members (ASs) can be added to a confederation.

NOTE: The `confederation` command is configured in the **config>router** context.

Use the following CLI syntax to configure a confederation:

CLI Syntax: `config>router# confederation` *`confed-as-num`* `members` *`member-as-num`*

When 4-byte AS number support is not disabled on router, the confederation and any of its members can be assigned an AS number in the range from 1 to 4294967295. The following example displays a confederation configuration command usage:

Example: `config>router>#  confederation 1000 members 100 200 300`

The following example displays the confederation configuration:

```
ALA-B>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
        interface "system"
            address 10.10.10.103/32
        exit
        interface "to-104"
            shutdown
            address 10.0.0.103/24
            port 1/1/1
        exit
        autonomous-system 100
        confederation 1000 members 100 200 300
        router-id 10.10.10.103
#------------------------------------------
ALA-B>config>router#
```

BGP Configuration Management Tasks

This section discusses the following BGP configuration management tasks:

- Modifying an AS Number on page 716
- Modifying a Confederation Number on page 717
- Modifying the BGP Router ID on page 717
- Modifying the Router-Level Router ID on page 718
- Deleting a Neighbor on page 719
- Deleting Groups on page 720

Modifying an AS Number

You can modify an AS number on a router but the new AS number will not be used until the BGP instance is restarted either by administratively disabling or enabling the BGP instance or by rebooting the system with the new configuration.

Since the AS number is defined in the **config>router** context, not in the BGP configuration context, the BGP instance is not aware of the change. Re-examine the plan detailing the autonomous sytem(s), the SRs belonging to each group, group names, and peering connections. Changing an AS number on a router could cause configuration inconsistencies if associated **peer-as** values are not also modified as required. At the group and neighbor levels, BGP will re-establish the peer relationships with all peers in the group with the new AS number.

Use the following CLI syntax to change an autonomous system number:

CLI Syntax: `config>router# autonomous-system` *`autonomous-system`*

CLI Syntax:
```
config>router# bgp
 group name
    neighbor ip-addr
          peer-as asn
```

Example:
```
config>router# autonomous-system 400
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.10.10.103
config>router>bgp>group# peer-as 400
config>router>bgp>group# exit
```

Modifying a Confederation Number

Modifying a confederation number will cause BGP to restart automatically. Changes immediately take effect.

Modifying the BGP Router ID

Changing the router ID number in the BGP context causes the new value to overwrite the router ID configured on the router level, system interface level, or the value inherited from the MAC address. It triggers an immediate reset of all peering sessions.

Example: `config>router>bgp# router-id` 10.0.0.123

This example displays the BGP configuration with the BGP router ID specified:

```
ALA-B>config>router>bgp# info detail
----------------------------------------------
     no shutdown
     no description
     no always-compare-med
     ibgp-multipath
     . . .
     router-id 10.0.0.123
----------------------------------------------
ALA-B>config>router>bgp#
```

Modifying the Router-Level Router ID

Changing the router ID number in the `config>router` context causes the new value to overwrite the router ID derive from the system interface address, or the value inherited from the MAC address.

When configuring a new router ID, protocols are not automatically restarted with the new router ID. The next time a protocol is (re) initialized the new router ID is used. An interim period of time can occur when different protocols use different router IDs. To force the new router ID, issue the `shutdown` and `no shutdown` commands for each protocol that uses the router ID or restart the entire router.

Use the following CLI syntax to change a router ID:

CLI Syntax: `config>router# router-id` *`router-id`*

Example:
```
config>router# router-id 10.10.10.104
config>router# no shutdown
config>router>bgp# shutdown
config>router>bgp# no shutdown
```

The following example displays the router ID configuration:

```
ALA-A>config>router# info
#------------------------------------------
# IP Configuration
#------------------------------------------
        interface "system"
            address 10.10.10.104/32
        exit
        interface "to-103"
            address 10.0.0.104/24
            port 1/1/1
        exit
        autonomous-system 100
        router-id 10.10.10.104

#------------------------------------------
ALA-B>config>router#
```

Deleting a Neighbor

In order to delete a neighbor, you must shut down the neighbor before issuing the `no neighbor ip-addr` command.

Use the following CLI syntax to delete a neighbor:

CLI Syntax:
```
config>router# bgp
 group name
    no neighbor ip-address
    shutdown
       no peer-as asn
       shutdown
```

Example:
```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# no neighbor 10.0.0.103
```

The following example displays the "`headquarters1`" configuration with the `neighbor 10.0.0.103` removed.

```
ALA-B>config>router>bgp# info
----------------------------------------------
            group "headquarters1"
              description "HQ execs"
                local-address 10.0.0.104
                neighbor 10.0.0.5
                    passive
                    peer-as 300
                exit
          exit
----------------------------------------------
ALA-B>config>router>bgp#
```

Deleting Groups

In order to delete a group, the neighbor configurations must be shut down first. After each neighbor is shut down, you must shut down the group before issuing the `no group name` command.

Use the following CLI syntax to shut down a peer and neighbor and then delete a group:

CLI Syntax:
```
config>router# bgp
 no group name
 shutdown
    no neighbor ip-address
    shutdown
       shutdown
```

Example:
```
config>router# bgp
config>router>bgp# group headquarters1
config>router>bgp>group# neighbor 10.0.0.105
config>router>bgp>group>neighbor# shutdown
config>router>bgp>group>neighbor# exit
config>router>bgp>group# neighbor 10.0.0.103
config>router>bgp>group# shutdown
config>router>bgp>group# exit
config>router>bgp# no group headquarters1
```

If you try to delete the group without shutting down the peer-group, the following message appears:

```
ALA-B>config>router>bgp# no group headquarters1
MINOR: CLI BGP Peer Group should be shutdown before deleted. BGP Peer
Group not deleted.
```