



BGP

This chapter describes how to configure the ASA to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

This chapter includes the following sections:

- [Information About BGP, page 29-1](#)
- [Licensing Requirements for BGP, page 29-3](#)
- [Guidelines and Limitations, page 29-3](#)
- [Configuring BGP, page 29-4](#)
- [Monitoring BGP, page 29-25](#)
- [Configuration Example for BGP, page 29-26](#)
- [Feature History for BGP, page 29-27](#)

Information About BGP

BGP is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP). This section includes the following topics:

- [When to Use BGP, page 29-1](#)
- [Routing Table Changes, page 29-2](#)

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- **Weight** -- This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
- **Local preference** -- The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- **Multi-exit discriminator** -- The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- **Origin** -- The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
 - **IGP**- The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
 - **EGP**-The route is learned via the Exterior Border Gateway Protocol (EBGP).
 - **Incomplete**- The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- **AS_path** -- When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- **Next hop** -- The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS.
- **Community** -- The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
 - **no-export**- Do not advertise this route to EBGP peers.
 - **no-advertise**- Do not advertise this route to any peer.
 - **internet**- Advertise this route to the Internet community; all routers in the network belong to it.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

Licensing Requirements for BGP

Model	License Requirement
ASAv	Standard or Premium License.
All other models	Base License.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

Supported in single and multiple context mode.

Firewall Mode Guidelines

Does not support transparent firewall mode. BGP is supported only in router mode.

Failover Guidelines

Supports Stateful Failover in single and multiple context mode.



Note When you delete and reapply the BGP configuration in the user context allow a delay of 60 seconds, to enable the slave/ standby ASA unit to sync.

Clustering Guidelines

Does not support clustering.

IPv6 Guidelines

Does not support IPv6.

Graceful Restart Guidelines

Does not support graceful restart.

Configuring BGP

This section describes how to enable the BGP process on your system. After you have enabled BGP, see the following topics to learn how to customize the BGP process on your system.

- [Task Flow for Configuring BGP, page 29-4](#)
- [Enabling BGP, page 29-5](#)
- [Defining the Best Path for a BGP Routing Process, page 29-6](#)
- [Configuring Policy Lists, page 29-8](#)
- [Configuring AS Path Filters, page 29-9](#)
- [Configuring Community Rules, page 29-10](#)
- [Configuring IPv4 Address Family Settings, page 29-10](#)

Task Flow for Configuring BGP

To configure BGP, perform the following steps:

-
- Step 1** In the CLI, enable BGP, and configure general BGP parameters.
 - Step 2** Define the best path for the BGP routing process and configure the best path configuration parameters.
 - Step 3** Add and configure policy lists.
 - Step 4** Add and configure AS path filters.
 - Step 5** Add and configure Community Rules.
 - Step 6** Configure IPv4 Address Family settings.
-

Enabling BGP

This section describes the steps required to enable BGP routing, establish a BGP routing process and configure general BGP parameters.

Detailed Steps

	Command	Purpose
Step 1	router bgp <i>as-num</i> Example: ciscoasa(config)# router bgp 2	Enables a BGP routing process, which places the ASA in router configuration mode. The <i>as-num</i> argument is the autonomous system number of the BGP routing process. Valid values are from 1-4294967295 and 1.0-XX.YY.
Step 2	bgp maxas-limit <i>number</i> Example: hostname(config-router)# bgp maxas-limit 15	Discards routes that have as-path segments that exceed a specified value. The <i>number</i> argument specifies the maximum number of autonomous system segments, allowed. Valid values are from 1 to 254.
Step 3	bgp log-neighbor-changes Example: hostname(config-router)# bgp log-neighbor-changes	Logs BGP neighbor resets.
Step 4	bgp transport path-mtu-discovery Example: hostname(config-router)# bgp transport path-mtu-discovery	Enables BGP to automatically discover the best TCP path MTU for each BGP session.
Step 5	bgp fast-external-fallover Example: hostname(config-router)# bgp fast-external-fallover	Enables BGP to terminate external BGP sessions of any directly adjacent peer if the link used to reach the peer goes down; without waiting for the hold-down timer to expire.
Step 6	bgp enforce-first-as Example: hostname(config-router)# bgp enforce-first-as	Allows a BGP routing process to discard updates received from an external BGP (eBGP) peers that do not list their autonomous system (AS) number as the first AS path segment in the AS_PATH attribute of the incoming route.

	Command	Purpose
Step 7	bgp asnotation dot Example: hostname(config-router)# bgp asnotation dot	Changes the default display and regular expression match format of BGP 4-byte autonomous system numbers from asplain (decimal values) to dot notation.
Step 8	timers bgp <i>keepalive holdtime</i> [<i>min-holdtime</i>] Example: hostname(config-router)# timers bgp 80 120	Adjusts BGP network timers. The argument <i>keepalive</i> is the frequency (in seconds) with which the ASA sends keepalive messages to its peer. The default value 60 seconds. The argument <i>holdtime</i> is the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. The default is 180 seconds. (Optional) The argument <i>min-holdtime</i> is the interval (in seconds) after not receiving a keepalive message from a neighbor, that the ASA declares a neighbor dead.

Defining the Best Path for a BGP Routing Process

This section describes the steps required to configure the BGP best path. For more information on the best path, see [BGP Path Selection, page 29-3](#).

Detailed Steps

	Command	Purpose
Step 1	router bgp <i>as-num</i> Example: ciscoasa(config)# router bgp 2	Enables a BGP routing process, which places the ASA in router configuration mode. The <i>as-num</i> argument is the autonomous system number of the BGP routing process.
Step 2	bgp default local-preference <i>number</i> Example: ciscoasa(config-router)# bgp default local-preference 500	Changes the default local preference value. The default value is 100. The <i>number</i> argument is any value between 0 and 4294967295. Higher values indicate higher preference.
Step 3	bgp always-compare-med Example: ciscoasa(config-router)# bgp always-compare-med	Enables Multi Exit Discriminator (MED) comparison among paths learned from neighbors in different autonomous systems.
Step 4	bgp bestpath compare-routerid Example: hostname(config-router)# bgp bestpath compare-routerid	Compares between similar routes received from external BGP (eBGP) peers during the best path selection process and switches the best path to the route with the lowest router ID.
Step 5	bgp deterministic-med Example: hostname(config-router)# bgp deterministic-med	Selects the best MED path advertised from the neighboring AS.
Step 6	bgp bestpath med missing-as-worst Example: hostname(config-router)# bgp bestpath med missing-as-worst	Sets a path with a missing MED attribute as the least preferred path.

Configuring Policy Lists

When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. This section describes the steps required to configure policy lists.

Detailed Steps

	Command	Purpose
Step 1	<p>policy-list <i>policy_list_name</i> {permit deny}</p> <p>Example: <pre>ciscoasa(config)# policy-list Example-policy-list1 permit</pre></p>	<p>Enables the policy-map configuration mode and allows you to create a BGP policy list.</p> <p>The <i>policy_list_name</i> argument indicates the name of the configured policy.</p> <p>The permit keyword allows access for matching conditions.</p> <p>The deny keyword denies access for matching conditions.</p>
Step 2	<p>match interface [...<i>interface_name</i>]</p> <p>Example: <pre>ciscoasa(config-policy-list)# match interface outside</pre></p>	<p>Distributes routes that have their next hop out one of the interfaces specified.</p> <p>The <i>interface_name</i> argument indicates one or more interfaces.</p>
Step 3	<p>match ip {address next-hop route-source}</p> <p>Example: <pre>hostname(config-policy-list)# match ip address hostname(config-policy-list)# match ip next-hop hostname(config-policy-list)# match ip route-source</pre></p>	<p>Redistributes routes by matching either or all of the following: the destination address, next hop router address, and router/access server source.</p>
Step 4	<p>match as-path</p> <p>Example: <pre>hostname(config-policy-list)# match as-path</pre></p>	<p>Matches a BGP autonomous system path.</p>
Step 5	<p>match community {<i>community-list_name</i> exact-match}</p> <p>Example: <pre>hostname(config-policy-list)# match community ExampleCommunity1</pre></p>	<p>Matches a BGP community.</p> <p>The <i>community-list_name</i> argument indicates one or more community lists.</p> <p>The exact-match keyword indicates that an exact match is required. All of the communities and only those communities specified must be present.</p>

	Command	Purpose
Step 6	<code>match metric</code> Example: hostname(config-policy-list)# match metric	Redistributes routes with the metric specified.
Step 7	<code>match tag</code> Example: hostname(config-policy-list)# match tag	Redistributes routes in the routing table that match the specified tags.

Configuring AS Path Filters

An AS path filter allows you to filter the routing update message by using access lists and look at the individual prefixes within an update message. If a prefix within the update message matches the filter criteria then that individual prefix is filtered out or accepted depending on what action the filter entry has been configured to carry out. This section describes the steps required to configure AS path filters.



Note

The **as-path access-lists** are not the same as the regular firewall ACLs.

Detailed Steps

	Command	Purpose
Step 1	<code>as-path access-list acl-number</code> <code>{permit deny} regexp</code> Example: ciscoasa(config)# as-path access-list 35 permit testaspath	Configures an autonomous system path filter using a regular expression in the global configuration mode. The <i>acl-number</i> argument is a number that specifies the AS-path access-list number. Valid values are from 1 to 500. The <i>regexp</i> argument represents a regular expression that defines the AS-path filter. The autonomous system number is expressed in the range from 1 to 65535.

Configuring Community Rules

A community is a group of destinations that share some common attribute. You can use community lists to create groups of communities to use in a match clause of a route map. Just like an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded. This section describes the steps required to configure community rules.

Detailed Steps

Command	Purpose
<p>Step 1</p> <pre>community-list {standard community list-name {deny permit} [community-number] [AA:NN] [internet] [no-advertise][no-export]} {expanded expanded list-name {deny permit} regexp}</pre> <p>Example: ciscoasa(config)# community-list standard excomm1 permit 100 internet no-advertise no-export</p>	<p>Creates or configures a BGP community list and controls access to it.</p> <ul style="list-style-type: none"> The argument <i>community list-name</i> allows you to define a named standard community list. The argument <i>community-number</i> specifies a community as a 32-bit number from 1 to 4294967200. A single community can be entered or multiple communities can be entered, each separated by a space. The argument <i>AA:NN</i> specifies an autonomous system number and network number entered in the 4-byte new community format. This value is configured with two 2-byte numbers separated by a colon. A number from 1 to 65535 can be entered for each 2-byte number. A single community can be entered or multiple communities can be entered, each separated by a space. The argument <i>expanded list-name</i> configures a named expanded community list. The <i>regexp</i> argument represents a regular expression that defines the AS-path filter. The autonomous system number is expressed in the range from 1 to 65535. <p>Note Regular expressions can be used only with expanded community lists.</p>

Configuring IPv4 Address Family Settings

The IPv4 settings for BGP can be set up from the IPv4 family option within the BGP configuration setup. The IPv4 family section includes subsections for General settings, Aggregate address settings, Filtering settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv4 family.

This section describes how to customize the BGP IPv4 family settings and includes the following topics:

- [Configuring IPv4 Family General Settings, page 29-11](#)
- [Configuring IPv4 Family Aggregate Address Settings, page 29-14](#)
- [Configuring IPv4 Family Filtering Settings, page 29-14](#)

- [Configuring IPv4 Family BGP Neighbor Settings, page 29-15](#)
- [Configuring IPv4 Network Settings, page 29-21](#)
- [Configuring Redistribution Settings, page 29-22](#)
- [Configuring Route Injection Settings, page 29-23](#)

Configuring IPv4 Family General Settings

This section describes the steps required to configure the general IPv4 settings.

Detailed Steps

	Command	Purpose
Step 1	router bgp <i>as-num</i> Example: ciscoasa(config)# router bgp 2	Enables a BGP routing process, which places the router in router configuration mode. The <i>as-num</i> argument is the autonomous system number of the BGP routing process.
Step 2	address-family ipv4 [unicast] Example: hostname(config-router)# address-family ipv4[unicast]	Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes. The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.
Step 3	bgp router-id <i>A.B.C.D</i> Example: hostname(config-router-af)# bgp router-id 10.86.118.3	(Optional) Configures a fixed router ID for the local BGP routing process. The argument <i>A.B.C.D</i> specifies a router identifier in the form of an IP address. Note If you do not specify a router ID, it is automatically assigned.
Step 4	distance bgp <i>external-distance internal-distance local-distance</i> Example: hostname(config-router-af)# distance bgp 80 180 180	Configures the administrative distance for BGP routes. <ul style="list-style-type: none"> The argument <i>external-distance</i> specifies administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The argument <i>internal-distance</i> specifies administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The argument <i>local-distance</i> specifies administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255.
Step 5	table-map {WORD <i>route-map_name</i> } Example: hostname(config-router-af)# table-map example1	Modifies metric and tag values when the IP routing table is updated with BGP learned routes. The argument <i>route-map_name</i> specifies the route map name from the route-map command.
Step 6	default-information originate Example: hostname(config-router-af)# default-information originate	Configures a BGP routing process to distribute a default route (network 0.0.0.0).
Step 7	auto-summary Example: hostname(config-router-af)# auto-summary	Configures automatic summarization of subnet routes into network-level routes.

	Command	Purpose
Step 8	bgp suppress-inactive Example: hostname(config-router-af)# bgp suppress-inactive	Suppresses the advertisement of routes that are not installed in the routing information base (RIB).
Step 9	synchronization Example: hostname(config-router-af)# synchronization	Synchronizes between BGP and your Interior Gateway Protocol (IGP) system.
Step 10	bgp redistribute-internal Example: hostname(config-router-af)# bgp redistribute-internal	Configures iBGP redistribution into an IGP, such as OSPF.
Step 11	bgp scan-time scanner-interval Example: hostname(config-router-af)# bgp scan-time 15	Configures scanning intervals of BGP routers for next hop validation. The argument <i>scanner-interval</i> specifies scanning interval of BGP routing information. Valid values are from 5 to 60 seconds. The default is 60 seconds.
Step 12	bgp nexthop trigger {delay seconds enable} Example: hostname(config-router-af)# bgp nexthop trigger delay 15	Enables you to configure BGP next-hop address tracking. <ul style="list-style-type: none"> • The trigger keyword specifies the use of BGP next-hop address tracking. Use this keyword with the delay keyword to change the next-hop tracking delay. Use this keyword with the enable keyword to enable next-hop address tracking. • The delay keyword changes the delay interval between checks on updated next-hop routes installed in the routing table. • The <i>seconds</i> argument specifies the delay in seconds. Range is from 0 to 100. Default is 5. • The enable keyword enables BGP next-hop address tracking.
Step 13	maximum-paths {number_of_paths ibgp number_of_paths} Example: router(config-router-af)# maximum-paths ibgp 2	Controls the maximum number of parallel iBGP routes that can be installed in a routing table. The <i>number_of_paths</i> argument specifies the number of routes to install to the routing table. Valid values are between 1 and 3. Note If the ibgp keyword is not used, then the <i>number_of_paths</i> argument controls the maximum number of parallel EBGp routes.

Configuring IPv4 Family Aggregate Address Settings

This section describes the steps required to define the aggregation of specific routes into one route.

Detailed Steps

	Command	Purpose
Step 1	<pre>router bgp <i>as-num</i></pre> <p>Example: ciscoasa(config)# router bgp 2</p>	<p>Enables a BGP routing process, which places the ASA in router configuration mode.</p> <p>The <i>as-num</i> argument is the autonomous system number of the BGP routing process.</p>
Step 2	<pre>address-family ipv4 [<i>unicast</i>]</pre> <p>Example: hostname(config-router)# address-family ipv4 [<i>unicast</i>]</p>	<p>Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes.</p> <p>The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.</p>
Step 3	<pre>aggregate-address <i>address mask</i> [<i>as-set</i>][<i>summary-only</i>][<i>suppress-map</i> <i>map-name</i>][<i>advertise-map</i> <i>map-name</i>][<i>attribute-map</i> <i>map-name</i>]</pre> <p>Example: hostname(config-router-af) aggregate-address 10.86.118.0 255.255.255.0 as-set summary-only suppress-map example1 advertise-map example1 attribute-map example1</p>	<p>Creates an aggregate entry in a BGP database.</p> <p>The argument <i>address</i> specifies the aggregate address.</p> <p>The argument <i>mask</i> specifies the aggregate mask.</p> <p>The argument <i>map-name</i> specifies the route map.</p>

Configuring IPv4 Family Filtering Settings

This section describes the steps required to filter routes or networks received in incoming BGP updates.

Detailed Steps

	Command	Purpose
Step 1	<code>router bgp <i>as-num</i></code> Example: <code>ciscoasa(config)# router bgp 2</code>	Enables a BGP routing process, which places the ASA in router configuration mode. The <i>as-num</i> argument is the autonomous system number of the BGP routing process.
Step 2	<code>address-family ipv4 [unicast]</code> Example: <code>hostname(config-router)# address-family ipv4 [unicast]</code>	Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes. The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.
Step 3	<code>distribute-list <i>acl-number</i> in out[]</code> Example: <code>hostname(config-router-af)# distribute-list ExampleAcl in bgp 2</code>	Filters routes or networks received in incoming or advertised in outgoing BGP updates. The argument <i>acl-number</i> specifies IP access list number. The access list defines which networks are to be received and which are to be suppressed in routing updates. TBD.

Configuring IPv4 Family BGP Neighbor Settings

This section describes the steps required to define BGP neighbors and neighbor settings.

**Note**

You cannot add neighbors that support graceful restart, because ASA 9.2.1 does not support graceful restart.

Detailed Steps

	Command	Purpose
Step 1	router bgp <i>as-num</i> Example: ciscoasa(config)# router bgp 2	Enables a BGP routing process, which places the router in router configuration mode. The <i>as-num</i> argument is the autonomous system number of the BGP routing process.
Step 2	address-family ipv4 [unicast] Example: hostname(config-router)# address-family ipv4[unicast]	Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes. The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.
Step 3	neighbor ip-address remote-as <i>as-number</i> Example: hostname(config-router-af)# neighbor 10.86.118.12 remote-as 3	Adds an entry to the BGP neighbor table. The argument <i>ip-address</i> specifies the IP address of the neighbor. The argument <i>as-number</i> specifies the autonomous system to which the neighbor belongs.
Step 4	neighbor ip-address shutdown Example: hostname(config-router-af)# neighbor 10.86.118.12 shutdown 3	(Optional) Enables you to disable a neighbor or peer group. The argument <i>ip-address</i> specifies the IP address of the neighbor.
Step 5	neighbor ip-address activate Example: hostname(config-router-af)# neighbor 10.86.118.12 activate	Exchanges information with a BGP neighbor. The argument <i>ip-address</i> specifies the IP address of the neighbor.

Command	Purpose
<p>Step 6</p> <p>neighbor {<i>ip-address</i>} distribute-list {<i>access-list-name</i>} {in out}</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 distribute-list ExampleAcl in</pre></p>	<p>Distributes BGP neighbor information as specified in an access list.</p> <ul style="list-style-type: none"> • The argument <i>ip-address</i> specifies the IP address of the neighbor. • The argument <i>access-list-number</i> specifies the number of a standard or extended access list. The range of a standard access list number is from 1 to 99. The range of an extended access list number is from 100 to 199. • The argument <i>expanded-list-number</i> specifies the number of an expanded access list number. The range of an expanded access list is from 1300 to 2699. • The argument <i>access-list-name</i> specifies the name of a standard or extended access list. • The argument <i>prefix-list-name</i> specifies the name of a BGP prefix list. • The keyword in implies that the access list is applied to incoming advertisements to that neighbor. • The keyword out implies that the access list is applied to outgoing advertisements to that neighbor.
<p>Step 7</p> <p>neighbor {<i>ip-address</i>} route-map <i>map-name</i> {in out}</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 route-map example1 in</pre></p>	<p>Applies a route map to incoming or outgoing routes.</p> <ul style="list-style-type: none"> • The argument <i>ip-address</i> specifies the IP address of the neighbor. • The argument <i>map-name</i> specifies the name of a route map. • The keyword in applies a route map to incoming routes. • The keyword out applies a route map to outgoing routes.
<p>Step 8</p> <p>neighbor {<i>ip-address</i>} prefix-list <i>prefix-list-name</i> {in out}</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 prefix-list NewPrefixList in</pre></p>	<p>Distributes BGP neighbor information as specified in a prefix list.</p> <ul style="list-style-type: none"> • The argument <i>ip-address</i> specifies the IP address of the neighbor. • The argument <i>prefix-list-name</i> specifies the name of a prefix list. • The keyword in implies that the prefix list is applied to incoming advertisements from that neighbor. • The keyword out implies that the prefix list is applied to outgoing advertisements to that neighbor.

	Command	Purpose
Step 9	<p>neighbor {<i>ip-address</i>} filter-list <i>access-list-number</i> {in out}</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 filter-list 5 in</pre></p>	<p>Sets up a filter list.</p> <ul style="list-style-type: none"> The argument <i>ip-address</i> specifies the IP address of the neighbor. The argument <i>access-list-name</i> specifies the number of an autonomous system path access list. You define this access list with the ip as-path access-list command. The keyword in implies that the access list is applied to incoming advertisements from that neighbor. The keyword out implies that the access list is applied to outgoing advertisements to that neighbor.
Step 10	<p>neighbor {<i>ip-address</i>} maximum-prefix <i>maximum</i> [<i>threshold</i>][restart <i>restart interval</i>][warning-only]</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 maximum-prefix 7 75 restart 12</pre></p>	<p>Controls the number of prefixes that can be received from a neighbor.</p> <ul style="list-style-type: none"> The argument <i>ip-address</i> specifies the IP address of the neighbor. The argument <i>maximum</i> specifies the maximum number of prefixes allowed from this neighbor. The argument <i>threshold</i> is an integer specifying at what percentage of maximum the router starts to generate a warning message. The range is from 1 to 100; the default is 75 (percent). The <i>restart interval</i> is an integer value (in minutes) that specifies the time interval after which the BGP neighbor restarts. The warning-only keyword allows the router to generate a log message when the maximum is exceeded, instead of terminating the peering.
Step 11	<p>neighbor {<i>ip-address</i>} default-originate[route-map <i>map-name</i>]</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 default-originate route-map example1</pre></p>	<p>Allows a BGP speaker (the local router) to send the default route 0.0.0.0 to a neighbor for use as a default route.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>map-name</i> is the name of the route-map. The route map allows route 0.0.0.0 to be injected conditionally.</p>
Step 12	<p>neighbor {<i>ip-address</i>} advertisement-interval <i>seconds</i></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 advertisement-interval 15</pre></p>	<p>Sets the minimum interval between the sending of BGP routing updates.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>seconds</i> is the time (in seconds). Valid values are from 0 to 600.</p>

	Command	Purpose
Step 13	<p>neighbor {<i>ip-address</i>} remove-private-as</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 remove-private-as</pre></p>	<p>Removes private autonomous system numbers from outbound routing updates.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p>
Step 14	<p>neighbor {<i>ip-address</i>}timers <i>keepalive</i> <i>holdtime</i> <i>min holdtime</i></p> <p>Example: <pre>router(config-router-af)# neighbor 10.86.118.12 timers 15 20 12</pre></p>	<p>Sets the timers for a specific BGP peer or peer group.</p> <ul style="list-style-type: none"> The argument <i>ip-address</i> specifies the IP address of the neighbor. The argument <i>keepalive</i> specifies the frequency (in seconds) with which the ASA sends keepalive messages to its peer. The default is 60 seconds. Valid values are from 0 to 65535. The argument <i>holdtime</i> specifies the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. The default is 180 seconds. The argument <i>min holdtime</i> specifies the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Please verify this.
Step 15	<p>neighbor {<i>ip-address</i>} password <i>string</i></p> <p>Example: <pre>router(config-router-af)# neighbor 10.86.118.12 password test</pre></p>	<p>Enables Message Digest 5 (MD5) authentication on a TCP connection between two BGP peers.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>string</i> is a case-sensitive password of up to 25 characters when the service password-encryption command is enabled and up to 81 characters when the service password-encryption command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces.</p> <p>Note You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.</p>
Step 16	<p>neighbor {<i>ip-address</i>} send-community[both standard extended]</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 send-community</pre></p>	<p>Specifies that communities attributes should be sent to a BGP neighbor.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The both keyword specifies that both standard and extended communities will be sent.</p> <p>The standard keyword specifies that only standard communities will be sent.</p> <p>The extended keyword specifies that only extended communities will be sent.</p>
Step 17	<p>neighbor {<i>ip-address</i>}next-hop-self</p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 next-hop-self</pre></p>	<p>Configures the router as the next hop for a BGP-speaking neighbor or peer group.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p>

Command	Purpose
<p>Step 18 <code>neighbor {ip-address} ebgp-multihop [ttl]</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 ebgp-multihop 5</pre></p>	<p>Accepts and attempts BGP connections to external peers residing on networks that are not directly connected.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>ttl</i> specifies time-to-live in the range from 1 to 255 hops.</p>
<p>Step 19 <code>neighbor {ip-address} disable-connected-check</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 disable-connected-check</pre></p>	<p>Disables connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p>
<p>Step 20 <code>neighbor {ip-address} ttl-security hops hop-count</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 ttl-security hops 15</pre></p>	<p>Secures a BGP peering session and configures the maximum number of hops that separate two external BGP (eBGP) peers.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>hop-count</i> is the number of hops that separate the eBGP peers. The TTL value is calculated by the router from the configured hop-count argument. Valid values are from 1 to 254.</p>
<p>Step 21 <code>neighbor {ip-address} weight number</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 weight 30</pre></p>	<p>Assigns a weight to a neighbor connection.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>number</i> is the weight to assign to a neighbor connection. Valid values are from 0 to 65535.</p>
<p>Step 22 <code>neighbor {ip-address} version number</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 version 4</pre></p>	<p>Configures the ASA to accept only a particular BGP version.</p> <p>The argument <i>ip-address</i> specifies the IP address of the neighbor.</p> <p>The argument <i>number</i> specifies the BGP version number. The version can be set to 2 to force the software to use only Version 2 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.</p>

Command	Purpose
<p>Step 23 <code>neighbor {ip-address} transport {connection-mode{active passive} path-mtu-discovery[disable]}</code></p> <p>Example: <pre>asa3(config-router-af)# neighbor 10.86.118.12 transport path-mtu-discovery</pre></p>	<p>Enables a TCP transport session option for a BGP session.</p> <ul style="list-style-type: none"> • The argument <i>ip-address</i> specifies the IP address of the neighbor. • The keyword connection-mode specifies the type of connection (active or passive). • The keyword path-mtu-discovery enables TCP transport path maximum transmission unit (MTU) discovery. TCP path MTU discovery is enabled by default. • The keyword disable disables TCP path MTU discovery.
<p>Step 24 <code>neighbor {ip-address} local-as [autonomous-system-number[no-prepend]]</code></p> <p>Example: <pre>hostname(config-router-af)# neighbor 10.86.118.12 local-as 5 no-prepend replace-as</pre></p>	<p>Customizes the AS_PATH attribute for routes received from an external Border Gateway Protocol (eBGP) neighbor.</p> <ul style="list-style-type: none"> • The argument <i>ip-address</i> specifies the IP address of the neighbor. • (Optional) The argument <i>autonomous-system-number</i> specifies the number of an autonomous system to prepend to the AS_PATH attribute. The range of values for this argument is any valid autonomous system number from 1 to 4294967295 or 1.0 to XX.YY. • Using the keyword no-prepend does not prepend the local autonomous system number to any routes received from the eBGP neighbor.

Configuring IPv4 Network Settings

This section describes the steps required to define the networks to be advertised by the BGP routing process.

Detailed Steps

	Command	Purpose
Step 1	<pre>router bgp <i>as-num</i></pre> <p>Example: ciscoasa(config)# router bgp 2 </p>	<p>Enables a BGP routing process, which places the ASA in router configuration mode.</p> <p>The <i>as-num</i> argument is the autonomous system number of the BGP routing process.</p>
Step 2	<pre>address-family ipv4 [unicast]</pre> <p>Example: hostname(config-router)# address-family ipv4 [unicast]</p>	<p>Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes.</p> <p>The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.</p>
Step 3	<pre>network {<i>network-number</i> [mask <i>network-mask</i>]} [route-map <i>map-tag</i>]</pre> <p>Example: hostname(config-router-af)# network 10.86.118.13 mask 255.255.255.255 route-map example1</p>	<p>Specifies the networks to be advertised by the BGP routing processes.</p> <ul style="list-style-type: none"> • The argument <i>network-number</i> specifies the network that BGP will advertise. • (Optional) The argument <i>network-mask</i> specifies the network or subnetwork mask with mask address. • (Optional) The argument <i>map-tag</i> specifies the identifier of a configured route map. The route map should be examined to filter the networks to be advertised. If not specified, all networks are advertised.

Configuring Redistribution Settings

This section describes the steps required to define the conditions for redistributing routes from another routing domain into BGP.

Detailed Steps

	Command	Purpose
Step 1	<pre>router bgp <i>as-num</i></pre> <p>Example: <pre>ciscoasa(config)# router bgp 2</pre></p>	<p>Enables a BGP routing process, which places the ASA in router configuration mode.</p> <p>The <i>as-num</i> argument is the autonomous system number of the BGP routing process.</p>
Step 2	<pre>address-family ipv4 [unicast]</pre> <p>Example: <pre>hostname(config-router)# address-family ipv4 [unicast]</pre></p>	<p>Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes.</p> <p>The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.</p>
Step 3	<pre>redistribute <i>protocol</i> [<i>process-id</i>] [<i>metric</i>] [route-map [<i>map-tag</i>]]</pre> <p>Example: <pre>hostname(config-router-af)# redistribute ospf 2 route-map example1 match external</pre></p>	<p>Redistributes routes from another routing domain into a BGP autonomous system.</p> <ul style="list-style-type: none"> • The argument <i>protocol</i> specifies the source protocol from which routes are being redistributed. It can be one of the following: Connected, EIGRP, OSPF, RIP or Static. • The argument <i>process-id</i> specifies a name for the specific routing process. • The argument <i>metric</i> specifies the metric for the redistributed route. • The argument <i>map-tag</i> specifies the identifier of a configured route map. The route map should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed.

Configuring Route Injection Settings

This section describes the steps required to define the routes to be conditionally injected into the BGP routing table.

Detailed Steps

	Command	Purpose
Step 1	<pre>router bgp <i>as-num</i></pre> <p>Example: ciscoasa(config)# router bgp 2 </p>	<p>Enables a BGP routing process, which places the ASA in router configuration mode.</p> <p>The <i>as-num</i> argument is the autonomous system number of the BGP routing process.</p>
Step 2	<pre>address-family ipv4 [unicast]</pre> <p>Example: hostname(config-router)# address-family ipv4[unicast]</p>	<p>Enables you to enter address family or router scope address family configuration mode to configure a routing session using standard IP Version 4 (IPv4) address prefixes.</p> <p>The keyword unicast specifies IPv4 unicast address prefixes. This is the default, even if not specified.</p>
Step 3	<pre>bgp inject-map <i>inject-map</i> exist-map <i>exist-map</i> [copy-attributes]</pre> <p>Example: hostname(config-router-af)# bgp inject-map example1 exist-map example2 copy-attributes</p>	<p>Configures conditional route injection to inject more specific routes into a BGP routing table.</p> <p>The argument <i>inject-map</i> specifies the name of the route map that specifies the prefixes to inject into the local BGP routing table.</p> <p>The argument <i>exist-map</i> specifies the name of the route map containing the prefixes that the BGP speaker will track.</p> <p>(Optional) The key word copy-attributes configures the injected route to inherit attributes of the aggregate route.</p>

Monitoring BGP

You can use the following commands to monitor the BGP routing process. For examples and descriptions of the command output, see the command reference. Additionally, you can disable the logging of neighbor change messages and neighbor warning messages.

To monitor or disable various BGP routing statistics, enter one of the following commands:

Command	Purpose
Monitoring BGP Routing	
<code>show bgp [ip-address [mask [longer-prefixes [injected] shorter-prefixes [length]]] prefix-list name route-map name]</code>	Displays the entries in the BGP routing table.
<code>show bgp cidr-only</code>	Displays routes with non-natural network masks (that is, classless interdomain routing, or CIDR).
<code>show bgp community community-number [exact-match] [no-advertise] [no-export]</code>	Display routes that belong to specified BGP communities.
<code>show bgp community-list community-list-name [exact-match]</code>	Displays routes that are permitted by the BGP community list.
<code>show bgp filter-list access-list-number</code>	Displays routes that conform to a specified filter list.
<code>show bgp injected-paths</code>	Displays all the injected paths in the BGP routing table.
<code>show bgp ipv4 unicast</code>	Displays entries in the IP version 4 (IPv4) BGP routing table for unicast sessions.
<code>show bgp neighbors ip_address</code>	Displays information about BGP and TCP connections to neighbors.
<code>show bgp paths [LINE]</code>	Displays all the BGP paths in the database.
<code>show bgp pending-prefixes</code>	Displays prefixes that are pending deletion.
<code>show bgp prefix-list prefix_list_name [WORD]</code>	Displays routes that match a specified prefix list.
<code>show bgp regexp regexp</code>	Displays routes that match the autonomous system path regular expression.
<code>show bgp replication [index-group ip-address]</code>	Displays update replication statistics for BGP update groups.
<code>show bgp rib-failure</code>	Displays BGP routes that failed to install in the Routing Information Base (RIB) table.
<code>show bgp route-map map-name</code>	Displays entries in the BGP routing table, based on the route map specified.
<code>show bgp summary</code>	Display the status of all BGP connections.
<code>show bgp system-config</code>	Display the system context specific BGP configuration in multi-context mode. Note This command is available in all user contexts in multi-context mode.

Command (continued)	Purpose (continued)
<code>show bgp update-group</code>	Display information about the BGP update groups.
Disabling BGP Log Messages	
<code>no bgp log-neighbor-changes</code>	Disables the logging of neighbor change messages. Enter this command in router configuration mode for the BGP routing process.
	Note By default, neighbor changes are logged.

Configuration Example for BGP

This example shows how to enable and configure BGP with various optional processes.

-
- Step 1** To enable BGP, enter the following commands:
- ```
ciscoasa(config)# router bgp 2
```
- Step 2** To enable you to discard routes that have a number of as-path segments that exceed the specified value:
- ```
ciscoasa(config-router)# bgp maxas-limit 15
```
- Step 3** To enable logging of BGP neighbor resets:
- ```
ciscoasa(config-router)# bgp log-neighbor-changes
```
- Step 4** To enable BGP to automatically discover the best TCP path MTU for each BGP session:
- ```
ciscoasa(config-router)# bgp transport path-mtu-discovery
```
- Step 5** To enable BGP to terminate external BGP sessions of any directly adjacent peer if the link used to reach the peer goes down; without waiting for the hold-down timer to expire:
- ```
ciscoasa(config-router)# bgp fast-external-fallover
```
- Step 6** To adjust BGP timers:
- ```
ciscoasa(config-router)# timers bgp 80 120
```
-

Feature History for BGP

Table 29-1 lists each feature change and the platform release in which it was implemented.

Table 29-1 Feature History for BGP

Feature Name	Platform Releases	Feature Information
BGP Support	9.2(1)	<p>Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Border Gateway Protocol.</p> <p>We introduced the following commands: router bgp, bgp maxas-limit, bgp log-neighbor-changes, bgp transport path-mtu-discovery, bgp fast-external-falover, bgp enforce-first-as, bgp asnotation dot, timers bgp, bgp default local-preference, bgp always-compare-med, bgp bestpath compare-routerid, bgp deterministic-med, bgp bestpath med missing-as-worst, policy-list, match as-path, match community, match metric, match tag, as-path access-list, community-list, address-family ipv4, bgp router-id, distance bgp, table-map, bgp suppress-inactive, bgp redistribute-internal, bgp scan-time, bgp nexthop, aggregate-address, neighbor, bgp inject-map, show bgp, show bgp cidr-only, show bgp all community, show bgp all neighbors, show bgp community, show bgp community-list, show bgp filter-list, show bgp injected-paths, show bgp ipv4 unicast, show bgp neighbors, show bgp paths, show bgp pending-prefixes, show bgp prefix-list, show bgp regexp, show bgp replication, show bgp rib-failure, show bgp route-map, show bgp summary, show bgp system-config, show bgp update-group, clear route network, maximum-path, network.</p> <p>We modified the following commands: show route, show route summary, show running-config router, clear config router, clear route all, timers lsa arrival, timers pacing, timers throttle, redistribute bgp.</p>

