



BGP

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

- [About BGP, on page 1](#)
- [Guidelines for BGP, on page 4](#)
- [Configure BGP, on page 5](#)
- [Monitoring BGP, on page 22](#)
- [History for BGP, on page 23](#)

About BGP

BGP is an inter and intra 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).

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).

BGP can also be used for carrying routing information for IPv6 prefix over IPv6 networks.



Note When a BGPv6 device joins the cluster, it generates a soft traceback when logging level 7 is enabled.

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.



Note AS loop detection is done by scanning the full AS path (as specified in the AS_PATH attribute), and checking that the AS number of the local system does not appear in the AS path. By default, EBGp advertises the learned routes to the same peer to prevent additional CPU cycles on the ASA in performing loop checks and to avoid delays in the existing outgoing update tasks.

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.
- Determine if multiple paths require installation in the routing table for [BGP Multipath, on page 3](#).
- 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.

BGP Multipath

BGP Multipath allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix. Traffic to the destination prefix is then shared across all installed paths.

These paths are installed in the table together with the best path for load-sharing. BGP Multipath does not affect best-path selection. For example, a router still designates one of the paths as the best path, according to the algorithm, and advertises this best path to its BGP peers.

In order to be candidates for multipath, paths to the same destination need to have these characteristics equal to the best-path characteristics:

- Weight
- Local preference
- AS-PATH length
- Origin code
- Multi Exit Discriminator (MED)
- One of these:

- Neighboring AS or sub-AS (before the addition of the BGP Multipaths)
- AS-PATH (after the addition of the BGP Multipaths)

Some BGP Multipath features put additional requirements on multipath candidates:

- The path should be learned from an external or confederation-external neighbor (eBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric.

These are the additional requirements for internal BGP (iBGP) multipath candidates:

- The path should be learned from an internal neighbor (iBGP).
- The IGP metric to the BGP next hop should be equal to the best-path IGP metric, unless the router is configured for unequal-cost iBGP multipath.

BGP inserts up to n most recently received paths from multipath candidates into the IP routing table, where n is the number of routes to install to the routing table, as specified when you configure BGP Multipath. The default value, when multipath is disabled, is 1.

For unequal-cost load balancing, you can also use BGP Link Bandwidth.



Note The equivalent next-hop-self is performed on the best path that is selected among eBGP multipaths before it is forwarded to internal peers.

Guidelines for BGP

Context Mode Guidelines

- Supported in single and multiple context mode.
- Only one Autonomous System (AS) number is supported for all contexts.

Firewall Mode Guidelines

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

IPv6 Guidelines

Supports IPv6. Graceful restart is not supported for IPv6 address family.

Additional Guidelines

- The system does not add route entry for the IP address received over PPPoE in the CP route table. BGP always looks into CP route table for initiating the TCP session, hence BGP does not form TCP session. Thus, BGP over PPPoE is not supported.
- To avoid adjacency flaps due to route updates being dropped if the route update is larger than the minimum MTU on the link, ensure that you configure the same MTU on the interfaces on both sides of the link.

- The BGP table of the member unit is not synchronized with the control unit table. Only its routing table is synchronized with the control unit routing table.

Configure BGP

This section describes how to enable and configure the BGP process on your system.

Procedure

-
- | | |
|---------------|--|
| Step 1 | Enable BGP, on page 5. |
| Step 2 | Define the Best Path for a BGP Routing Process, on page 6. |
| Step 3 | Configure Policy Lists, on page 7. |
| Step 4 | Configure AS Path Filters, on page 8. |
| Step 5 | Configure Community Rules, on page 9. |
| Step 6 | Configure IPv4 Address Family Settings, on page 10. |
| Step 7 | Configure IPv6 Address Family Settings, on page 16. |
-

Enable BGP

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

Procedure

-
- | | |
|---------------|--|
| Step 1 | For single-mode, in ASDM, choose Configuration > Device Setup > Routing > BGP > General . |
| Note | For multi-mode, in ASDM choose Configuration > Context Management > BGP. After enabling BGP, switch to a security context and enable BGP by choosing Configuration > Device Setup > Routing > BGP > General . |
| Step 2 | Check the Enable BGP Routing check box. |
| Step 3 | In the AS Number field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to XX.YY. |
| Step 4 | (Optional) Check the Limit the number of AS numbers in the AS_PATH attribute of received routes check box to restrict the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254. |
| Step 5 | (Optional) Check the Log neighbor changes check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability. |

- Step 6** (Optional) Check the Use TCP path MTU discovery check box to use the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
- Step 7** (Optional) Check the Enable fast external failover check box to reset the external BGP session immediately upon link failure.
- Step 8** (Optional) Check the Enforce that first AS is peer's AS for EBGp routes check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
- Step 9** (Optional) Check the Use dot notation for AS numbers check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65535 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
- Step 10** Specify the timer information in the Neighbor timers area:
- In the Keepalive interval field, enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.
 - In the Hold Time field, enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. The default value is 180 seconds.
 - (Optional) In the Min. Hold Time field, enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.
- Note** A hold time of less than 20 seconds increases the possibility of peer flapping.
- Step 11** (Optional) In the Non Stop Forwarding section do the following:
- Check the Enable Graceful Restart check box to enable ASA peers to avoid a routing flap following a switchover.
 - In the Restart Time field, enter the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
 - In the Stale Path Time field, enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
- Step 12** Click **OK**.
- Step 13** Click **Apply**.

Define 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, on page 3](#).

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Best Path**.
The Best Path configuration pane appears.

- Step 2** In the Default Local Preference field, specify a value between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.
- Step 3** Check the Allow comparing MED from different neighbors check box to allow the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
- Step 4** Check the Compare router-id for identical EBGp paths check box to compare similar paths received from external BGP peers during the best path selection process and switch the best path to the route with the lowest router ID.
- Step 5** Check the Pick the best MED path among paths advertised from the neighboring AS check box to enable MED comparison among paths learned from confederation peers.add a new network entry. The comparison between MEDs is made only if no external autonomous systems are there in the path.
- Step 6** Check the Treat missing MED as the least preferred one check box to consider the missing MED attribute as having a value of infinity, making this path the least desirable; therefore, a path with a missing MED is least preferred.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.
-

Configure 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.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Policy Lists**.
- Step 2** Click **Add**.
- The Add Policy List dialog box appears. From this dialog box, you can add a policy list name, its redistribution access (that is, permit or deny), match interfaces, specify IP addresses, match the AS path, match community names list, match metrices, and match tag numbers.
- Step 3** In the **Policy List Name** field, enter a name for the policy list.
- Step 4** Click the **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 5** Check the **Match Interfaces** check box to distribute routes that have their next hop out of one of the interfaces specified and do one of the following:
- In the **Interface** field, enter the interface name.
 - In the Interface field, click the ellipses to manually browse and locate the interface. Choose one or more interfaces, click **Interface**, then click **OK**.
- Step 6** In the Specify IP area, configure the following:
- a) Check the Match Address check box to redistribute any routes that have a destination network number address that is permitted by a standard access list or prefix list, and performs policy routing on packets.

Specify an access list / prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

- b) Check the Match Next Hop check box to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified.

Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

- c) Check the Match Route Source check box to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list.

Specify an access list/ prefix list or click the ellipses to manually browse and locate an access list. Choose one or more access lists, click Access List, then click OK.

Step 7 Check the Match AS Path check box to match a BGP autonomous system path.

Specify an AS path filter or click the ellipses to manually browse and locate the AS Path Filter. Choose one or more AS Path Filters, click AS Path Filter, then click OK.

Step 8 Check the Match Community Names List check box to match a BGP community.

- a) Specify a community rule or click the ellipses to manually browse and locate the community rules. Choose one or more community rules, click Community Rules, then click OK.
b) Check the Match the specified community exactly check box to match a specific BGP community.

Step 9 Check the **Match Metrics** check box to redistribute routes with the metric specified. If you specify more than one metric, the routes can be matched with either metric.

Step 10 Check the **Match Tag Numbers** check box to redistribute routes in the routing table that match the specified tags. If you specify more than one tag number, routes can be matched with either metric.

Step 11 Click OK.

Step 12 Click Apply.

Configure 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.

Procedure

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > BGP > AS Path Filters**.

Step 2 Click Add.

The Add Filter dialog box appears. From this dialog box, you can add a filter name, its redistribution access (that is, permit or deny), and regular expression.

- Step 3** In the Name field, enter a name for the AS Path Filter.
- Step 4** Click the **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 5** Specify the regular expression. Click **Build** to build regular expression.
- Step 6** Click **Test** to test if a regular expression matches a string of your choice.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.

Configure 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.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > Community Rules > .**
- Step 2** Click **Add**.

The Add Community Rule dialog box appears. From this dialog box, you can add a rule name, rule type, its redistribution access (that is, permit or deny) and specific communities.
- Step 3** In the Rule Name field, enter a name for the community rule.
- Step 4** Click **Standard** or **Expanded** radio button to indicate the community rule type.
- Step 5** Click **Permit** or **Deny** radio button to indicate the redistribution access.
- Step 6** To add Standard Community Rules:
 - a) In the Communities field, specify a community number. Valid values are from 1 to 4294967200.
 - b) (Optional) Check the Internet (well-known community) check box to specify the Internet community. Routes with this community are advertised to all peers (internal and external).
 - c) (Optional) Check the Do not advertise to any peers (well-known community) check box to specify the no-advertise community. Routes with this community are not advertised to any peer (internal or external).
 - d) (Optional) Check the Do not export to next AS (well-known community) check box to specify the no-export community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.
- Step 7** To add expanded community rules:
 - a) In the **Regular Expression** field, enter a regular expression. Alternately, Click **Build** to build regular expression.
 - b) Click **Test** to examine if the regular expression built, matches a string of your choice.
- Step 8** Click **OK**.
- Step 9** Click **Apply**.

Configure 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.

Configure IPv4 Family General Settings

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

Procedure

-
- | | |
|----------------|--|
| Step 1 | In ASDM, choose Configuration > Device Setup > Routing > BGP > IPv4 Family . |
| Step 2 | Click General .

The General IPv4 family BGP parameters configuration pane is displayed. |
| Step 3 | Specify External, Internal and Local distances in the Administrative Distances area. |
| Step 4 | Choose a route map name from the Learned Routes Map drop-down list. Click Manage to add and configure route maps. |
| Step 5 | (Optional) Check the Generate Default Route check box to configure a BGP routing process to distribute a default route (network 0.0.0.0). |
| Step 6 | (Optional) Check the Summarize subnet routes into network-level routes check box to configure automatic summarization of subnet routes into network-level routes. |
| Step 7 | (Optional) Check the Advertise inactive routes check box to advertise routes that are not installed in the routing information base (RIB). |
| Step 8 | (Optional) Check the Redistribute iBGP into an IGP check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF. |
| Step 9 | (Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the Scanning Interval field. Valid values are from 5 to 60 seconds. |
| Step 10 | (Optional) Check the Enable address tracking check box to enable BGP next hop address tracking. Specify the delay interval between checks on updated next-hop routes installed in the routing table in the Delay Interval field. |
| Step 11 | (Optional) Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the Number of paths field and check the iBGP multipaths check box. |
| Step 12 | Click Apply . |
-

Configure IPv4 Family Aggregate Address Settings

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

Procedure

-
- | | |
|---------------|---|
| Step 1 | In ASDM, choose Configuration > Device Setup > Routing > BGP > IPv4 Family . |
|---------------|---|

- Step 2** Click **Aggregate Address**.
The Aggregate Address parameters configuration pane is displayed.
- Step 3** Click **Add**.
The Add Aggregate Address pane is displayed.
- Step 4** Specify a network object in the **Network** field.
- Step 5** Check the **Generate autonomous system set path information** check box to generate autonomous system set path information.
- Step 6** Check the **Filters all more- specific routes from the updates** check box to filter all more-specific routes from updates.
- Step 7** Choose a route-map from the Attribute Map drop-down list. Click **Manage** to add or configure a route map.
- Step 8** Choose a route-map from the Advertise Map drop-down list. Click **Manage** to add or configure a route.
- Step 9** Choose a route-map from the Suppress Map drop-down list. Click **Manage** to add or configure a route.
- Step 10** Click **OK**.
- Step 11** Specify a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60.
- Step 12** Click **Apply**.

Configure IPv4 Family Filtering Settings

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

Procedure

- Step 1** Choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Filtering**.
The Define filters for BGP updates pane is displayed.
- Step 3** Click **Add**.
The Add Filter pane is displayed.
- Step 4** Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.
- Step 5** Choose a standard access list from the Access List drop-down list. Click **Manage** to add a new ACL.
- Step 6** For outbound filters, you can optionally specify what types of route are distributed.
- a) Choose an option from the Protocol drop-down list.
You can choose a routing protocol, such as **BGP**, **EIGRP**, **OSPF**, or **RIP**.
Choose **Connected** to filter on peers and networks learned through connected routes.
Choose **Static** to filter on peers and networks learned through static routes.
 - b) If you chose BGP, EIGRP, or OSPF, also choose the **Process ID** for that protocol.
- Step 7** Click **OK**.

Step 8 Click **Apply**.

Configure IPv4 Family BGP Neighbor Settings

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

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing BGP > IPv4 Family**.
- Step 2** Click **Neighbor**.
- Step 3** Click **Add**.
- Step 4** Click **General** in the left pane.
- Step 5** Enter a BGP neighbor IP address in the **IP Address** field. This IP address is added to the BGP neighbor table.
- Step 6** Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.
- Step 7** (Optional) Enter a description for the BGP neighbor in the **Description** field.
- Step 8** (Optional) Check the **Shutdown neighbor administratively** check box to disable a neighbor or peer group.
- Step 9** (Optional) Check the **Enable address family** check box to enable communication with the BGP neighbor.
- Step 10** (Optional) Check the **Global Restart Functionality for this peer** check box to enable or disable the Border Gateway Protocol (BGP) graceful restart capability for a ASA neighbor or peer group.
- Note** This option is enabled when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.
- Step 11** Click **Filtering** in the left pane.
- Step 12** (Optional) Choose the appropriate incoming or outgoing access control list in the Filter routes using an access list area, to distribute BGP neighbor information. Click **Manage** to add an ACL and ACEs as required.
- Step 13** (Optional) Choose the appropriate incoming or outgoing route maps in the Filter routes using a route map area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.
- Step 14** (Optional) Choose the appropriate incoming or outgoing prefix list in the Filter routes using a prefix list area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.
- Step 15** (Optional) Choose the appropriate incoming or outgoing AS path filter in the Filter routes using AS path filter area, to distribute BGP neighbor information. Click **Manage** to configure AS path filters.
- Step 16** (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box to control the number of prefixes that can be received from a neighbor.
- Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum prefixes** field.
 - Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold level** field. Valid values are integers between 1 to 100. The default value is 75.
 - (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:
 - Click **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field.

- Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.

- Step 17** Click **Routes** in the left pane.
- Step 18** Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field.
- Step 19** (Optional) Check the **Generate Default** route check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route.
- Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the Route map drop-down list. Click **Manage** to add and configure a route map.
- Step 20** (Optional) To add conditionally advertised routes do the following:
- a) Click **Add** in the Conditionally Advertised Routes section.
 - b) Choose the route map that will be advertised from the Advertise Map drop-down list, if the conditions of the exist map or the non-exist map are met.
 - c) Do one of the following:
 - Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - d) Click **Ok**.
- Step 21** (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.
- Step 22** Click **Timers** in the left pane.
- Step 23** (Optional) Check the Set timers for the BGP peer check box to set the keepalive frequency, hold time and minimum hold time.
- Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor, in the **Keepalive frequency** field. Valid values are between 0 and 65535. The default value is 60 seconds.
 - Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the **Hold time** field. The default value is 180 seconds.
 - (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the Min Hold time field.
- Note** A hold time of less than 20 seconds increases the possibility of peer flapping.
- Step 24** Click **Advanced** in the left pane.
- Step 25** (Optional) Check the Enable Authentication check box to enable MD5 authentication on a TCP connection between two BGP peers.
- Choose an encryption type from the Encryption Type drop-down list.
 - Enter a password in the **Password** field. Reenter the password in the **Confirm Password** field.
- The password is case-sensitive and can be up to 25 characters long, when the service password-encryption command is enabled and up to 81 characters long, when the service password-encryption command is

not enabled. The string can contain any alphanumeric characters, including spaces. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

Step 26 (Optional) Check the **Send Community Attribute to this neighbor** check box.

Step 27 (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

Step 28 Do one of the following:

- Click **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected.
- (Optional) Enter the time-to-live in the **TTL hops** field. Valid values are between 1 and 255.
- (Optional) Check the **Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.
- Click **Limit number of TTL hops to neighbor** to enable you to secure a BGP peering session.
- Enter the maximum number of hops that separate eBGP peers in the **TTL hops** field. Valid values are between 1 and 254.

Step 29 (Optional) Enter a weight for the BGP neighbor connection in the **Weight** field.

Step 30 Choose the BGP version that the ASA will accept from the **BGP version** drop-down list.

Note 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.

Step 31 (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.

Step 32 Choose the TCP connection mode from the **TCP transport mode** drop-down list.

Step 33 Click **Migration** in the left pane

Step 34 (Optional) Check the **Customize the AS number** for routes received from the neighbor check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.

- Enter the local autonomous system number in the **Local AS Number** field. Valid values are between 1 and 65535.
- (Optional) Check the **Do not prepend local AS number** for routes received from neighbor check box. The local AS number will not be prepended to any routes received from eBGP peer.
- (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box. The AS number from the local routing process is not prepended.
- (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box.

Step 35 Click **OK**.

Step 36 Click **Apply**.

Configure IPv4 Network Settings

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

Procedure

-
- | | |
|---------------|---|
| Step 1 | In ASDM, choose Configuration > Device Setup > Routing > BGP > IPv4 Family . |
| Step 2 | Click Networks .

The Define networks to be advertised by the BGP routing process configuration pane appears. |
| Step 3 | Click Add .

The Add Network pane is displayed. |
| Step 4 | Specify the network that BGP will advertise in the Address field. |
| Step 5 | (Optional) Choose a network or subnetwork mask from the Netmask drop-down list. |
| Step 6 | Choose a route map that should be examined to filter the networks to be advertised from the Route Map drop-down list. Click Manage to configure or add a route map. |
| Step 7 | Click OK . |
| Step 8 | Click Apply . |
-

Configure IPv4 Redistribution Settings

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

Procedure

-
- | | |
|---------------|---|
| Step 1 | In ASDM, choose Configuration > Device Setup > Routing > BGP > IPv4 Family > . |
| Step 2 | Click Redistribution .

The Redistribution pane is displayed. |
| Step 3 | Click Add .

The Add Redistribution pane is displayed. |
| Step 4 | Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list. |
| Step 5 | Choose a process ID for the source protocol from the Process ID drop-down list. |
| Step 6 | (Optional) Enter a metric for the redistributed route in the Metric field. |
| Step 7 | Choose a route map that should be examined to filter the networks to be redistributed from the Route Map drop-down list. Click Manage to configure or add a route map. |
| Step 8 | Check one or more of the Internal , External and NSSA External Match check boxes to redistribute routes from an OSPF network.

This step is only applicable for redistribution from OSPF networks. |

- Step 9** Click **OK**.
- Step 10** Click **Apply**.

Configure IPv4 Route Injection Settings

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

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family > .**
- Step 2** Click **Route Injection**.
The Route Injection pane is displayed.
- Step 3** Click **Add**.
The Add Conditionally injected route pane is displayed.
- Step 4** Choose the route map that specifies the prefixes to inject into the local BGP routing table from the **Inject Map** drop-down list.
- Step 5** Choose the route map containing the prefixes that the BGP speaker will track from the **Exist Map** drop-down list.
- Step 6** Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.

Configure IPv6 Address Family Settings

The IPv6 settings for BGP can be set up from the IPv6 family option within the BGP configuration setup. The IPv6 family section includes subsections for General settings, Aggregate address settings and Neighbor settings. Each of these subsections enable you to customize parameters specific to the IPv6 family.

This section describes how to customize the BGP IPv6 family settings.

Configure IPv6 Family General Settings

This section describes the steps required to configure the general IPv6 settings.

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **General**.
The general IPv6 family BGP parameters configuration pane is displayed.

- Step 3** Specify external, internal and local distances in the **Administrative Route Distances** area.
- Step 4** (Optional) Check the **Generate Default Route** check box to configure a BGP routing process to distribute a default route (network 0.0.0.0).
- Step 5** (Optional) Check the **Advertise inactive routes** check box to advertise routes that are not installed in the routing information base (RIB).
- Step 6** (Optional) Check the **Redistribute iBGP into an IGP** check box to configure iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
- Step 7** (Optional) Enter a scanning interval (in seconds) for BGP routers for next-hop validation in the **Scanning Interval** field. Valid values are from 5 to 60 seconds.
- Step 8** (Optional) Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table in the **Number of paths** field.
- Step 9** (Optional) Check the **iBGP multipaths** check box and specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table in the **Number of paths** field.
- Step 10** Click **Apply**.

Configure IPv6 Family Aggregate Address Settings

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

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Aggregate Address**.
The Aggregate Address parameters configuration pane is displayed.
- Step 3** Click **Add**.
The Add Aggregate Address pane is displayed.
- Step 4** Specify an IPv6 address in the **IPv6/Address Mask** field. Alternately, browse to add a network object.
- Step 5** Check the **Generate autonomous system set path information** check box to generate autonomous system set path information. The path advertised for this route will be an AS_SET consisting of all elements contained in all paths that are being summarized.

Note Do not use this form of the aggregate-address command when aggregating many paths, because this route must be continually withdrawn and updated as autonomous system path reachability information for the summarized routes changes.
- Step 6** Check the **Filters all more-specific routes from the updates** check box to filter all more-specific routes from updates. This not only creates the aggregate route but also suppresses advertisements of more-specific routes to all neighbors.
- Step 7** Choose a route-map from the **Attribute Map** drop-down list. Click **Manage** to add or configure a route map. This allows attributes of the aggregate route to be changed.
- Step 8** Choose a route-map from the **Advertise Map** drop-down list. Click **Manage** to add or configure a route. This selects specific routes that will be used to build different components of the aggregate route.
- Step 9** Choose a route-map from the **Suppress Map** drop-down list. Click **Manage** to add or configure a route. This creates the aggregate route but suppresses advertisement of specified routes.

- Step 10** Click **OK** .
- Step 11** Specify a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. This specifies the interval at which the routes will be aggregated. The default value is 30 seconds.
- Step 12** Click **Apply** .

Configure IPv6 Family BGP Neighbor Settings

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

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Neighbor** .
- Step 3** Click **Add** .
- Step 4** Click **General** in the left pane.
- Step 5** Enter a BGP neighbor IPv6 address in the **IPv6 Address** field. This IPv6 address is added to the BGP neighbor table.
- Step 6** Enter the autonomous system to which the BGP neighbor belongs in the **Remote AS** field.
- Step 7** (Optional) Enter a description for the BGP neighbor in the **Description** field.
- Step 8** (Optional) Check the **Shutdown neighbor administratively** check box to disable a neighbor or peer group.
- Step 9** (Optional) Check the **Enable address family** check box to enable communication with the BGP neighbor.
- Step 10** Click **Filtering** in the left pane.
- Step 11** (Optional) Choose the appropriate incoming or outgoing route maps in the **Filter routes using a route map** area, to apply a route map to incoming or outgoing routes. Click **Manage** to configure a route map.
- Step 12** (Optional) Choose the appropriate incoming or outgoing prefix list in the **Filter routes using a prefix list** area, to distribute BGP neighbor information. Click **Manage** to configure prefix lists.
- Step 13** (Optional) Choose the appropriate incoming or outgoing AS path filter in the **Filter routes using AS path filter** area, to distribute BGP neighbor information. Click **Manage** to configure AS path filters.
- Step 14** (Optional) Check the **Limit the number of prefixes allowed from the neighbor** check box to control the number of prefixes that can be received from a neighbor.
- Step 15** Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum prefixes** field.
- Step 16** Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold level** field. Valid values are integers between 1 to 100. The default value is 75.
- Step 17** (Optional) Check the **Control prefixes received from a peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:
- Click **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the **Restart interval** field.
 - Click **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
- Step 18** Click **Routes** in the left pane.

- Step 19** Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field.
- Step 20** (Optional) Check the **Generate Default route** check box to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route.
- Step 21** Choose the route map that allows the route 0.0.0.0 to be injected conditionally from the **Route map** drop-down list. Click **Manage** to add and configure a route map.
- Step 22** (Optional) To add conditionally advertised routes do the following:
- Click **Add** in the Conditionally Advertised Routes section.
 - Choose the route map that will be advertised from the **Advertise Map** drop-down list, if the conditions of the exist map or the non-exist map are met.
 - Do one of the following:
 - Click **Exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - Click **Non-exist Map** and choose a route map. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
 - Click **Ok**.
- Step 23** (Optional) Check the **Remove private autonomous system (AS) numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.
- Step 24** Click **Timers** in the left pane.
- Step 25** (Optional) Check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time and minimum hold time.
- Step 26** Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor in the **Keepalive frequency** field. Valid values are between 0 and 65535. The default value is 60 seconds.
- Step 27** Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the **Hold time** field. The default value is 180 seconds.
- Step 28** (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead, in the **Min Hold time** field.
- Note** A hold time of less than 20 seconds increases the possibility of peer flapping.
- Step 29** Click **Advanced** in the left pane.
- Step 30** (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.
- Step 31** Choose an encryption type from the **Encryption Type** drop-down list.
- Step 32** Enter a password in the **Password** field. Reenter the password in the **Confirm Password** field.
- The password is case-sensitive and can be up to 25 characters long, when the service password-encryption command is enabled and up to 81 characters long, when the service password-encryption command is not enabled. The string can contain any alphanumeric characters, including spaces. You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.
- Step 33** (Optional) Check the **Send Community Attribute to this neighbor** check box.
- Step 34** (Optional) Check the **Use ASA as next hop for neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.
- Step 35** Do one of the following:

- Click **Allow connections with neighbor that is not directly connected** to accept and attempt BGP connections to external peers residing on networks that are not directly connected.
 - (Optional) Enter the time-to-live in the **TTL hops** field. Valid values are between 1 and 255.
 - (Optional) Check the **Disable connection verification** check box to disable connection verification to establish an eBGP peering session with a single-hop peer that uses a loopback interface.
- Click **Limit number of TTL hops to neighbor** to enable you to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254

Step 36 (Optional) Enter a weight for the BGP neighbor connection in the **Weight** field.

Step 37 Choose the BGP version that the ASA will accept from the BGP version drop-down list.

Note 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.

Step 38 (Optional) Check the **TCP Path MTU Discovery** check box to enable a TCP transport session for a BGP session.

Step 39 Choose the **TCP connection mode** from the TCP transport mode drop-down list.

Step 40 Click **Migration** in the left pane

Step 41 (Optional) Check the **Customize the AS number for routes received from the neighbor** check box to customize the AS_PATH attribute for routes received from an eBGP neighbor.

- Enter the local autonomous system number in the Local AS Number field. Valid values are between 1 and 65535.
- (Optional) Check the Do not prepend local AS number for routes received from neighbor check box. The local AS number will not be prepended to any routes received from eBGP peer.
- (Optional) Check the Replace real AS number with local AS number in routes received from neighbor check box. The AS number from the local routing process is not prepended.
- (Optional) Check the Accept either real AS number or local AS number in routes received from neighbor check box.

Step 42 Click **OK**.

Step 43 Click **Apply**.

Configure IPv6 Network Settings

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

Procedure

Step 1 In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.

Step 2 Click **Networks**.

The Define the networks to be advertised by the BGP routing process configuration pane appears.

- Step 3** Click **Add**.
The Add Network pane is displayed.
- Step 4** (Optional) In the Prefix Name field, specify the prefix name for the DHCPv6 Prefix Delegation client (see [Enable the IPv6 Prefix Delegation Client](#)).
- Step 5** In the **IPv6 Address/mask** field, specify the network that BGP will advertise.
If you specify the **Prefix Name**, enter the subnet prefix and mask; the advertised network consists of the delegated prefix + the subnet prefix.
- Step 6** Choose a route map that should be examined to filter the networks to be advertised from the **Route Map** drop-down list. Optionally, click **Manage** to configure or add a route map.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.

Configure IPv6 Redistribution Settings

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

Procedure

- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv6 Family**.
- Step 2** Click **Redistribution**.
- Step 3** Click **Add**.
The Add Redistribution pane is displayed.
- Step 4** In the Source Protocol drop-down list, choose the protocol from which you want to redistribute routes into the BGP domain.
- Step 5** In the Process ID drop-down list, choose a process ID for the source protocol. This is available only for OSPF source protocol.
- Step 6** (Optional) In the Metric field, enter a metric for the redistributed route.
- Step 7** In the Route Map drop-down list, choose a route map that should be examined to filter the networks to be redistributed. Click **Manage** to configure or add a route map.
- Step 8** Check one or more of the Match check-boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.
This step is only applicable for redistribution from OSPF networks.
- Step 9** Click **OK**.
- Step 10** Click **Apply**.

Configure IPv6 Route Injection Settings

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

Procedure

-
- Step 1** In ASDM, choose **Configuration > Device Setup > Routing > BGP > IPv4 Family**.
- Step 2** Click **Route Injection**.
- Step 3** Click **Add**.
- The Add Conditionally injected route pane is displayed.
- Step 4** In the **Inject Map** drop-down list, choose the route map that specifies the prefixes to inject into the local BGP routing table.
- Step 5** In the **Exist Map** drop-down list, choose the route map containing the prefixes that the BGP speaker will track.
- Step 6** Check the **Injected routes will inherit the attributes of the aggregate route** check box to configure the injected route to inherit attributes of the aggregate route.
- Step 7** Click **OK**.
- Step 8** Click **Apply**.
-

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 various BGP routing statistics, enter one of the following commands:



Note To disable BGP Log messages, enter the **no bgp log-neighbor-changes** command in the router configuration mode. This disables the logging of neighbor change messages. Enter this command in router configuration mode for the BGP routing process. By default, neighbor changes are logged.

- **Monitoring > Routing > BGP Neighbors**

Each row represents one BGP neighbor. For each neighbor, the list includes the IP address, the AS number, the router ID, the state (active, idle and so on), the uptime, graceful restart capability, the restart time and the stalepath time.

- **Monitoring > Routing > BGP Routes**

Each row represents one BGP route. For each route, the list includes the status code, IP address, the next hop address, the route metric, the local preference values, the weight and the path.

History for BGP

Table 1: Feature History for BGP

Feature Name	Platform Releases	Feature Information
BGP Support	9.2(1)	Support was added for routing data, performing authentication, and redistributing and monitoring routing information using the Border Gateway Protocol. We introduced the following screens: Configuration > Device Setup > Routing > BGP Monitoring > Routing > BGP Neighbors, Monitoring > Routing > BGP Routes We modified the following screens: Configuration > Device Setup > Routing > Static Routes> Add > Add Static Route Configuration > Device Setup > Routing > Route Maps> Add > Add Route Map
BGP support for ASA clustering	9.3(1)	We added support for L2 and L3 clustering. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > General
BGP support for nonstop forwarding	9.3(1)	We added support for Nonstop Forwarding. We modified the following screens: Configuration > Device Setup > Routing > BGP > General, Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor, Monitoring > Routing > BGP Neighbors
BGP support for advertised maps	9.3(1)	We added support for BGPv4 advertised map. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv4 Family > Neighbor > Add BGP Neighbor > Routes
BGP support for IPv6	9.3(2)	We added support for IPv6. We introduced the following screen: Configuration > Device Setup > Routing > BGP > IPv6 Family
IPv6 network advertisement for delegated prefixes	9.6(2)	The ASA now supports the DHCPv6 Prefix Delegation client. The ASA obtains delegated prefix(es) from a DHCPv6 server. The ASA can then use these prefixes to configure other ASA interface addresses so that Stateless Address Auto Configuration (SLAAC) clients can autoconfigure IPv6 addresses on the same network. You can configure the BGP router to advertise these prefixes. We modified the following screen: Configuration > Device Setup > Routing > BGP > IPv6 Family > Networks

