

BGP

This section describes how to configure the Firewall Threat Defense to route data, perform authentication, and redistribute routing information using the Border Gateway Protocol (BGP).

- BGP, on page 1
- Requirements and prerequisites for BGP, on page 5
- Guidelines for BGP, on page 5
- Configure BGP, on page 6

BGP

Border Gateway Protocol (BGP) is an inter- and intra-autonomous system routing protocol that

- operates within and between autonomous systems (networks or groups of networks under common administration with common routing policies),
- exchanges routing information for the Internet, and
- serves as the protocol used between Internet service providers (ISPs).

BGP usage scenarios

BGP usage scenarios are network configurations that

- enable ISPs to exchange customer and ISP routes between autonomous systems using External BGP (EBGP),
- allow service providers to exchange routes within an autonomous system using Interior BGP (IBGP), and
- support IPv6 prefix routing over IPv6 networks.

BGP implementation contexts

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.

BGP path selection

BGP path selection is a routing process that selects only one path as the best path when BGP receives multiple advertisements for the same route from different sources, puts the selected path in the IP routing table, and propagates the path to its neighbors using specific criteria in order of precedence.

BGP path selection criteria

BGP uses these criteria, in the order presented, to select a path for a destination:

1. If the path specifies a next hop that is inaccessible, drop the update. Otherwise, prefer the path with the largest weight.
2. 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.
3. 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).
4. 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.
5. If the paths are still the same, prefer the path through the closest IGP neighbor.
6. Determine if multiple paths require installation in the routing table for BGP multipath, on page 2.
7. If both paths are external, prefer the path that was received first (the oldest one).
8. 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.
9. Prefer the path that comes from the lowest neighbor address.

BGP multipath

BGP multipath is a BGP feature that

- allows installation into the IP routing table of multiple equal-cost BGP paths to the same destination prefix,
- enables traffic to the destination prefix to be shared across all installed paths, and
- does not affect best path selection algorithm.

BGP multipath requirements

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.

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 in 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.

Routing tables

A routing table is a network database that

- records paths to destination networks based on exchanged routing information
- updates dynamically as neighbors share new or modified routing routes, and
- determines the optimal path to a destination without requiring periodic updates.

BGP attributes for route selection

Routes learned via BGP have several attributes that determine the best path when multiple routes to a destination exist. The most important BGP attributes include:

- Weight: A Cisco-defined attribute that is local to a router. It is not advertised to neighboring routers. Among multiple routes to the same destination, the route with the highest weight is preferred.
- Local preference: This attribute selects an exit point from the local autonomous system (AS) and is shared throughout the local AS. If multiple exit points exist, the one with the highest local preference attribute is chosen 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. However, when the next hop is in the same subnet the peering address of the EBGP peer, the next hop is not modified. This behavior is referred to as the third-party next hop.

  Use the **next-hop-self** command when redistributing VPN-advertised routes to IBGP as peers to ensure that the routes are redistributed with the correct next hop IP.
- 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.

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 device in performing loop checks and to avoid delays in the existing outgoing update tasks.

Requirements and prerequisites for BGP

This reference provides the requirements and prerequisites necessary for configuring BGP on supported devices.

Model support

Firewall Threat Defense

Firewall Threat Defense Virtual

Guidelines for BGP

Context mode guidelines

Consider these context mode requirements when implementing BGP.

- Supported in single and multiple context modes.
- 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.

Additional guidelines

Consider these additional operational requirements and limitations.

- For BGP, the next hop IP address for the routes is the network IP address and not 0.0.0.0.
- The system does not add a route entry for the IP address received over PPPoE in the CP route table. BGP always looks into the CP route table for initiating the TCP session, hence BGP does not form TCP session.

  Thus, BGP over PPPoE is not supported.
- BGP is not supported on management-only or BVI interfaces.
- 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.

- BGP with PATH MTU (PMTU) can cause adjacency flaps if MTU discovery fails, especially with ECMP routing. Hence, be cautious while using BGP, PMTU, and ECMP as packet drops can occur if MTU discovery fails due to any reason.
- 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.
- Enable the graceful timeout feature when you use BGP on Firewall Threat Defense devices that are configured in high availability mode. Enabling this feature reduces packet drops that occur due to issues related to BGP convergence.
- When you configure a route-based site-to-site VPN using static or dynamic VTI interfaces, ensure that the TTL hop count is greater than one if you use BGP as the routing protocol.
- On Firewall Threat Defense devices, changing a large number of BGP neighbors from iBGP to eBGP in a single deployment can fail. During deployment, Cloud-Delivered Firewall Management Center can remove and re-create the same peer in immediate succession, and the deployment can fail if the previous peer deletion is still in progress. To avoid this issue, split the changes across multiple deployments. Remove or modify a subset of neighbors, deploy the configuration, and then apply additional neighbor changes in subsequent deployments.

Configure BGP

Configure BGP routing protocol to establish peering relationships and control route advertisement between autonomous systems.

BGP configuration involves multiple components including basic settings, neighbor relationships, route filtering, and redistribution policies.

Procedure

Step 1 Configure basic BGP settings, on page 7

Step 2 Configure BGP general settings, on page 9

Step 3 Configure BGP Neighbor Settings, on page 11

Step 4 Configure BGP aggregate address settings, on page 15

Step 5 Configure BGPv4 filtering settings, on page 16

Note The Filtering section is applicable only to IPv4 settings

Step 6 Configure BGP network settings, on page 17

Step 7 Configure BGP redistribution settings, on page 18

Step 8 Configure BGP route injection settings, on page 19

Step 9 Configure BGP route import and export, on page 19

BGP is configured with all necessary settings, enabling dynamic routing and proper route exchange between autonomous systems.

Configure basic BGP settings

Enable and customize BGP settings to control routing behavior and improve network operations.

You can set many basic settings for BGP.

For a device using virtual routing, the basic settings described in this section must be configured in the **BGP** page under **General Settings**. For more information, refer to Modifications to the Cloud-Delivered Firewall Management Center web interface routing page.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select the **Routing** tab and do the following:

a) (For a virtual-router-aware device) Under **General Settings**, click **BGP**.

b) Check the **Enable BGP** check box to enable the BGP routing process.

c) In the **AS Number** field, enter the autonomous system (AS) number for the local BGP process. The AS number internally includes multiple autonomous numbers. Valid values are from 1 to 4294967295 or from 1.0 to 65535.65535 when a dot notation is used. The AS number is a uniquely assigned value that identifies each network on the Internet.

d) In the **Router Id** drop-down list, choose **Automatic** or **Manual** (appears for non-cluster and a cluster in spanned etherchannel mode) or **Cluster Pool** (appears for a cluster in individual interface mode). If you choose **Automatic**, the highest-level IP address on the Firewall Threat Defense device is used as the router ID. If you choose **Manual**, enter the IP address in the **IP Address** field. If you choose **Cluster Pool**, enter the cluster pool value in the **Cluster Pool** field. For more information on creating the cluster pool address, refer to Configure address pools..

e) To use a fixed router ID, choose **Manual** and enter an IPv4 address in the **IP Address** field. The default value is **Automatic**. For a virtual router-aware device, you can override the router ID settings in the **Virtual Routers** > **BGP** page.

Step 3 (Optional) Edit the various BGP settings, starting with **General**. The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network. Click **Edit** (✎) to edit the settings in the group:

a) Enter a **Scanning Interval** for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.

b) In the **Number of AS numbers in AS_PATH attribute of received routes** field, enter the maximum number of AS numbers allowed in the AS_PATH attribute of received routes. An AS_PATH attribute is a sequence of intermediate AS numbers between source and destination routers that form a directed route for packets to travel. Valid values are between 1 and 254. The default value is **None**.

c) 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. This is enabled by default.

d) Check the **Use TCP path MTU discovery** check box to use the Path MTU determining technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation. This is enabled by default.

e) Check the **Reset session upon failover** check box to reset the external BGP session immediately upon link failure. This is enabled by default.

f) Check the **Enforce the 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 misconfigured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system. This is enabled by default.

g) Check the **Use dot notation for AS number** 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 1 to 65535 are represented as decimal numbers, and AS numbers greater than 65535 are represented using dot notation. This is disabled by default.

Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

h) Check the **BGP community new format (AA:NN)** check box to configure BGP to display community values in the AA:NN format, where AA represents the autonomous system number, and NN represents the community number, each being 2 bytes long. This format conforms to RFC-1997. By default, BGP communities are displayed as a single 32-bit decimal number.

i) Click **OK**.

Step 4 (Optional) Edit the **Best Path Selection** section:

a) Enter a value for **Default local preference** 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.

b) 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. This is disabled by default.

c) 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. This is disabled by default.

d) Check the **Pick the best-MED path among paths advertised by neighbor AS** check box to enable MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path. This is disabled by default.

e) Check the **Treat missing MED as the least preferred path** check box to consider the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred. This is disabled by default.

f) Click **OK**.

Step 5 (Optional) Edit the **Neighbor Timers** section:

a) Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message in the **Keepalive interval** field. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.

b) Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Hold time** field. The default value is 180 seconds. Specify a value from 0 to 65535.

c) (Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured in the **Min Hold time** field. Specify a value from 3 to 65535.

Note A hold time of less than 20 seconds increases the possibility of peer flapping.

d) Click **OK**.

Step 6 In the **Next Hop** section, optionally check the **Address tracking** check box to enable BGP next hop address tracking and enter the **Delay Interval** between checks on updated next-hop routes installed in the routing table. Click **OK**.

Note The **Next Hop** section is applicable only to IPv4 settings.

Step 7 (Optional) Edit the **Graceful Restart** section:

Note This section is available only when the Firewall Threat Defense device is in failover or spanned cluster mode. This is done so that there is no drop in packets in the traffic flow, when one of the devices in the failover setup fails.

a) Check the **Enable Graceful Restart** check box to enable Firewall Threat Defense peers to avoid a routing flap following a switchover.
b) Specify the time duration that Firewall Threat Defense peers will wait to delete stale routes before a BGP open message is received in the **Restart Time** field. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
c) Enter the time duration that the Firewall Threat Defense will wait before deleting stale routes after an end of record (EOR) message is received from the restarting Firewall Threat Defense in the **Stalepath Time** field. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.
d) Click **OK** and then click **Save**.

Step 8 To view the BGP basic settings, from the virtual routers drop-down, select the desired router, and then click **BGP**.

This page displays the basic settings that are configured in the **Settings** page. You can edit the router ID settings on this page.

To edit the router ID settings, modify the IP address in the **IP Address** fields. The modified value overrides the router ID settings that were configured on the **BGP** page under **General Settings**.

The device is configured with basic BGP settings, enabling BGP routing, specifying the AS number, router ID, and any custom parameters you set.

Configure BGP general settings

Configure route maps, administrative route distances, synchronization, next-hop, and packet forwarding for BGP to optimize routing performance in your network.

The defaults for these settings are appropriate in most cases, but you can adjust them to fit the needs of your network.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select the **Routing** tab.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down, select the virtual router for which you are configuring BGP.

Step 4 Choose **BGP** > **IPv4** or **IPv6**.

Step 5 In the **General** tab, update the following sections:

a) In the **Setting** section, enter or select a **RouteMap** object and click **OK**.

Note The **RouteMap** field is applicable only to IPv4 settings.

b) In the **Administrative Route Distances** section, update the following as required, and click **OK**:

- **External**: Enter the 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 default value is 20.
- **Internal**: Enter 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 default value is 200.
- **Local**: Enter administrative distance for local BGP routes. Local routes are routes originated by the local BGP process, such as networks specified with the network router configuration command or routes redistributed into BGP from another routing process. The range of values for this argument is from 1 to 255. The default value is 200.

c) In the **Routes and Synchronization** section, update the following as required, and click **OK**:

- (Optional) **Generate default routes**: Check the check box of this option to configure default-information originate.
- (Optional) **Summarize subnet routes into network level routes**: Check the check box of this to configure automatic summarization of subnet routes into network-level routes. This check box is applicable only to IPv4 settings.
- (Optional) **Advertise inactive routes**: Check the check box of this to advertise routes that are not installed in the routing information base (RIB).
- (Optional) **Synchronize between BGP and IGP systems**: Check the check box of this to enable synchronization between BGP and your Interior Gateway Protocol (IGP) system. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems.
- (Optional) **Redistribute IBGP into IGP**: Check the check box of this to configure IBGP redistribution into an interior gateway protocol (IGP), such as OSPF.

d) In the **Forward Packets Over Multiple Paths** section, update the following as required and click **OK**:

- (Optional) **Number of Paths**: Enter the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.
- (Optional) **IBGP Number of Paths**: Enter the maximum number of parallel internal Border Gateway Protocol (IBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Step 6 Click **Save**.

BGP general settings are configured with your specified route maps, administrative distances, synchronization options, and multi-path forwarding parameters.

Configure BGP Neighbor Settings

Define and customize BGP neighbor connections on your device for efficient routing and advanced peer management.

A BGP router must connect with each of its peers before exchanging updates. These peers are called BGP neighbors. Perform this procedure to define BGP IPv4 or IPv6 neighbors and neighbor settings.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device. Select the **Routing** tab.

(For a virtual router-aware device) From the virtual routers drop-down list, choose the virtual router for which you are configuring BGP.

Step 2 Choose **BGP** > **IPv4** or **IPv6**. Click **Neighbor**.

Step 3 Click **Add** to define BGP neighbors and neighbor settings.

a) Enter the BGP neighbor **IP address**. This IP address is added to the BGP neighbor table. When you are configuring BGP IPv6 on static VTI, enter the virtual tunnel IP address of the neighbor. Choose the BGP neighbor **Interface**.

Note The **Interface** field is only applicable to IPv6 settings.

b) Enter the autonomous system to which the BGP neighbor belongs, in the **Remote AS** field. Check the **Enabled address** check box to enable communication with this BGP neighbor. Further neighbor settings will be configured only if the **Enabled address** check box is selected.

c) (Use this option only when the routing does not result in a loop.) Check the **AS Override** check box to enable overriding of the AS number of the originating router with the AS number of the sending BGP router.

Note Loop prevention in BGP is achieved by verifying the AS number in the AS path. BGP rejects route updates when the AS path attribute contains its own AS number. However, in circumstances where the same ASN is used by discontiguous network segments, to establish end-to-end reachability, you can opt to override the AS number with that of the sending BGP router.

d) (Optional) Check the **Shutdown administratively** check box to disable a neighbor or peer group.
e) (Optional) Check the **Configure graceful restart (failover / spanned mode)** check box to enable configuration of the BGP graceful restart capability for this neighbor. After selecting this option, you must check the **Enable graceful restart** check box to specify whether graceful restart should be enabled or disabled for this neighbor.

Note
- The graceful restart is enabled only when the device is in HA mode or when L2 cluster (all nodes from the same network) is configured.
- The graceful restart option for BGPv6 is enabled only on Firewall Threat Defense Version 7.3+.
- If you configure graceful restart only at General Settings and not at BGP IPv6, the global General Settings configuration persists.
- If you configure graceful restart at General Settings and also at BGP IPv6, the global General Settings configuration is overridden by the BGP IPv6 configuration settings.

f) (Optional) To enable configuration of the BFD support for BGP, from the **BFD Fallover** drop-down list, choose the BFD type—single-hop, multi-hop, or auto-detect-hop. This selection registers the BGP neighbor to receive forwarding path detection failure messages from BFD. Choose None if you do not want to have BFD support.
g) (Optional) Enter a **Description** for the BGP neighbor.
h) (Optional) From the **Update Source** drop-down list, choose an interface to source the BGP packets.

You can choose a loopback address as this interface to overcome path failures. You can also choose any physical, port-channel, or a sub-interface.

Step 4 (Optional) In **Filtering Routes**, use access lists, route maps, prefix lists and AS path filters as required, to distribute BGP Neighbor information. Update the following sections:

a) Choose or select the appropriate incoming or outgoing **Access List** to distribute BGP neighbor information.

Note Access lists are only applicable to IPv4 settings.

b) Choose or select the appropriate incoming or outgoing **Route Maps** to apply a route map to incoming or outgoing routes.
c) Choose or select the appropriate incoming or outgoing **Prefix List** to distribute BGP neighbor information.
d) Choose or select the appropriate incoming or outgoing **AS path filter** to distribute BGP neighbor information.

e) Check the check box of **Limit the number of prefixes allowed from the neighbor** 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 and 100. The default value is 75.

f) Check the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following:

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

g) Click **OK**.

Step 5 (Optional) In **Routes**, specify miscellaneous Neighbor route parameter. Proceed to update the following:

a) Enter the minimum interval (in seconds) between the sending of BGP routing updates in the **Advertisement Interval** field. Valid values are between 1 and 600.

b) Check the **Remove private AS numbers from outbound routing updates** check box to exclude the private AS numbers from being advertised on outbound routes.

c) Check the **Generate default routes** 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. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the **Route map** field.

d) To add conditionally advertised routes, click Add Row +. In the Add Advertised Route dialog box, do the following:

1. Add or choose a route map in the **Advertise Map** field, that will be advertised if the conditions of the exist map or the non-exist map are met.
2. Click **Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
3. Click **Non-Exist Map** and choose a route map from the Route Map Object Selector. This route map is compared with the routes in the BGP table, to determine whether the advertise map route is advertised.
4. Click **OK**.

Step 6 In **Timers**, check the **Set timers for the BGP peer** check box to set the keepalive frequency, hold time, and minimum hold time.

- **Keepalive interval**: Enter the frequency (in seconds) with which Firewall Threat Defense sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
- **Hold time**: Enter the interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.

- **Min hold time**: (Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that Firewall Threat Defense declares a peer dead. Valid values are between 3 and 65535. The default value is 3 seconds.

Note A hold time of less than 20 seconds increases the possibility of peer flapping.

Step 7 In **Advanced**, update the following:

a) (Optional) Check the **Enable Authentication** check box to enable MD5 authentication on a TCP connection between two BGP peers.

1. Choose an encryption type from the **Enable Encryption** drop-down list.
2. 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.

Note You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.

b) (Optional) Select the **Send Community attribute to this neighbor** check box to specify that community attributes should be sent to the BGP neighbor

c) (Optional) Select the **Use FTD as next hop for this neighbor** check box to configure the router as the next-hop for a BGP speaking neighbor or peer group.

d) Select the **Disable Connection Verification** check box to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.

e) Select **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. However, for eBGP multi-hop neighborship configuration, you must enter the **TTL hops** value. Alternatively, select **Limited number of TTL hops to neighbor**, 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.

f) (Optional) Select the **Use TCP MTU path discovery** check box to enable a TCP transport session for a BGP session.

g) Choose the TCP connection mode from the **TCP Transport Mode** drop-down list. Options are Default, Active, or Passive.

h) (Optional) Enter a **Weight** for the BGP neighbor connection.

i) Select the **BGP Version** that Firewall Threat Defense will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.

Step 8 Update **Migration**, only if AS migration is considered.

Note The AS migration customization should be removed after transition has been completed.

a) (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.
b) Enter the local autonomous system number in the **Local AS number** field. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
c) (Optional) Check the **Do not prepend local AS number to routes received from neighbor** check box to prevent the local AS number from being prepended to any routes received from eBGP peer.
d) (Optional) Check the **Replace real AS number with local AS number in routes received from neighbor** check box to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
e) (Optional) Check the **Accept either real AS number or local AS number in routes received from neighbor** check box to configure the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Step 9 Click **OK**, and then click **Save**.

The BGP neighbor and its associated advanced settings are configured. The device can now establish and manage BGP sessions as specified, ensuring effective routing and network reachability.

Configure BGP aggregate address settings

Route aggregation combines the attributes of several different routes so that only a single route is advertised, reducing the number of routes that need to be advertised.

BGP neighbors store and exchange routing information and the amount of routing information increases as more BGP speakers are configured. Route aggregation is the process of combining the attributes of several different routes so that only a single route is advertised. Aggregate prefixes use the classless interdomain routing (CIDR) principle to combine contiguous networks into one classless SET of IP addresses that can be summarized in routing tables. As a result fewer routes need to be advertised. Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 Click **Add Aggregate Address**.

Step 5 Enter a value for the aggregate timer (in seconds) in the **Aggregate Timer** field. Valid values are 0 or any value between 6 and 60. The default value is 30.

Step 6 Click (+)**Add** and update the **Add Aggregate Address** dialog box:

a) **Network**: Enter an IPv4 address or select the desired network/hosts objects.
b) **Attribute Map**: (Optional) Enter or select the route map used to SET the attribute of the aggregate route.
c) **Advertise Map**: (Optional) Enter or select the route map used to select the routes to create AS_SET origin communities.
d) **Suppress Map**: (Optional) Enter or select the route map used to select the routes to be suppressed.
e) **Generate AS SET path Information**: (Optional) Check the check box to enable generation of autonomous system SET path information.
f) **Filter all routes from updates**: (Optional) Check the check box to filter all more-specific routes from updates.
g) Click **OK**.

The BGP aggregate address settings are configured and fewer routes need to be advertised due to route aggregation.

What to do next

- For BGPv4 settings, proceed to Configure BGPv4 filtering settings, on page 16.
- For BGPv6 settings, proceed to Configure BGP network settings, on page 17.

Configure BGPv4 filtering settings

Configure filtering settings to filter routes or networks received in incoming BGP updates and restrict routing information that the router learns or advertises.

Filtering settings are used to filter routes or networks received in incoming BGP updates. Filtering is used to restrict routing information that the router learns or advertises.

Before you begin

Filtering is only applicable for a BGP IPv4 routing policy.

Follow these steps to configure BGPv4 filtering settings:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.
Step 2 Select the **Routing** tab.
Step 3 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.
Step 4 Choose **BGP** > **IPv4**.
Step 5 Click **Filtering**.

Note The **Filtering** field is applicable only to IPv4 settings.

Step 6 Click ( )**Add** and update the **Add Filter** dialog box:

a) **Access List**: Choose an access control list that defines which networks are to be received and which are to be suppressed in routing updates.

b) **Direction**: (Optional) Choose a direction that specifies if the filter should be applied to inbound updates or outbound updates.

c) **Protocol**: (Optional) Choose the routing process for which you want to filter: None, BGP, Connected, OSPF, RIP, or Static.

d) **Process ID**: (Optional) Enter the process ID for the OSPF routing protocol.

e) Click **OK**.

Step 7 Click **Save**.

BGPv4 filtering settings are configured to filter routes or networks based on the specified access control list and filtering criteria.

Configure BGP network settings

Configure BGP network settings to control which networks are advertised by the BGP routing process and to apply route map filters for network advertisement.

Network settings are used to add networks that will be advertised by the BGP routing process and route maps that will be examined to filter the networks to be advertised.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select the **Routing** tab.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 4 Choose **BGP** > **IPv4** or **IPv6**.

Step 5 Click **Networks**.

Step 6 Click **Add** and update the **Add Networks** dialog box:

a) **Network**: Choose the network to be advertised by the BGP routing processes.

Note For a network prefix to be advertised, a route to the device must exist on the routing table.

To add a new network object, refer to Create a network object.

b) (Optional) **Route Map**: Enter or choose a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. To add a new route map object, refer to Configure a route map.

c) Click **OK**.

Step 7 Click **Save**.

The BGP network settings are configured with the specified networks and route maps for advertisement filtering.

Configure BGP redistribution settings

Configure redistribution settings to define the conditions for redistributing routes from another routing domain into BGP.

Redistribution settings allow you to define the conditions for redistributing routes from another routing domain into BGP.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select the **Routing** tab.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 4 Choose **BGP** > **IPv4** or **IPv6**.

Step 5 Click **Redistribution**.

Step 6 Click **Add** and update the **Add Redistribution** dialog:

a) **Source Protocol**: Select the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

Note User-defined virtual routers do not support redistributing routes from RIP.

b) **Process ID**: Enter the identifier for the selected source protocol. Applies to the OSPF protocol. For devices using virtual routing, this drop-down lists the process ID assigned for the virtual router for which you are configuring the BGP settings.

c) **Metric**: (Optional) Enter a metric for the redistributed route.

d) **Route Map**: Enter or select a route map that should be examined to filter the networks to be redistributed.

If not specified, all networks are redistributed. To create a new route map object, click **Add** (+). For the procedure to add a new route map, see Configure Route Map Entry.

e) **Match**: The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more conditions. These options are enabled only when OSPF is chosen as the Source Protocol.

- Internal
- External 1
- External 2
- NSSA External 1

- NSSA External 2

f) Click **OK**.

The BGP redistribution settings are configured and routes from the specified source protocol are redistributed into the BGP domain according to the defined conditions.

Configure BGP route injection settings

Configure BGP route injection settings to control which routes are conditionally injected into the BGP routing table based on defined criteria.

Route injection settings allow you to define the routes to be conditionally injected into the BGP routing table.

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device.

Step 2 Select the **Routing** tab.

Step 3 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 4 Choose **BGP** > **IPv4** or **IPv6**.

Step 5 Click **Route Injection**.

Step 6 Click **Add** and update the **Add Route Injection** dialog box:

a) **Inject Map**: Enter or select the route map that specifies the prefixes to inject into the local BGP routing table. To create a new route map object, click **Add** (+). For the procedure to add a new route map, see Configure Route Map Entry.

b) **Exist Map**: Enter or select the route map containing the prefixes that the BGP speaker will track.

c) **Injected routes will inherit the attributes of the aggregate route**: Check this box to configure the injected route to inherit attributes of the aggregate route.

d) Click **OK**.

Step 7 Click **Save**.

The BGP route injection settings are configured and the specified routes will be conditionally injected into the BGP routing table according to the defined route maps.

Configure BGP route import and export

This task enables you to implement inter-virtual-router route leaking in BGP by importing or exporting routes using route target extended communities, allowing controlled sharing of routes between virtual routers.

In BGP, you can implement an inter-virtual-router route leak by importing or exporting routes using the route target extended community of the destination and source virtual routers respectively. You can use a route map to filter the desired route targets instead of leaking the entire routing table. You can also leak the routes of global virtual router to user-defined virtual routers and vice versa.

- You can configure BGP to leak routes between two user-defined virtual routers using the route target extended communities:
  - Tag the routes with the route targets from the source virtual router using route target export.
  - Import the routes that are matching the route targets into the destination virtual router using route target import.
  - Optionally, you can filter routes from source virtual router or to destination virtual router using export or import route maps respectively. You can configure route map with match extended community list for filtering the routes. Similarly, you can configure route map with set extended community route targets to tag the routes with the route target extended community.
- To import routes from the global virtual router to a user-defined virtual router, specify the IPv4/IPv6 route map in Global Virtual Router Import Route Map to import to the user-defined virtual router.
- To export routes from a user-defined virtual router to the global virtual router, in addition to exporting the route targets, you can also specify the Global Virtual Router Export Route Map to export from the user-defined virtual router.

The BGP inter-virtual-router route leaking supports both IPv4 and IPv6 prefixes.

Before you begin

- Create virtual routers.
- Enable BGP on the virtual router.
- Configure BGP, on page 6.

Follow these steps to configure BGP route import/export settings:

Procedure

Step 1 Choose **Devices** > **Device Management**, and edit the Firewall Threat Defense device. Select the **Routing** tab.

Step 2 (For a virtual-router-aware device) From the virtual routers drop-down, choose the virtual router for which you are configuring BGP.

Step 3 Choose **BGP** > **IPv4** or **IPv6**.

Step 4 (Supported for virtual routers only) Click **Route Import/Export**. In the **Route Targets Import** field, enter the route target extended community that you want to match for the routes to be imported. On deployment, the routes of the destination virtual router that match this value are imported to the source virtual router's BGP table.

Note

- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma-separated values.
- This value can range from 0:1 to 65534:65535.

Step 5 In the **Route Targets Export** field, enter the route target extended community to tag the source virtual router's routes with the route target value. On deployment, the routes of the source virtual router are tagged with this value.

Note
- The route target must be in **ASN:nn** format.
- You can enter multiple route targets as comma separated values.
- This value can range from 0:1 to 65534:65535.

Step 6 Route maps help you to narrow down the routes to be shared instead of leaking the entire routing table. Route map filtering is applied on the list of routes that are obtained with the specified route target values:

a) (Optional) Under **User Virtual Router**, choose the route map from the **Import Route Map** drop-down list to filter the routes at the destination virtual router.

Note The user virtual router import route map is effective only when the route targets import is configured.

b) (Optional) Under **User Virtual Router**, choose the route map from the **Export Route Map** drop-down list to filter the routes at the source virtual router before the routes are exported to other virtual routers.

Note You can use the match and set clauses in the route map with the route target extended community lists for filtering based on other criteria or tagging the routes with the route target community values. For more information, refer to Configure a route map.

Step 7 To share the routes between a user-defined virtual router and global virtual router, specify the route map under the **Global Virtual Router**:

a) To leak the global virtual router routes to the user-defined virtual router, select the route map from the **Import Route Map** drop-down list. The IPv4 or IPv6 route map is imported to the user-defined virtual router.

b) To leak the user-defined virtual router routes to the global virtual router, select the route map from the **Export Route Map** drop-down list. The IPv4 or IPv6 route map is exported to the global virtual router.

Note You must specify the route targets for export apart from specifying the route map.

Note You can use the match clause of the route map object to filter the routes for leaking. For more information, refer to Configure a route map.

Step 8 Follow the procedure (Step 3 to Step 8) to configure relevant BGP route import and export settings for other virtual routers as well. Click **Save**, then **Deploy** to apply your import/export settings.

When the packets flow into the ingress virtual router, BGP imports the routes from the destination virtual routers that have the matching route target value and if a route map is also configured, the routes are further filtered and used to identify the best path routes for routing the packets.