Configure Virtual Routers

For Releases 16.1R2 and later, except as noted.

A virtual router is a software object that functions like a hardware-based Layer 3 Internet Protocol (IP) router. A virtual router enables a computer to perform the functions of a physical router. Just as with a physical router, on a virtual router you configure static and dynamic routing protocols, including unicast and multicast protocols, router advertisements, and redistribution and import policies.

You configure all properties of a virtual router in the Configure Virtual Router popup window.

Set Up a Virtual Router

1. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
2. Select the Configuration tab in the top menu bar.
3. Select Networking > Virtual Routers in the left menu bar
4. Click the Add icon. The Configure Virtual Router popup window displays.

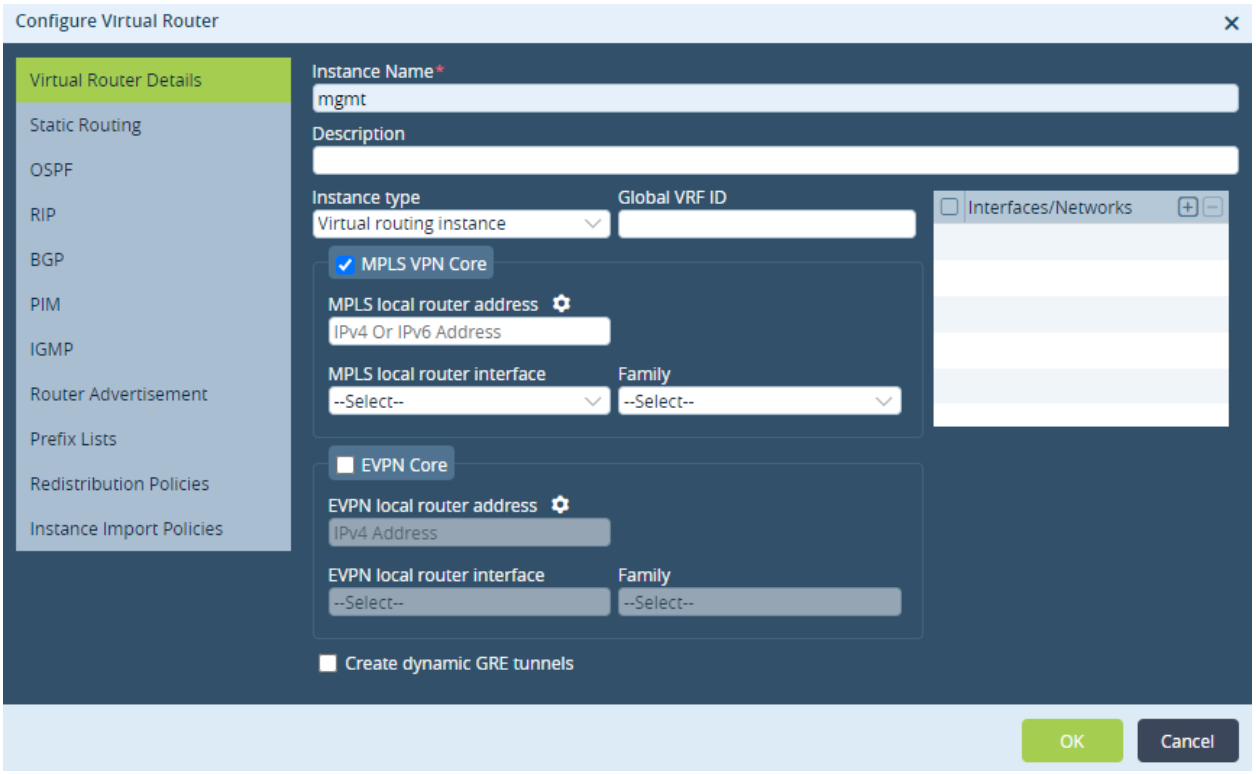

5. Select the Virtual Router Details tab. Enter information for the following fields.

Field	Description
Instance Name	Enter a unique name for the virtual router.
Description	Enter a text description for the interface.
Instance Type	Select the virtual router instance type: ◦ Virtual Routing Instance—Configure a simple VPN. This is the basic instance type. ◦ Virtual Routing Forwarding Instance—Configure a router for Layer 3 VPN.
Global VRF ID	Enter an ID for the global VRF.
MPLS VPN Core (Group of Fields)	For a virtual routing instance and for MPLS, click to configure the virtual router as the core router.
◦ MPLS Local Router Address	For a virtual routing instance and for MPLS, enter the local router's IPv4 or IPv6 address.
◦ MPLS Local Router Interface	Select the local router interface to use for MPLS.
◦ Family	Select the family to use for the virtual router.
EVPN Core (Group of Fields)	Ethernet VPN (EVPN) enables you to connect two or more Layer 2 domains over IP or MPLS Layer 3 underlay networks.
◦ EVPN Local Router Address	Enter the IP address of the local EVPN router.
◦ EVPN Local Router Interface	Select a local router interface for the EVPN core.
◦ Family	Select a family for the virtual router.
Create Dynamic GRE Tunnels	Click to create dynamic GRE tunnels.
Interfaces/Networks	Select one or more interfaces to assign to the routing instance.

6. Click OK.

Configure Static Routes

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select Static Routing in the left menu bar in the Configure Virtual Router window.

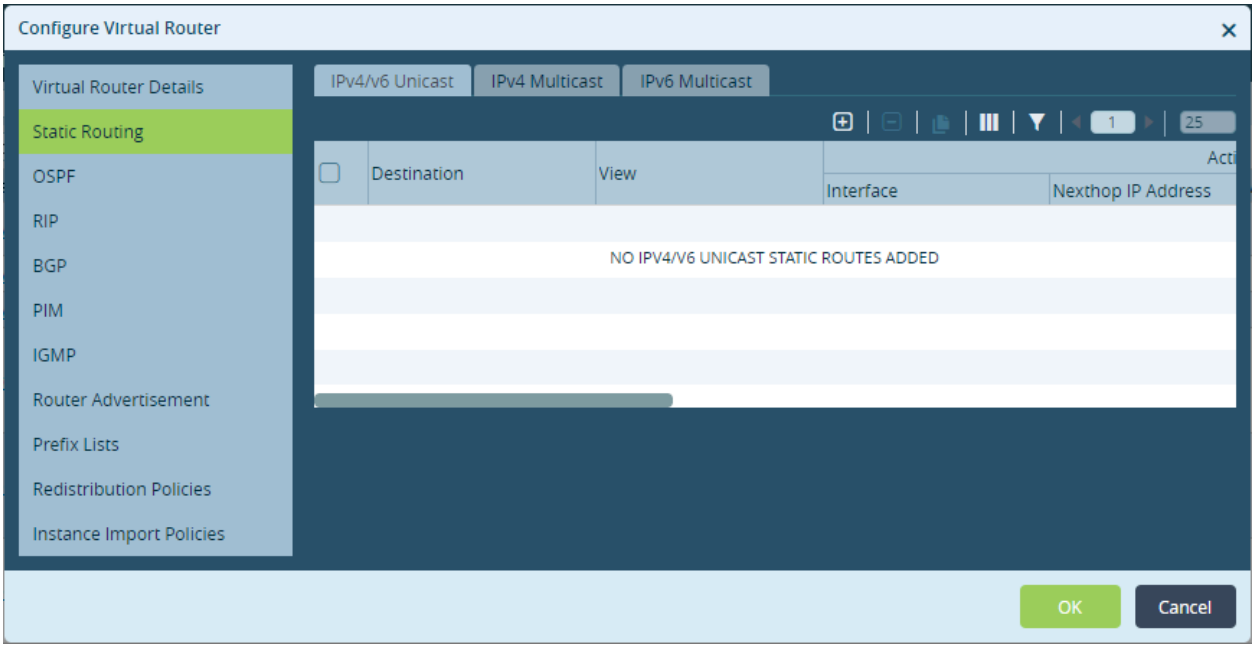

7. To add an IPv4 or IPv6 unicast static route, select the IPv4/IPv6 Unicast tab and then click the Add icon. Enter information for the following fields.

Updated: Tue, 14 Sep 2021 15:28:36 GMT

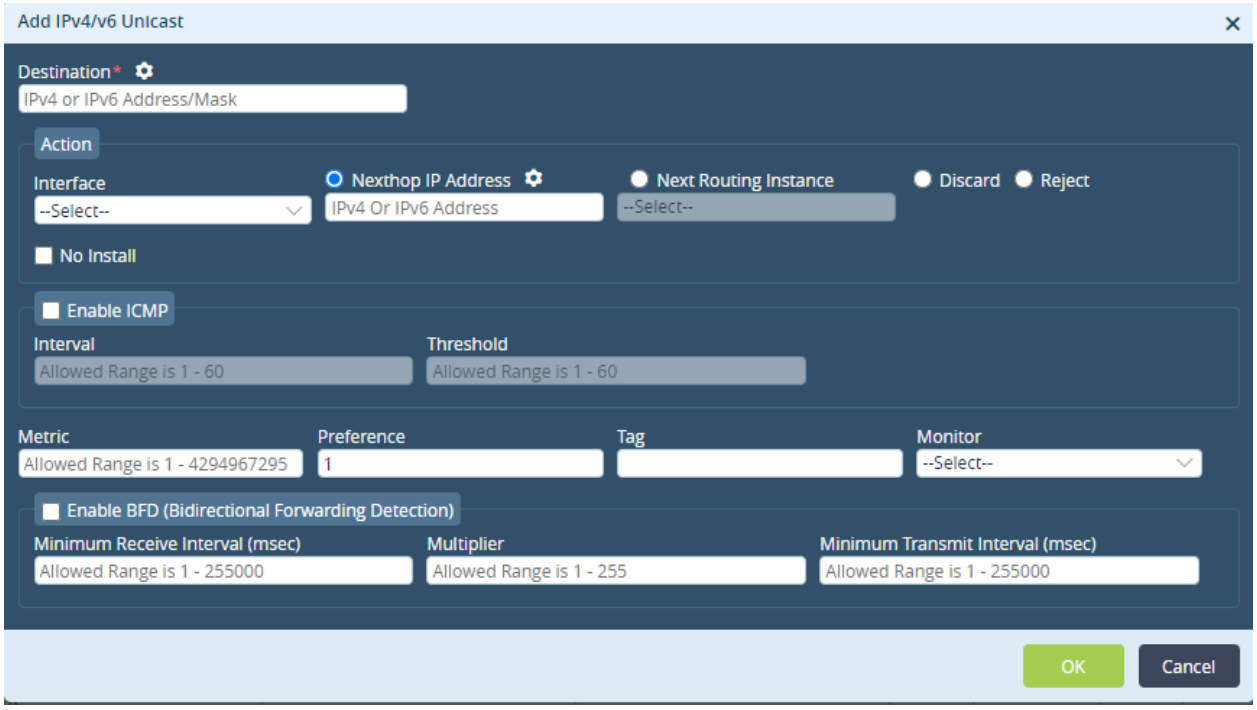

Field	Description
Destination	Enter the destination IP address or network.
Action (Group of Fields)	
◦ Next-Hop Interface	Select the next-hop interface towards the destination network.
◦ Next-Hop IP Address	Click to specify the IP address to use to reach the destination network.
◦ Next Routing Instance	Click to select the routing instance to use to reach the destination network.
◦ Discard	Click to discard the static route in the routing table.
◦ Reject	Click to configure the static route and take no action.
◦ No Install	Click to not install the static route in routing table.
Enable ICMP (Group of Fields)	Click to enable ICMP Enter these details: ◦ Interval—Indicates the time interval between ICMP packets. ◦ Threshold—Indicates the threshold for ICMP.
◦ Interval	Enter the time interval between ICMP packets.
◦ Threshold	Enter the threshold of ICMP packets.
Metric	Enter the cost to reach the destination network. *Range:* 1 through 4294967295
Preference	Enter the preference value of the static route. You can assign a preference for each route.
Tag	Enter a tag for the static route.

Monitor	Select the name of a liveness detection monitor that must be up for the static route to become active. To configure a monitor, see Configure IP SLA Monitor Objects.
Enable BFD (Group of Fields)	Click to enable BFD on the interface, to allow BFD to report when a static route becomes unavailable.
◦ Minimum Receive Interval	Enter the minimum time interval to receive routes, in milliseconds. *Range:* 1 through 255000 milliseconds
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval. *Range:* 1 through 255
◦ Minimum Transmit Interval	Enter the time after which routes can be retransmitted, in milliseconds. *Range:* 1 through 255000 milliseconds

8. Click OK

9. To add an IPv4 multicast static route, select the IPv4 Multicast tab and then click the Add icon. Enter information for the following fields.

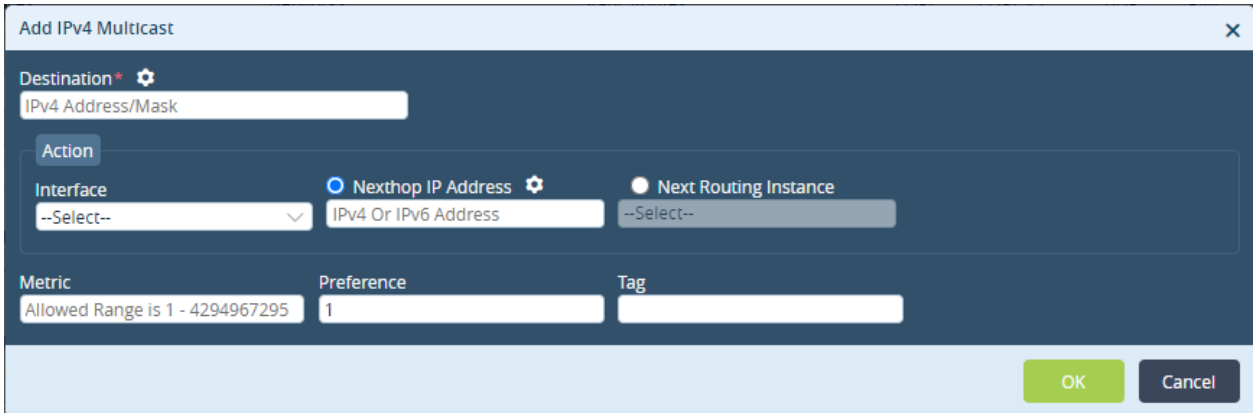

Field	Description
Destination	Enter the destination IP address or network.
Action (Group of Fields)	
◦ Interface	Select the interface towards the destination network.
◦ Next-Hop IP Address	Click to specify the IP address to use to reach the destination network.
◦ Next Routing Instance	Click to select the routing instance to use to reach the destination network.
Metric	Enter the cost to reach the destination network. *Range:* 1 through 4294967295
Preference	Enter the preference value of the IPv4 route.
Tag	Enter a tag for the IPv4route.

10. Click OK.

11. To add an IPv6 multicast static route, select the IPv6 Multicast tab and then click the Add icon. Enter information for the following fields.

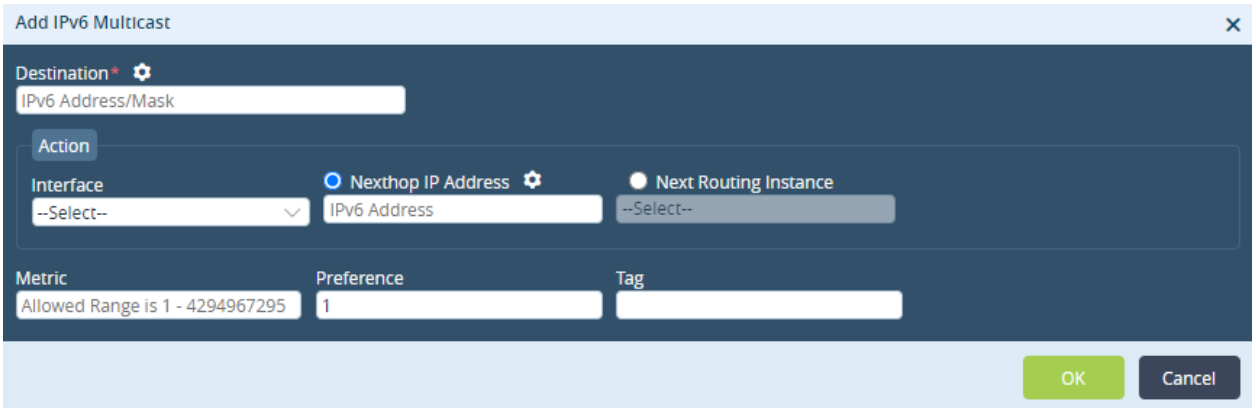

Field	Description
Destination	Enter the destination IP address or network.
Action (Group of Fields)	
◦ Interface	Select the interface towards the destination network.
◦ Next-Hop IP Address	Click to specify the IP address to use to reach the destination network.
◦ Next Routing Instance	Click to select the routing instance to use to reach the destination network.
Metric	Enter the cost to reach the destination network. *Range:* 1 through 4294967295
Preference	Enter the preference value of the IPv6 route.
Tag	Enter a tag for the IPv6 route.

12. Click OK. The static route displays in the Configure Virtual Router popup window.

Configure OSPF and OSPFv3

The open shortest path first (OSPF) is an interior gateway routing protocol (IGP) that uses a link-state routing algorithm OSPFv2 for IPv4 is defined in RFC 2328.

To configure OSPF and OSPFv3:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.

3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select OSPF in the left menu bar in the Configure Virtual Router window.

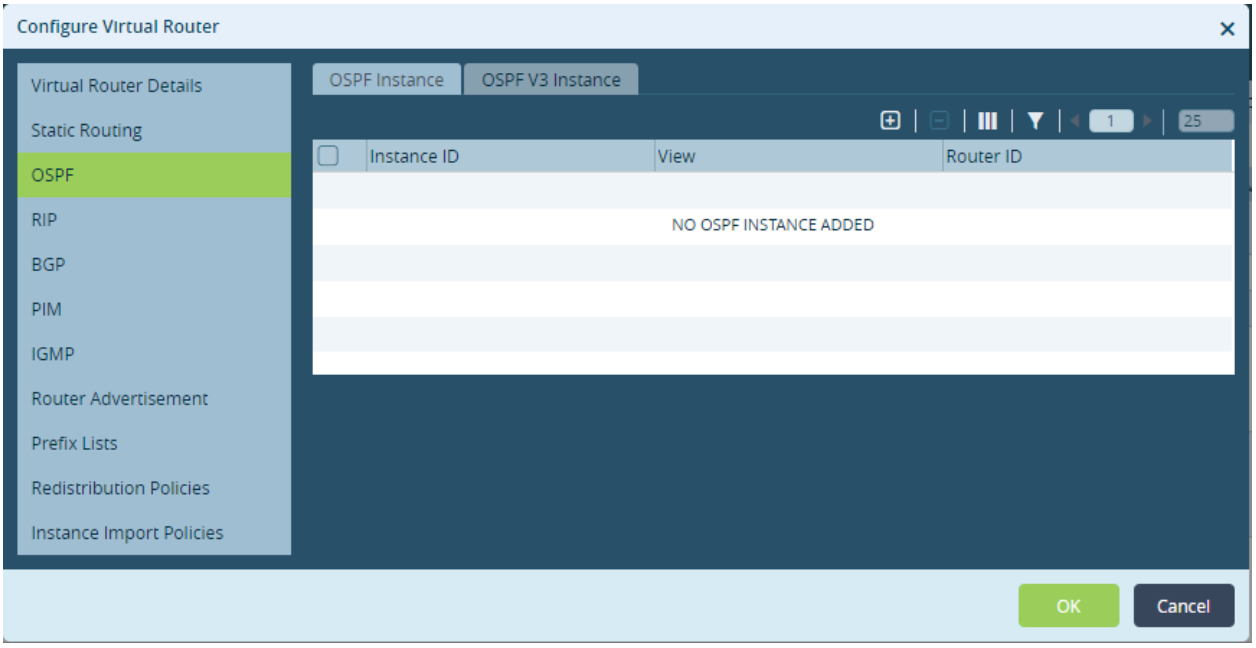

7. Click the Add icon. The Add OSPF Instance popup window displays. Enter information for the following fields.

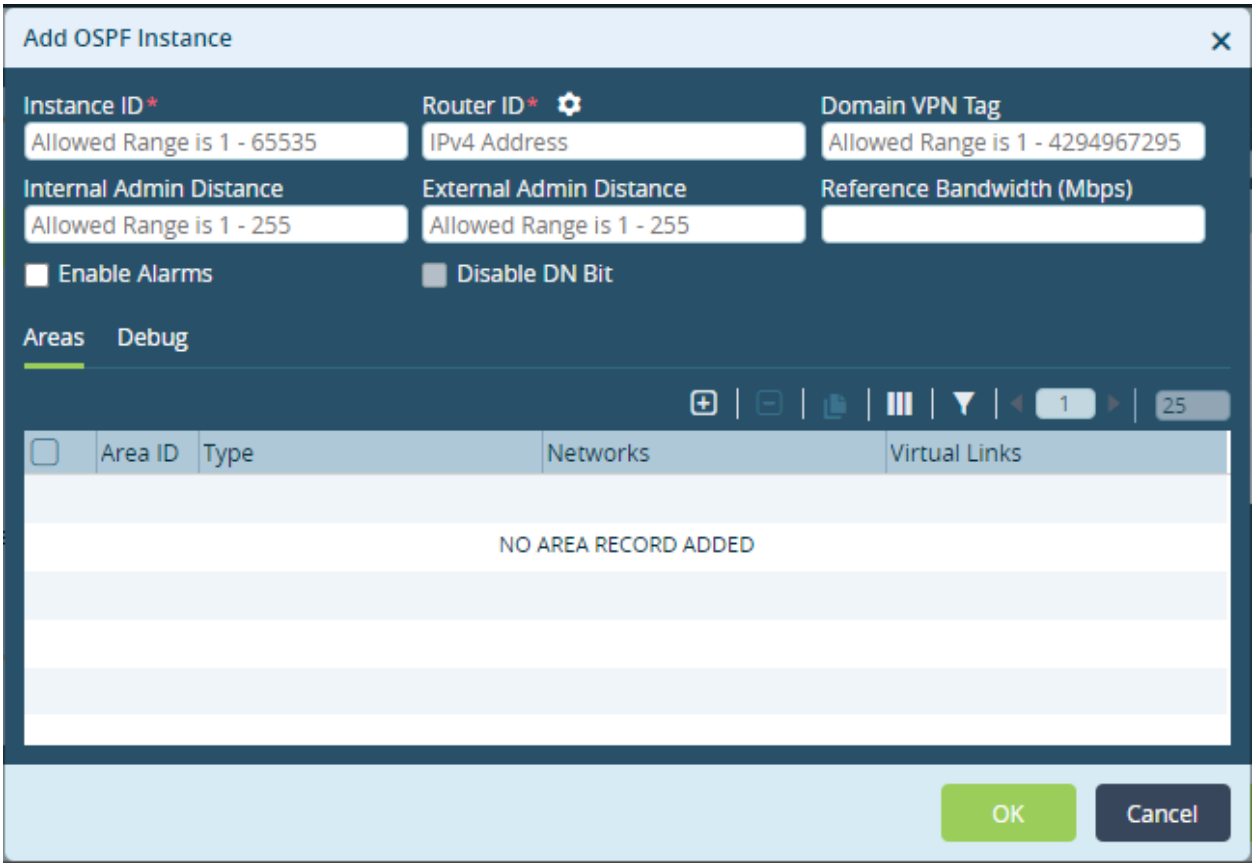

Field	Description
Instance ID	Enter the instance ID to assign to OSPF. *Range:* 1 through 65535
Router ID	Enter the router IP address to use for OSPF.
Domain VPN Tag	Enter the MPLS VPN tag attached to OSPF routes in this domain. Use this to enabled the OSPF PE-CE protocol on a PE router for external learned routes. *Range:* 1 through 4294967295
Internal Admin Distance	Enter the administrative distance for internal routes (routes learned within the routing domain). *Range:* 1 through 255
External Admin	Enter the administrative distance for external routes (routes learned from another routing

Field	Description
Distance	domain). *Range:* 1 through 255
Reference Bandwidth	Enter the reference bandwidth value to use when calculating the interface cost, in Mbps.
Enable Alarms	Click to enable the generation of alarms.
Disable DN Bit	Click to reset the DN bit When redistributing routes. The DN bit is used for loop prevention, so it is always enabled or set.
Areas Tab	Select the Areas tab to configure the OSPF area.

8. Click the Add icon to configure an OSPF area. An area is a collection of OSPF networks, routers, and links. Each area is assigned an ID. An area with zero as its ID is a backbone or normal area. Areas with non-zero IDs are non-backbone areas. Each area must be connected to the backbone area known as area 0. Areas communicate with other areas through the backbone area. Enter information in the following fields.

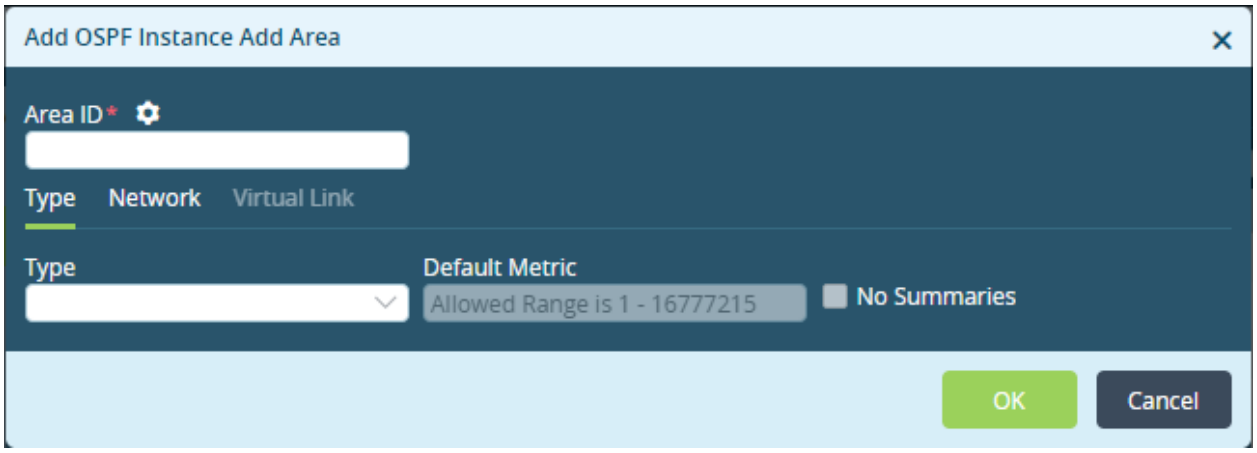

Field	Description
Area ID	Enter an ID for the area. A backbone area has an area ID of 0.0.0.0. Areas with non-zero IDs are non-backbone areas.
Type (Tab)	Select the Type tab to configure the OSPF area type.
◦ Type	Select the area type: ◦ Backbone—Backbone area is normal area.

Field	Description
	◦ Normal—For non-backbone area. ◦ NSSA—Not-so-stubby areas can import external routes into the OSPF routing domain and that can provide transit services to routing domains that are not part of the OSPF routing domain. ◦ Stub—External routes are not advertised.
Default Metric	For a stub or an NSSA area, enter the metric for the default route. *Range:* 1 through 16777215
◦ No Summaries	For all area types except Normal, click to have a border router not advertise routes from the area.

9. Select the Network tab to configure the network interface or IP address of the OSPF network. The list of configured networks displays.

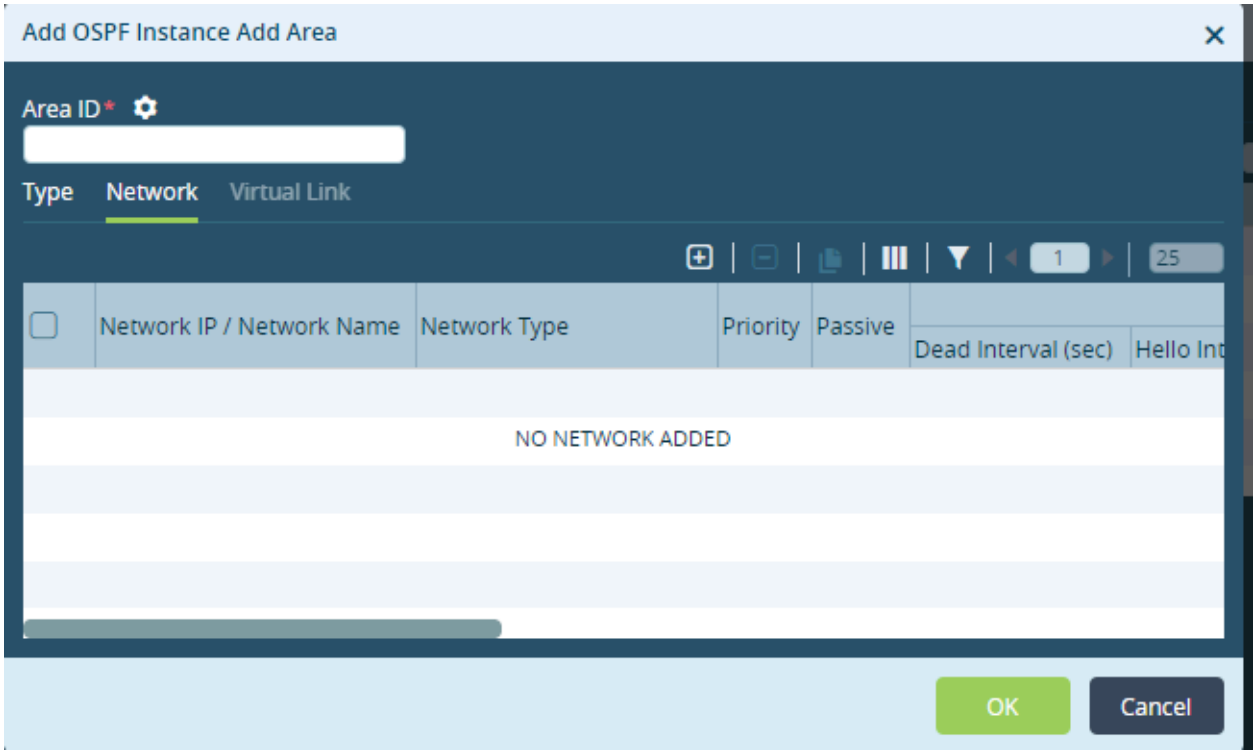

10. Click the Add icon, and enter information for the following fields.

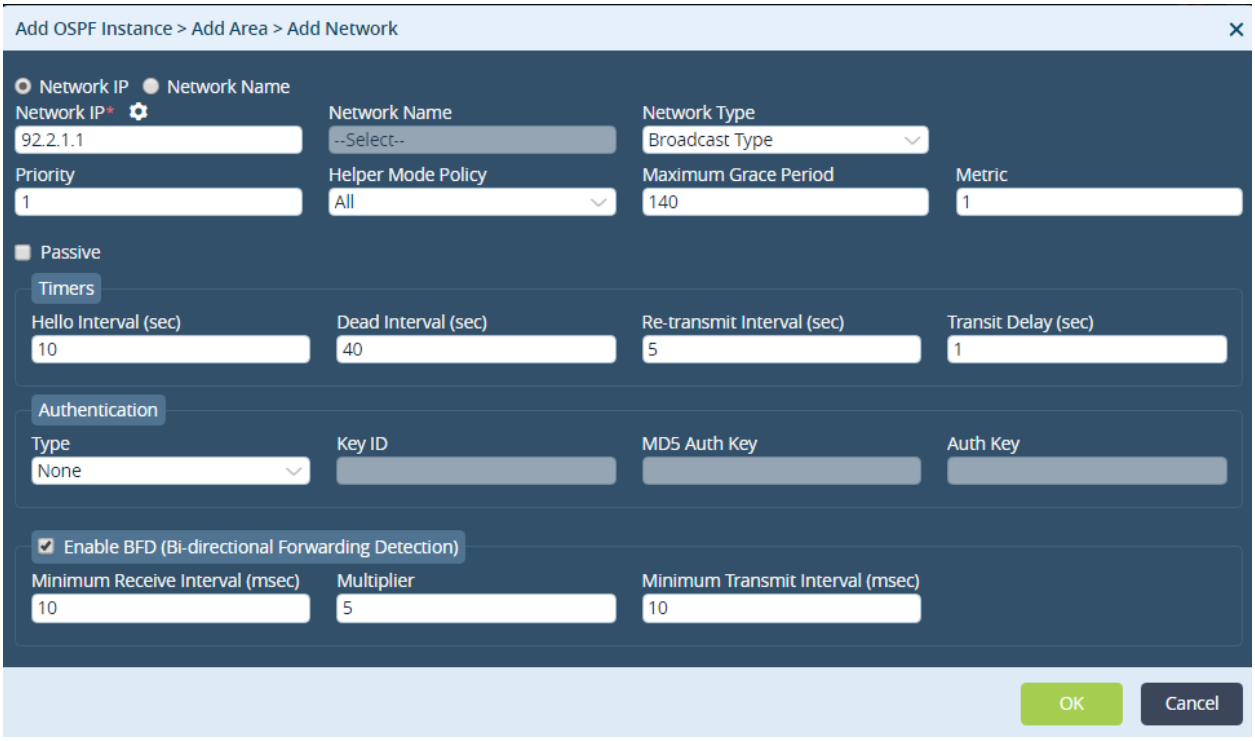

Field	Description
Network IP	Click and enter the IP address of the network. Click the Tool icon to parameterize the IP address.
Network Name	Click and enter the name of the network.
Network Type	Select the network type: ◦ Broadcast Type ◦ Loopback Type ◦ Point-to-Point Type
Priority	Enter a priority value to use in the election of the designated router (DR) and the backup designated router (BDR). On a multiaccess network, the OSPF router with the highest priority becomes the designated router, and the OSPF router with the second-highest priority becomes the backup router. If you set the priority to 0, the device does not participate in designated router and backup designated router election process
Helper Mode Policy	Select in which peer OSPF restart situation the local router should act as a helper: ◦ All—All OSPF restart situations ◦ Policy Reload—Software upgrade or reload of the peer router

Field	Description
	◦ Policy Software—Crash of the OSPF process on the peer router ◦ Policy Switch—Control plane switchover on the peer router ◦ Policy Unknown—OSPF issues a restart reason not signaled in the graceful restart (type 9) link-state advertisement (LS
Maximum Grace Period	Enter a value to signal how long, in seconds, a helper should help a router. When this period expires, the helper brings down the adjacency with the restarting router, flushes its LSAs from the database, and floods new LSAs to the rest of the network to inform the other routers that it has lost its adjacency to the neighbor.
Metric	Enter a value for the OSPF interface cost, which is used to calculate the total cost to reach a destination. *Range:* 1 through 65535 *Default:* 1
Passive	Click to indicate that the router is a passive listener. A passive router does not advertise itself. If you do not click, the router actively propagates messages.
Timers (Group of Fields)	
◦ Hello Interval	Enter the time, in seconds, between the transmission of hello packets that this interface sends to neighbor routers. *Range:* 1 through 255 seconds *Default:* 10 seconds
◦ Dead Interval	Enter the time period, in seconds, during which at least one hello packet must be received from a neighbor before the router declares that neighbor to be down. *Range:* 1 through 65535 seconds *Default:* 40 seconds
◦ Retransmit Interval	Enter the time, in seconds, between the retransmission of LSAs to adjacent routers for a given interface. *Range:* 1 through 3600 seconds *Default:* 5 seconds
◦ Transit Delay	Enter the delay in retransmitting a message, in seconds Enter the time, in seconds, for how often to transmit a link-state update (LSU) on the interface.

Field	Description
	Range: 1 through 3600 seconds *Default:* 1 second
Authentication (Group of Fields)	
◦ Type	Select how to authenticate OSPF router traffic: ◦ MD5—Use encrypted authentication ◦ Simple Password—Use simple password-based authentication.
◦ Key ID	For MD5, enter the key ID.
◦ MD5 Auth Key	For MD5, enter the authorization key.
◦ Auth Key	For password-based authentication, enter the password.
Enable BFD (Group of Fields)	Click to enable BFD for OSPF, When BFD is enabled, when OSPF goes down, the router is marked as being down. When a BFD session that supports OSPF goes down, the OSPF neighborship also goes down without waiting for the dead timer interval to expire.
◦ Minimum Receive Interval	Enter the time interval, in milliseconds, at which the BFD peer can receive control packets. *Range:* 1 through 255000 milliseconds *Default:* 150 milliseconds
◦ Multiplier	Enter the number of times that a BFD control packet can be missed before BFD declares the neighbor to be down.
◦ Minimum Transmit Interval	Enter the time interval, in milliseconds, at which this device can send BFD control packets. *Range:* 1 through 255000 milliseconds *Default:* 150 milliseconds

5. Select the Virtual Link tab to configure an OSPF virtual link. When you merge networks, non-backbone areas communicate with each other through a virtual link. The list of configure virtual links displays.

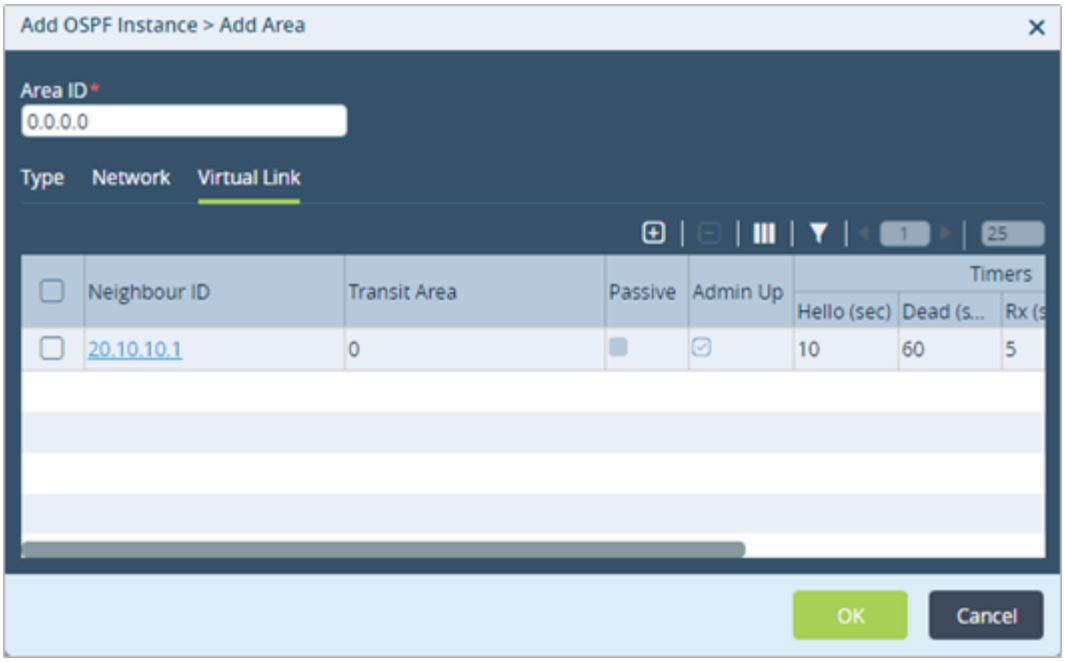

6. Click the ⊞ Add icon, and enter information for the following fields.

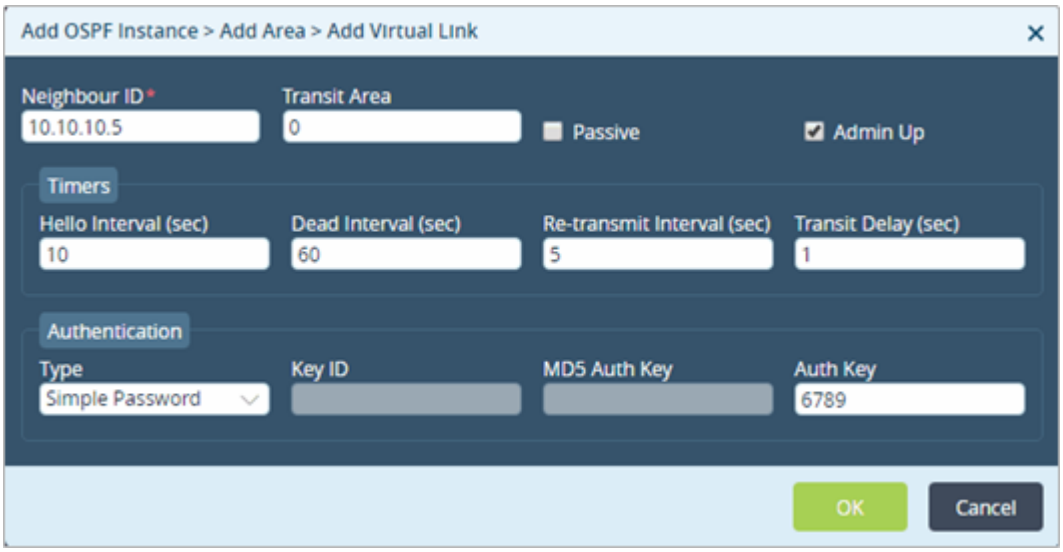

Field	Description
Neighbor ID	Enter the IP address of the neighboring area.
Transit Area	Enter the ID or IP address of the backbone area.

Field	Description
Passive	Click to mark the router as a passive listener. A passive router sends no advertisement messages.
Admin Up	Click to indicate that the administrative status of the link is up.
Timers (Group of Fields)	There are two timers for OSPF as below.
◦ Hello Interval	Enter the interval, in seconds after which router sends advertisement messages.
◦ Dead Interval	Enter the time to wait, in seconds, before the router declares a neighbor to be dead because it has received no advertisements within that amount of time.
◦ Retransmit Interval	Enter the retransmit interval, in seconds, after which the router can retransmit a message.
◦ Transmit Delay	Enter the delay, in seconds, for retransmitting a message.
Authentication (Group of Fields)	
◦ Type	Select how to authenticate router traffic: ◦ MD5—Use encrypted authentication ◦ Simple Password—Use password-based authentication
◦ Key ID	For MD5 authentication, enter the key ID.
◦ MD5 Auth key	For MD5 authentication, enter the authorization key.
◦ Auth Key	For password-based authentication, enter the password.

7. Click OK.

8. In the OSPFv3 Instance tab, click the Add icon to configure OSPFv3. Enter information for the following fields.

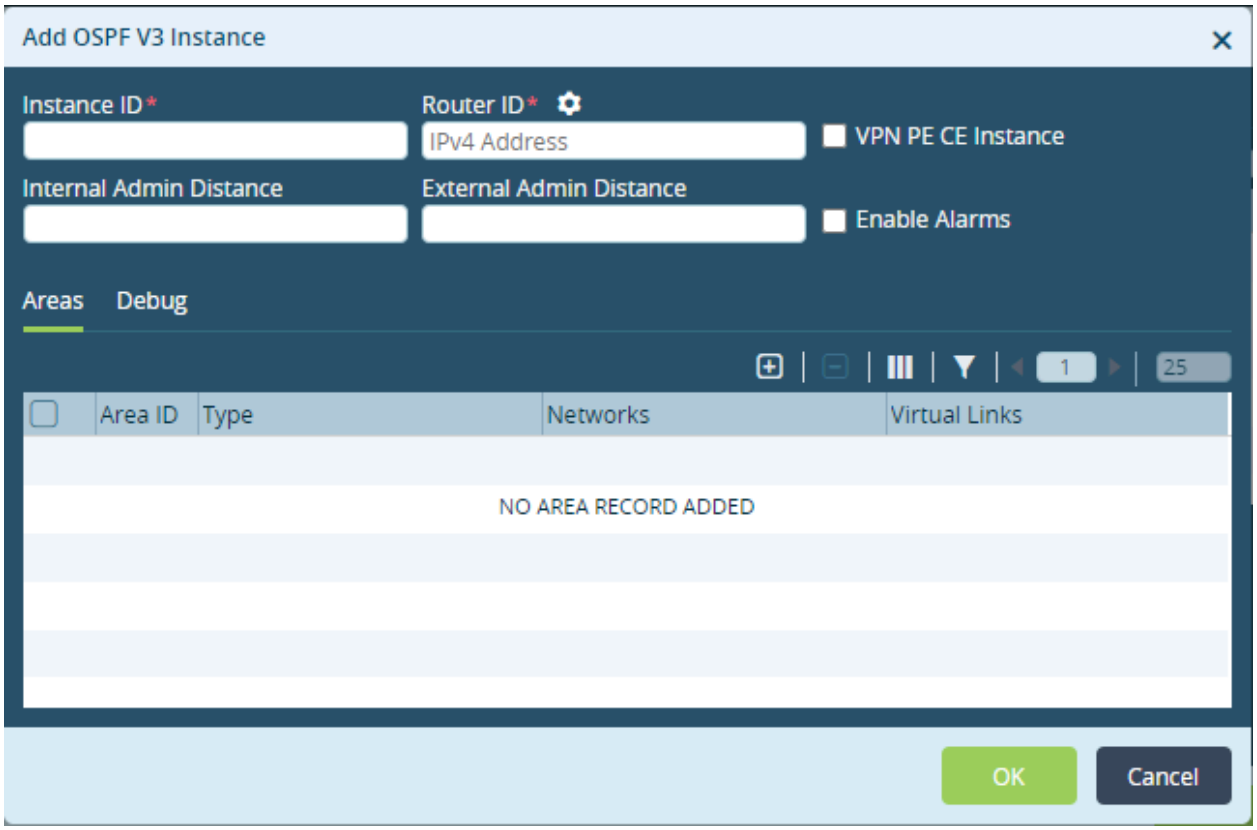

Field	Description
Interface ID	Enter the instance ID to assign to OSPFv3.
Router ID	Enter the router IP address to use for OSPFv3.
VPN PE CE Instance	Click to enable a VPN PE CE instance.
Internal Admin Distance	Enter the administrative distance for internal routes (routes learned within the routing domain).
External Admin Distance	Enter the administrative distance for external routes (routes learned from another routing domain).
Enable Alarms	Click to enable the alarm generation.

9. Click the Add icon to configure an OSPFv3 area. An area is a collection of OSPF networks, routers, and links. Each area is assigned an ID. An area with zero as its ID is a backbone or normal area. Areas with non-zero IDs are non-backbone areas. Each area must be connected to the backbone area known as area 0. Areas communicate with other areas through the backbone area. Enter information for the following fields.

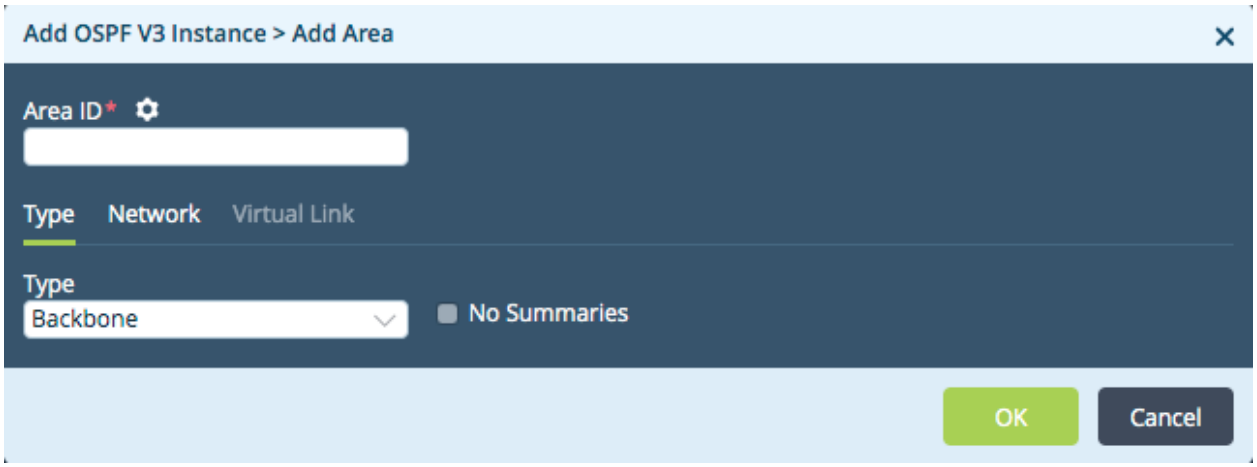

Field	Description
Area ID	Enter an ID for the area. A backbone area has an area ID of 0.0.0.0. Areas with non-zero IDs are non-backbone areas.
Type (Tab)	Select the Type tab to configure the OSPFv3 area type.
◦ Type	Select the area type: ◦ Backbone—Backbone area is normal area. ◦ NSSA—Not-so-stubby areas can import external routes into the OSPF routing domain and that can provide transit services to routing domains that are not part of the OSPF routing domain. ◦ STUB—External routes are not advertised.
◦ No Summaries	For all area types except Normal, click to have a border router not advertise routes from the area.

10. Select the Network tab to configure the network interface or IP address of the network running OSPFv3. The list of configured networks displays.

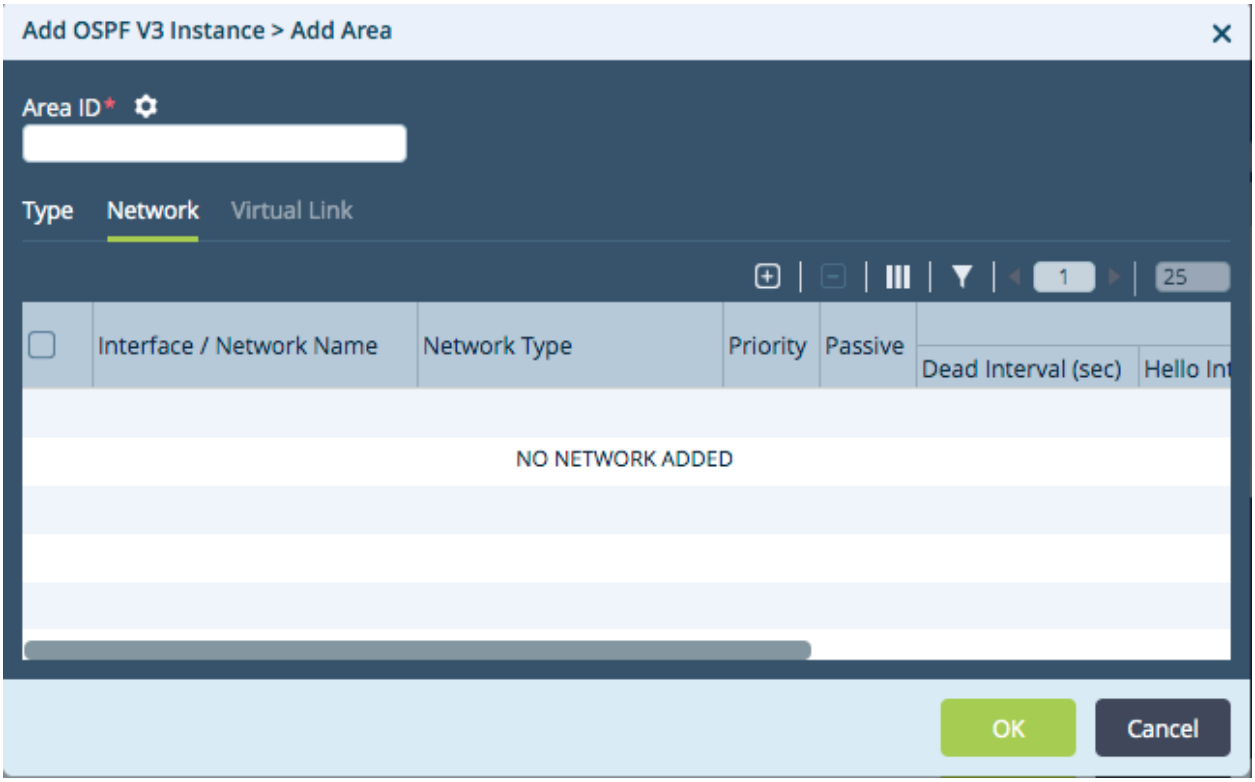

11. Click the Add icon, and enter information for the following fields.

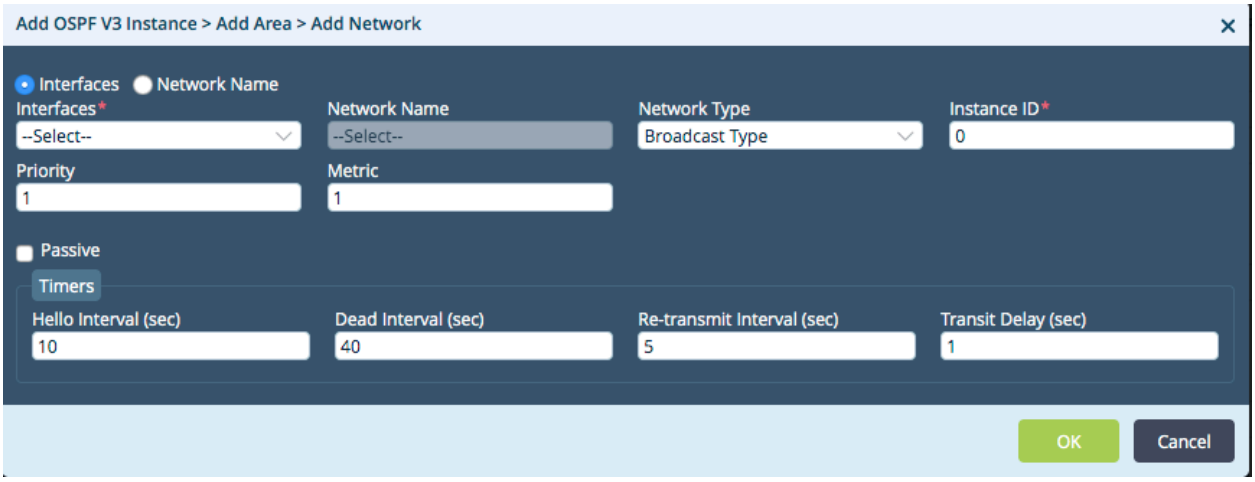

Field	Description
Interfaces	Select an interface name.
Network Name	Click to enter the name of the network.

Field	Description
Interfaces	Select an interface for the OSPFv3 area.
Network Name	Select the name of the network for the OSPFv3 area.
Network Type	Select the network type: ◦ Broadcast Type ◦ Loopback Type ◦ Point-to-Point Type
Instance ID	Enter the ID for the OSPFv3 instance.
Priority	Enter a value for the priority. A router with a higher priority propagates routes before other routers.
Metric	Enter a metric value to use to determine how to choose a route to advertise. The route can be chosen based on path length, bandwidth, hop count, load, path cost, MTU, and communication cost.
Passive	Click to indicate that the router is a passive listener. A passive router does not advertise itself. If you do not click, the router actively propagates messages.
Timers (Group of Fields)	
◦ Hello Interval	Enter the interval, in seconds, after which router advertises itself.
◦ Dead Interval	Enter the time to wait, in seconds, before declaring a router dead, in seconds, because the router does not advertise itself.
◦ Retransmit Interval	Enter the time after which the router can transmit a message, in seconds.
◦ Transit Delay	Enter the delay in retransmitting a message, in seconds.

12. Select the Virtual Link tab to configure an OSPFv3 virtual link. When you merge networks, non-backbone areas communicate with each other through a virtual link. The list of configured virtual links displays.

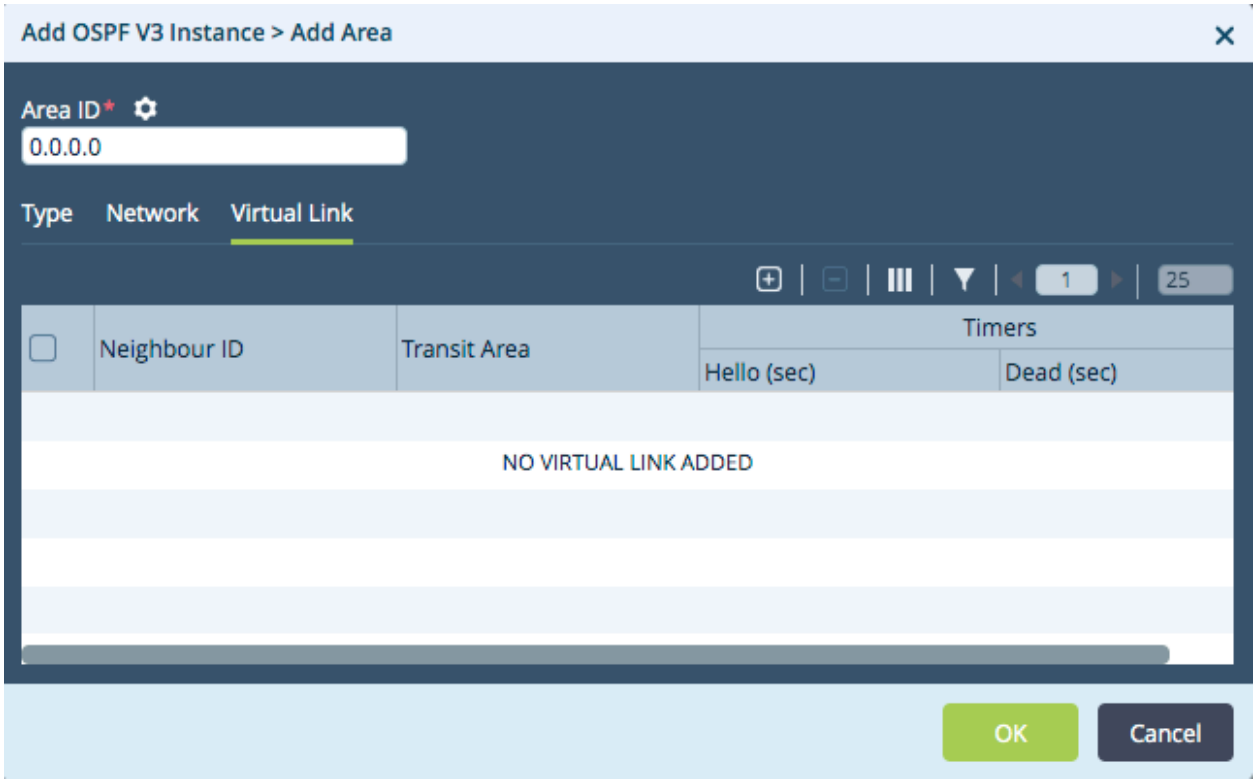

13. Click the Add icon, and add information for the following fields.

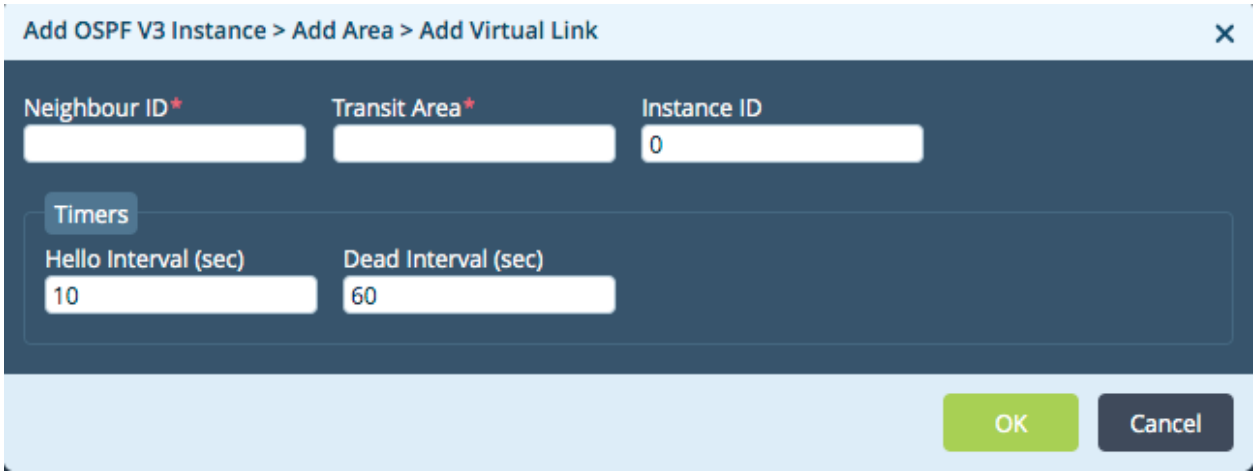

Field	Description
Neighbor ID	Enter the IP address of the neighboring area.

Field	Description
Transit Area	Enter the ID or IP address of the backbone area.
Instance ID	Enter the ID of the OSPFv3 instance.
Timers (Group of Fields)	There are two timers for OSPFv3 as below.
◦ Hello Interval	Enter the interval, in seconds after which router sends advertisement messages.
◦ Dead Interval	Enter the time to wait, in seconds, before the router declares a neighbor to be dead because it has received no advertisements within that amount of time.

14. Click OK. The Configure Virtual Router popup window displays the OSPF instance.

Configure RIP

Routing information protocol (RIP) is a distance-vector routing protocol. RIP employs hop counts as routing metrics and prevents routing loops by implementing a limit on the number of hops allowed in the source-to-destination path. The largest number of hops allowed for RIP is 15. This number, however, limits the size of networks supported by RIP.

To configure RIP:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select RIP in the left menu bar in the Configure Virtual Router window.

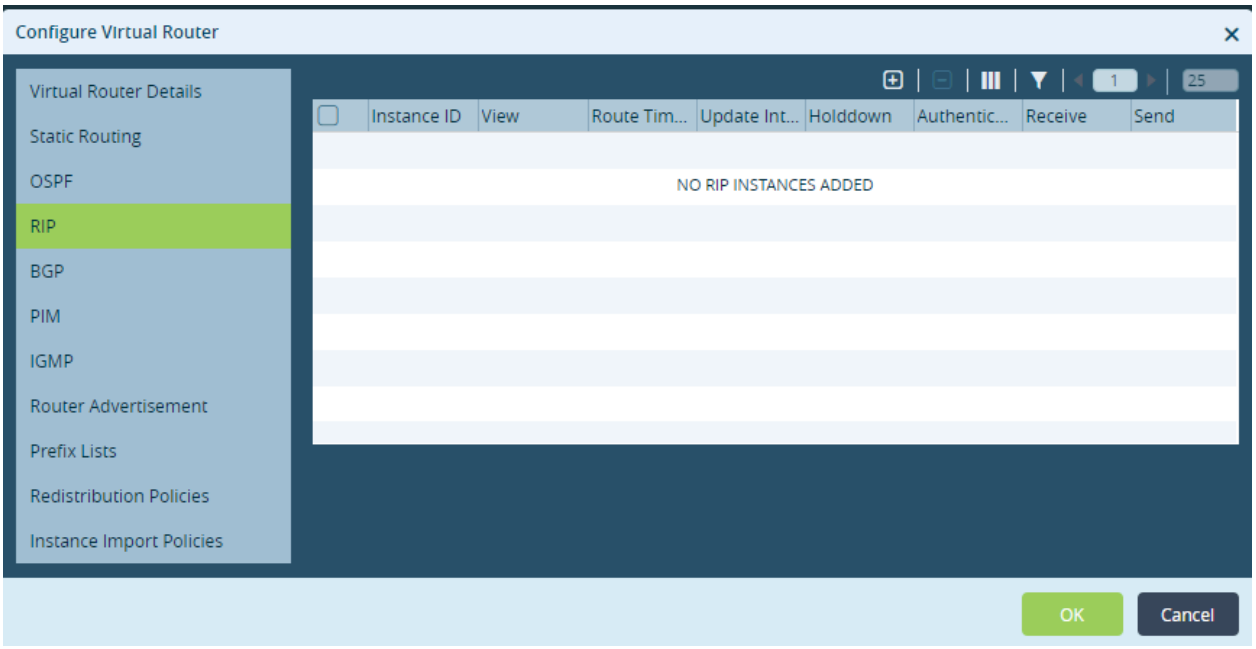

7. Click the Add icon. The Add RIP Instance popup window displays.

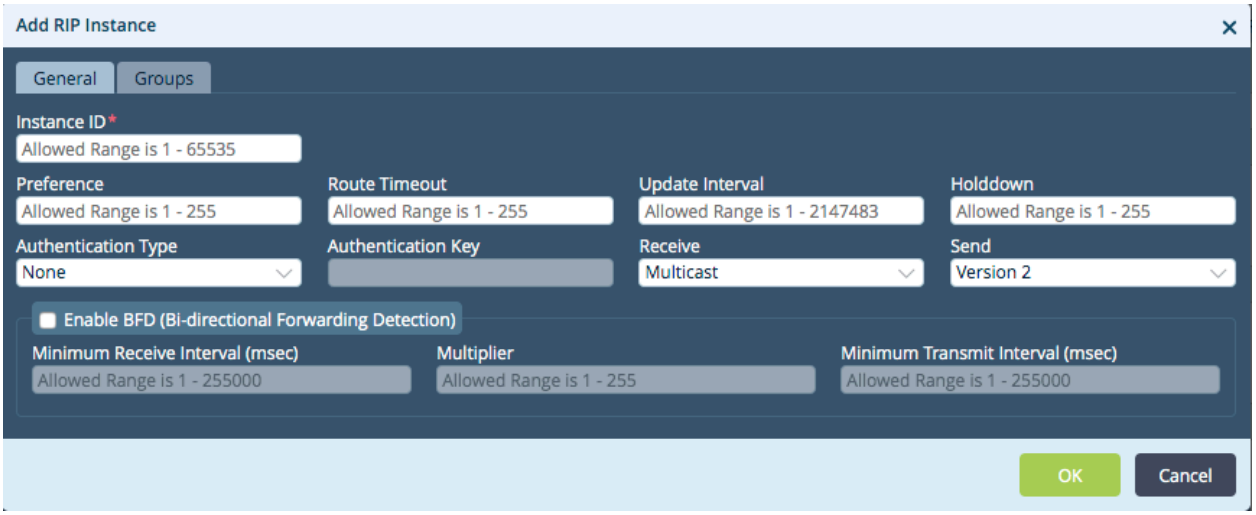

8. In the General tab, enter information for the following fields.

Field	Description
General Tab	Select the General tab to enter the details.
Instance ID	Enter the RIP instance ID number. *Range:* 1 through 65535

Field	Description
Preference	Enter a preference value for the instance. *Range:* 1 through 255
Route Timeout	Enter the time value for routes, in milliseconds. *Range:* 1 through 255
Update Interval	Enter the interval between gratuitous response messages in seconds. Response messages are broadcast to all RIP-enabled interfaces. *Range:* 1 through 2147483 *Default:* 30 seconds
Holddown	Enter a value for the hold-down timer. The hold-down timer is started for each route entry when the hop count changes from a lower value to higher value. During this time, no update can be made to the routing entry. *Range:* 1 through 255
Authentication Type	Select the type of authentication: ◦ MD5 ◦ None ◦ Simple password
Authentication Key	For simple password authentication, enter the password.
Receive	Select how to receive response message from neighboring routers: ◦ Multicast ◦ None
Send	Selection how to send request messages: ◦ None ◦ Version 2

Field	Description
Enable BFD (Group of Fields)	Click to enable BFD on the interface, to allow BFD to report when RIP becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds. *Range: 1 through 255000*
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval. *Range: 1 through 255*
◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds. *Range: 1 through 255000*

9. Select the Groups tab to configure information for a RIP group. The Add Group popup window displays. Click the the ⊞ Add icon to add a RIP group.
10. In the General tab, enter information for the following fields:

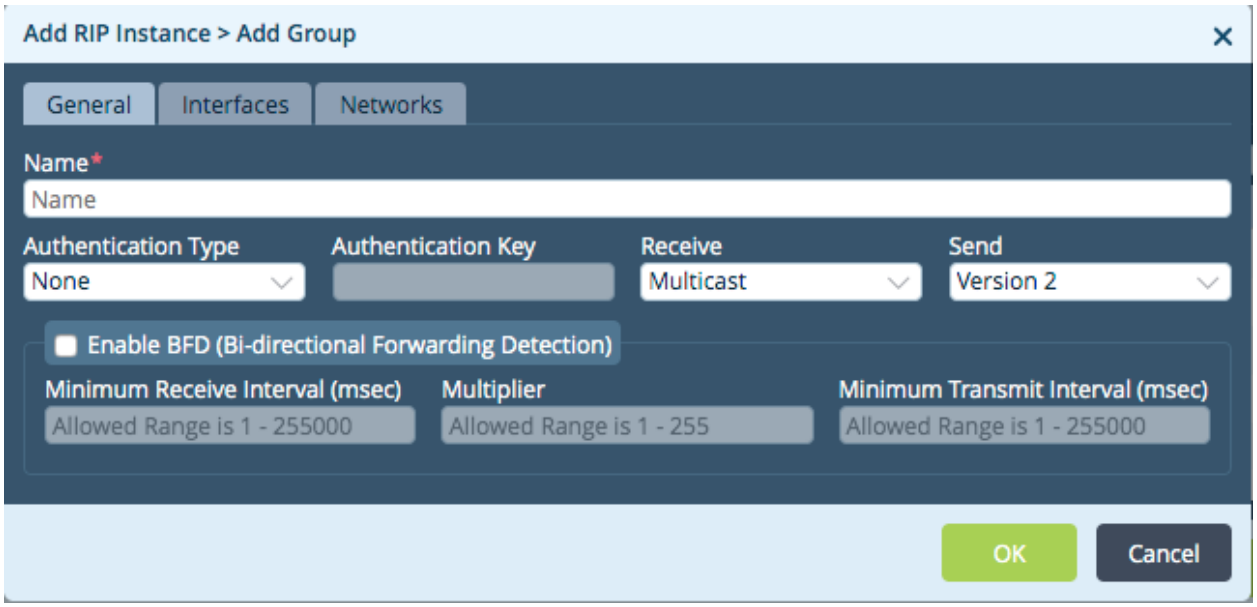

Field	Description
General Tab	Select the General tab to enter the details.
Name	Enter the name of the RIP group.
Authentication Type	Select the type of authentication: ◦ MD5 ◦ None ◦ Simple Password
Authentication Key	For simple password authentication, enter the password.
Receive	Select how to receive response message from neighboring routers: ◦ Multicast ◦ None
Send	Select how to send request messages: ◦ None ◦ Version 2
Enable BFD (Group of Fields)	Click to enable BFD on the interface, to allow BFD to report when RIP becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds. *Range:* 1 through 255000
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval. *Range:* 1 through 255
◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds. *Range:* 1 through 255000

11. Click OK.

12. Select the Interfaces tab to configure interfaces in the RIP group. Enter information for the following fields.

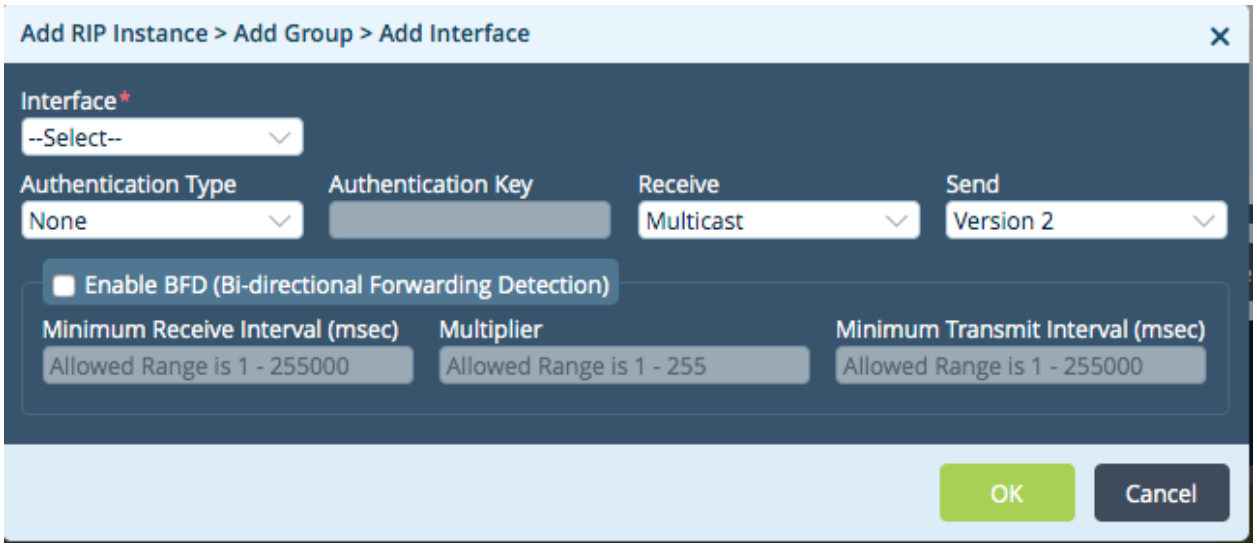

Field	Description
Interface	Select the interface to add to the RIP group.
Authentication Type	Type of authentication: ◦ MD5 ◦ None ◦ Simple Password
Authentication Key	For Simple Password authentication, enter the password. It can be up to 64 characters long.
Receive	Select how to receive response message from neighboring routers: ◦ Multicast ◦ None
Send	Selection how to send request messages: ◦ None ◦ Version 2
Enable BFD (Group of Fields)	Click to enable BFD on the interface, to allow BFD to report when RIP becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds. *Range:* 1 through 255000
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval. *Range:* 1 through 255
◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds. *Range:* 1 through 255000

13. Click OK. The Edit Group window displays the interface.

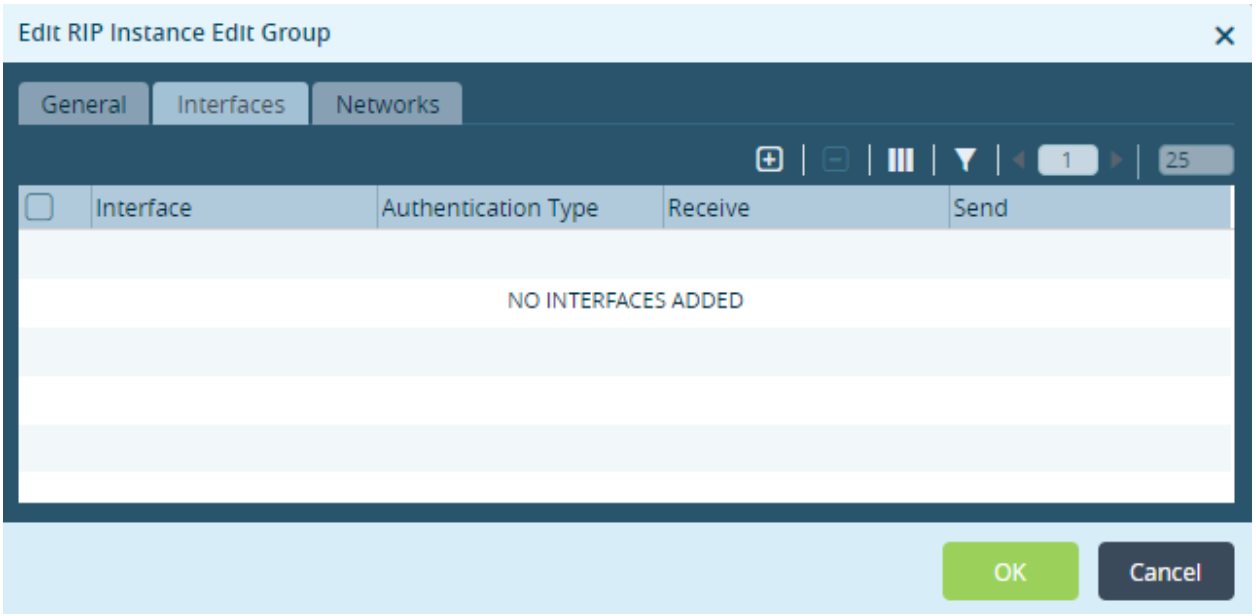

14. Select the Network tab to configure networks in the RIP group, and enter information for the following fields.

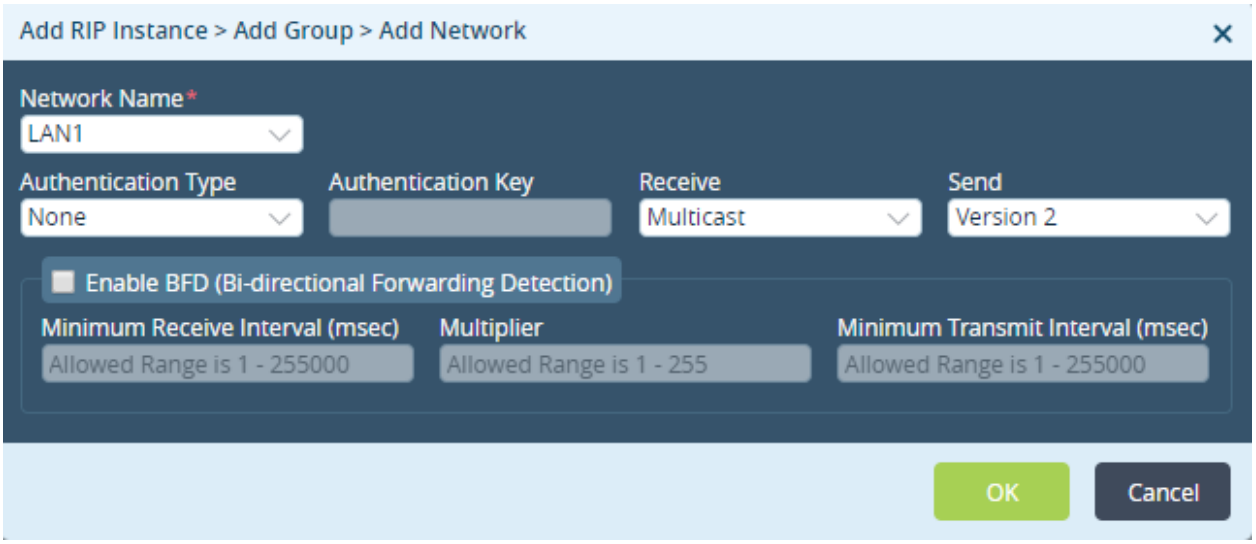

Field	Description
Network Name	Select the network to add to the RIP group.
Authentication Type	Type of authentication: ◦ MD5 ◦ None ◦ Simple Password

Field	Description
Authentication Key	For Simple Password authentication, enter the password.
Receive	Select how to receive response message from neighboring routers: ◦ Multicast ◦ None
Send	Selection how to send request messages: ◦ None ◦ Version 2
Enable BFD (Group of Fields)	Click to enable BFD on the interface, to allow BFD to report when RIP becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds. *Range:* 1 through 255000
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval. *Range:* 1 through 255
◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds. *Range:* 1 through 255000

15. Click OK. The list of configured networks displays.

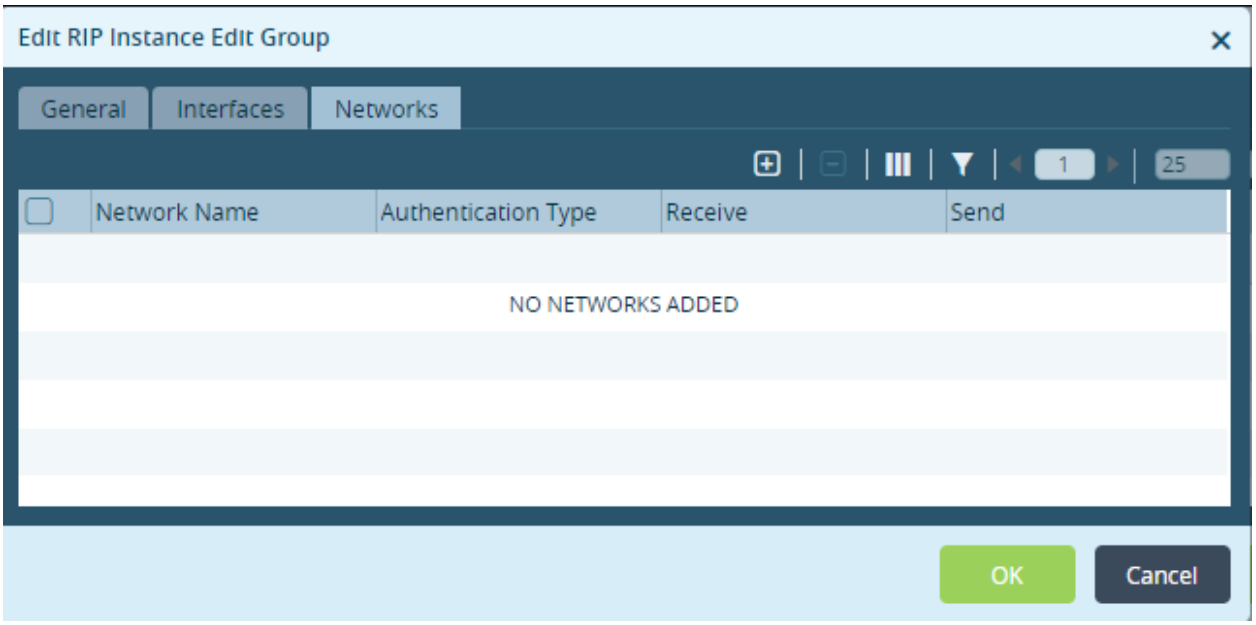

16. Click OK. The Configure Virtual Router popup window displays the RIP instance.

Configure BGP

Border gateway protocol (BGP) is an exterior gateway protocol (EGP) used for exchanging routing information between gateway hosts in a network. BGP is often the protocol used between gateway hosts on the Internet.

Configure BGP

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the BGP tab in the left menu bar.

Configure Basic BGP

1. If you are continuing from the previous section, skip to Step 7.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the BGP tab in the left menu bar. The main pane displays a list of the BGP instances that are already configured.
7. Click the Add icon. The Add BGP Instance popup window displays.
8. Select the General tab, and enter information for the following fields.

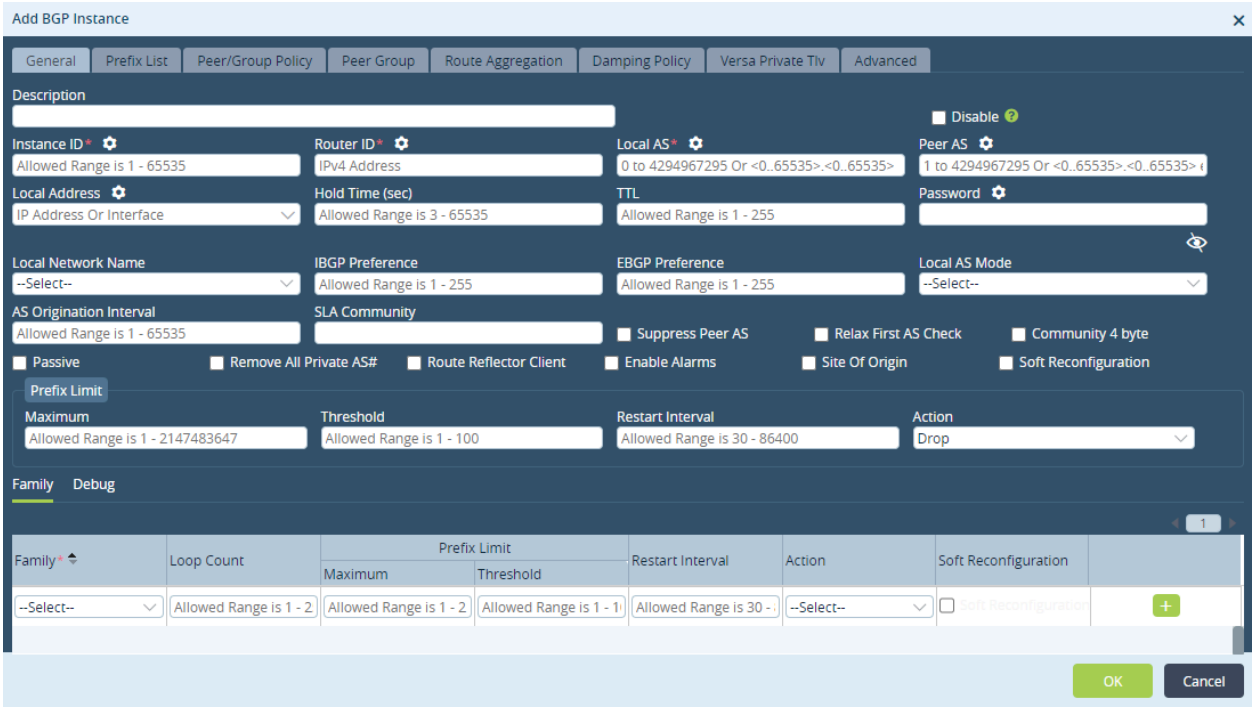

Field	Description
Description	Enter a description for the BGP instance.

Field	Description
Disable	Disable all static neighbors in the BGP instance
Instance ID	Enter the ID assigned of the BGP instance. A virtual router can have multiple BGP instances. *Range:* 1 through 65535
Router ID	Enter the IP address of the router. Click the Tool icon to parameterize the value for this field.
Local AS	Enter the local AS number. Click the Tool icon to parameterize the value for this field. *Range:* 0 through 4294967295, or (0 through 65535.0 through 65535)
Peer AS	Enter the remote peer's AS number. Click the Tool icon to parameterize the value for this field. *Range:* 0 through 4294967295, or (0 through 65535.0 through 65535)
Local Address	Select the IP address or interface of the BGP instance. Click the Tool icon to parameterize the value for this field.
Hold Time	Enter the hold time, in seconds, to negotiate with a peer. *Range:* 3 through 65535 seconds
TTL	Enter the time-to-live value, which is the number of hops that a packet can travel in a network before the packet expires. *Range:* 1 through 255 *Default for EBGP:* 64 (Note that you do not need to enable EBGP multihop.) *Default for IBGP:* 64
Password	Enter the password to authenticate the BGP instance.
Local Network Name	Select the network to which the BGP instance belongs. The drop-down lists the names of user-defined networks.
IBGP Preference	Enter the preference value to assign to routes learned from IBGP. *Range:* 1 through 255 *Default:* 200

Field	Description
EBGP Preference	Enter the preference value to assign to routes learned from EBGP. *Range:* 1 through 255 *Default:* 20
Local AS Mode	Select the BGP AS mode to use on the local device: ◦ 1—Peering session is established with the local AS configured in BGP instance or with a BGP group or neighbor. When importing routes, an AS number is not inserted in the AS path. When exporting routes, the selected local AS number is prepended to the AS path. ◦ 2—Peering session is established with local AS configured as a BGP group or neighbor. When importing routes, the local AS number of the group or neighbor is inserted in AS path. When exporting routes, the local AS number configured on the BGP group or neighbor and the local AS number configured for the BGP instance are prepended to the AS path. This is the default. ◦ 3—Peering session is established with the local AS configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS configured for the BGP group or neighbor and and the local AS number configured for the BGP instance are prepended to AS path. ◦ 4—Peering session is established with the local AS number configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS number configured for the BGP group or neighbor is prepended to the AS path. *Default:* 2
AS Origination Interval	Enter the minimum time that must elapse between successive advertisements of Update messages that report changes within the advertising BGP speaker's own AS. *Range:* 1 through 65535
SLA Community	Enter the extended community value to use for SLAs.
Suppress Peer AS	(For Releases 21.2.1 and later.) Click to suppress and not advertise routes received from an EBGP neighbor to another neighbor that is in the same AS as originating neighbor.
Relax First AS Check	(For Releases 21.2.1 and later.) Click to relax the check that the first (left-most) AS number in an AS path receives from an EBGP neighbor is same as neighbor's AS.
Community 4 Byte	(For Releases 21.2.1 and later.) Click to process the community using 4-byte AS

Field	Description
	numbers.
Passive	Click to have BGP only accept traffic, and not transmit routes.
Remove All Private AS Numbers	Click to remove all private AS numbers from a route's AS path of the route before sending the route to peers.
Route Reflector Client	Click to have an IBGP router function as a route reflector. A route reflector broadcasts the routes of all the other routers in the network
Enable Alarms	Click to enable alarm generation.
Site of Origin	Enter the site-of-origin community to use while redistributing routes.
Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.
Prefix Limit (Group of Fields)	
◦ Maximum	Enter the maximum number of received prefixes. *Range:* 1 through 2147483647
◦ Threshold	Enter the prefix limit number to reach before taking the configured action. *Range:* 1 through 100
◦ Restart Interval	Enter the time, in seconds, to re-establish a session that exceeds prefix limit. ◦ *Range: 30* through 86400
◦ Action	Select the action to take when the prefix limit exceeds: ◦ Drop ◦ Warn
Family (Group of Fields)	Configure the following fields, and then click the Add icon to add a BGP family.
◦ Family	Select the protocol for the family:

Field	Description
	◦ IPv4 Unicast—Applicable to BGP. ◦ IPv4 Multicast—Applicable to BGP. ◦ IPv4 Versa Private—Applicable to SD-WAN. ◦ IPv4 Layer 3 VPN Unicast —Applicable to Layer 3 VPN. ◦ IPv4 VPN Multicast—Applicable to Layer 3 VPN. ◦ IPv6 Unicast—Applicable to BGP. ◦ IPv6 Multicast—Applicable to BGP. ◦ IPv6 VPN Unicast—Applicable to Layer 3 VPN. ◦ IPv6 VPN Multicast—Applicable to Layer 3 VPN. ◦ L2VPN EVPN—Applicable for Layer 2 VPN.
◦ Loop Count	Enter a loop value. A family is considered to be a loop if the number of neighboring ASs is more than this loop value. *Range:* 1 through 255
◦ Prefix Limit	Enter the prefix limit that a BGP instance can receive per session from its peer. ◦ Maximum—Enter the maximum prefix limit. *Range:* 1 through 2147483647 ◦ Threshold—Enter the threshold prefix limit. *Range:* 1 through 100
◦ Restart Interval	Enter the time, in seconds, to re-establish a session that exceeds the prefix limit. *Range:* 30 through 86400
◦ Action	Select the action to take when the prefix limit is exceeded: ◦ Drop ◦ Warn
◦ Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.

9. Click OK.

Configure BGP Prefix Lists

Peer group policy uses prefix lists to change the attributes of routes, and to allow or deny advertising routes to the peer routers.

To configure a prefix list:

1. In the Configure Virtual Router > BGP window, select the Prefix List tab. The Add Prefix List popup window displays.
2. In the Prefix List Name field, enter a name for the prefix list.

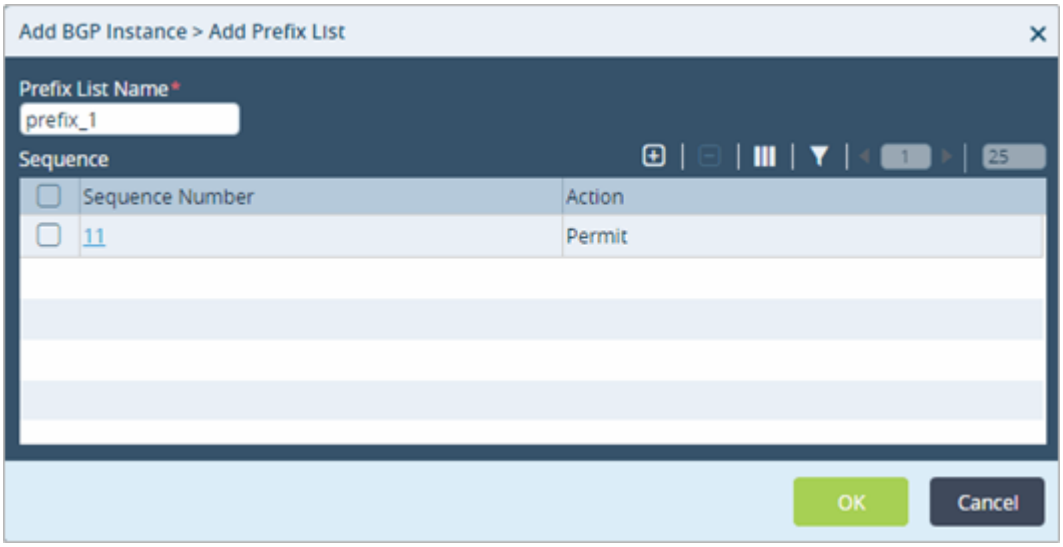

3. Click the Add icon to configure the prefix list. Enter information for the following fields.

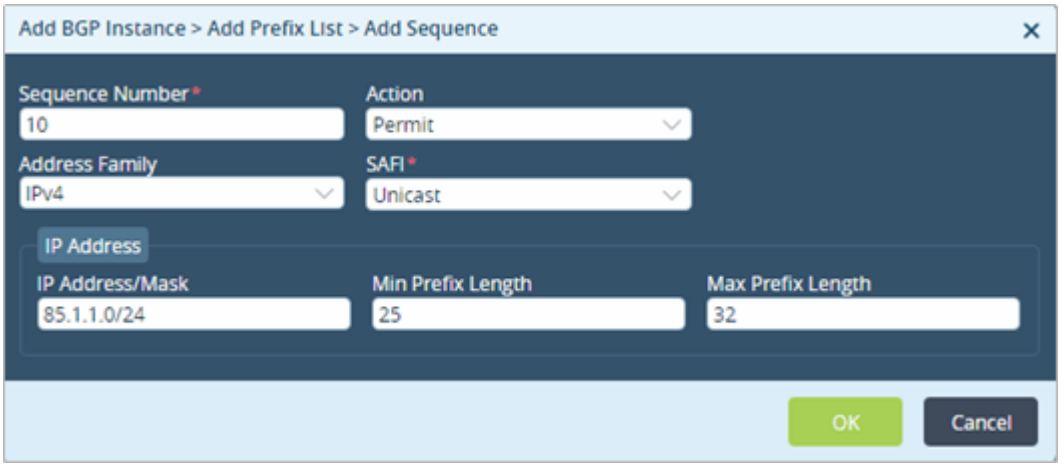

Field	Description
Sequence Number (Required)	Enter a number for the order or sequence number of the prefix list.
Action	Select the action to take on the routes: ◦ Deny—Select this to deny routes on this prefix list. ◦ Permit—Select this to allow routes on this prefix list.
Address Family	Select the broadcast family protocol of the route: ◦ IPv4 ◦ IPv6
SAFI (Required)	Select the subaddress family indicator.
IP Address (Group of Fields)	Configure an IP address to group the routes for this prefix list.
◦ IP Address/Mask	Enter the IP address of the routes grouped in this prefix list.
◦ Min Prefix Length	Enter the minimum number of prefix length to match. *Range:* 25 through 32
◦ Max Prefix Length	Enter the maximum number of prefix length to match. *Range:* 25 through 32

4. Click OK. The Add Prefix List window displays the prefix list.

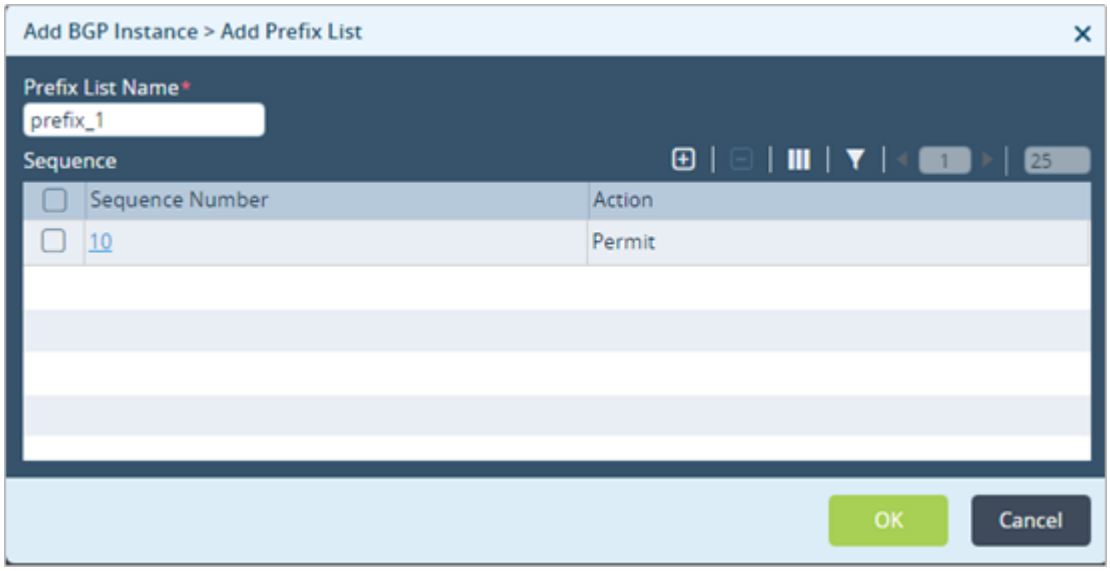

Configure Peer and Peer Group Policy

Peer group policy filters the routes defined in the prefix list. For matching routes, policy can change route attributes. The policy allows or denies advertising these routes to the peers.

To configure peer and peer group policy:

1. In the Configure Virtual Router > BGP window, select the Peer/Group Policy tab. The Add Add Peer/Group Policy popup window displays.
2. Click the Add icon, and enter a name for the peer group policy.

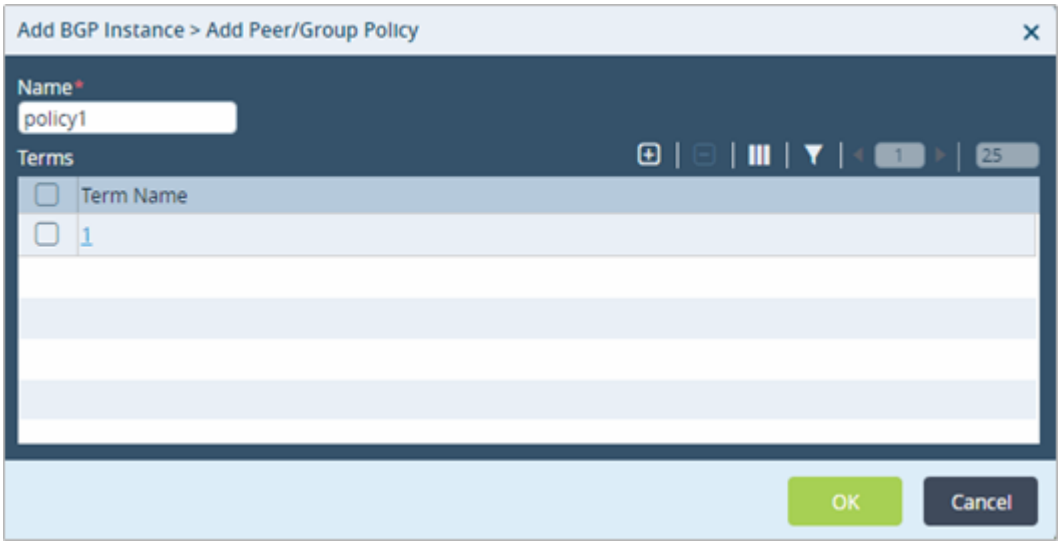

3. Click the Add icon, and enter information for the following fields.

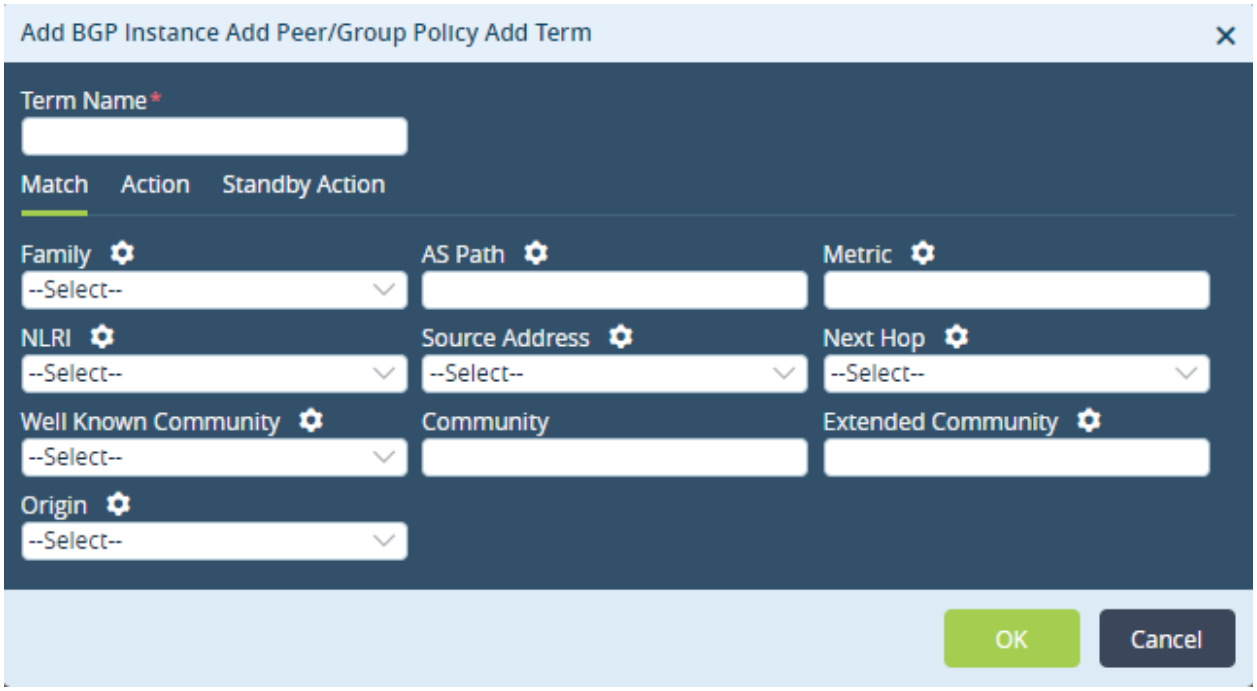

Field	Description
Term Name (Required)	Enter the name of the policy term. A policy executes terms in the order in which they are listed in the Term Name table.

4. Select the Match tab to enter information for match criteria. For all these fields, click the Tool icon to parameterize the value for this field.

Field	Description
Family	Select the protocol family to match: ◦ IPv4 Unicast—Use for BGP ◦ IPv4 Multicast—Use for BGP ◦ IPv4 Versa Private—Use for SD-WAN ◦ IPv4 Layer 3 VPN Unicast—Use for Layer 3 VPN ◦ IPv4 VPN Multicast—Use for Layer 3 VPN ◦ IPv6 Unicast—Use for BGP ◦ IPv6 Multicast—Use for BGP ◦ IPv6 VPN Unicast— Use for Layer 3 VPN ◦ IPv6 VPN Multicast—Use for Layer 3 VPN ◦ L2VPN EVPN—Use for Layer 2 VPN

Field	Description
AS Path	Enter the AS path number to match.
Metric	Enter the BGP multiple exit discriminator (MED) value to match.
NLRI	Select the network layer reachability information (NLRI) from the user-defined prefix lists to match.
Source Address	Select the source address of the prefix list from the list of user-defined prefix lists to match.
Next Hop	Select the IP address of the prefix list to as the next hop from the user-defined prefix lists to match.
Well-Known Community	(For Releases 21.2.1 and later.) Select a well-known community: ◦ no_advertise—A BGP speaker that receives a route containing this community value must not advertise the route to any external or internal peer. ◦ no_export—A BGP speaker that receives a route containing this community value must not advertise the route to its external BGP peers. However, the BGP speaker can advertise the route to its IBGP peers and to confederation peers in other member ASs within its local confederation. ◦ no_export_subconfed—A BGP speaker that receives a route containing this community value must not advertise the route to any external peer including peers in other member ASs within its confederation. If you select a well-known community, you cannot also configure a string in the Community field.
Community	Enter the BGP community string to match. A BGP community is a group of destinations with a common property. This path attribute in BGP update messages identifies community members and performs actions at a group level instead of to an individual level. BGP communities help identify and segregate BGP routes, enabling a smooth traffic flow. If you configure a community string, you cannot also select a community in the Well-Known Community field.
Extended Community	Enter the extended BGP community string to match. In an extended community, you can group a larger number of destinations than in a community.

Field	Description
Origin	Select the source of the route: ◦ Local EGP ◦ Remote IGP ◦ Unknown Heritage

5. Select the Action tab to enter information for the action to take for routes that meet the match criteria. For most of these fields, click the Parameterize icon to parameterize the value for this field.

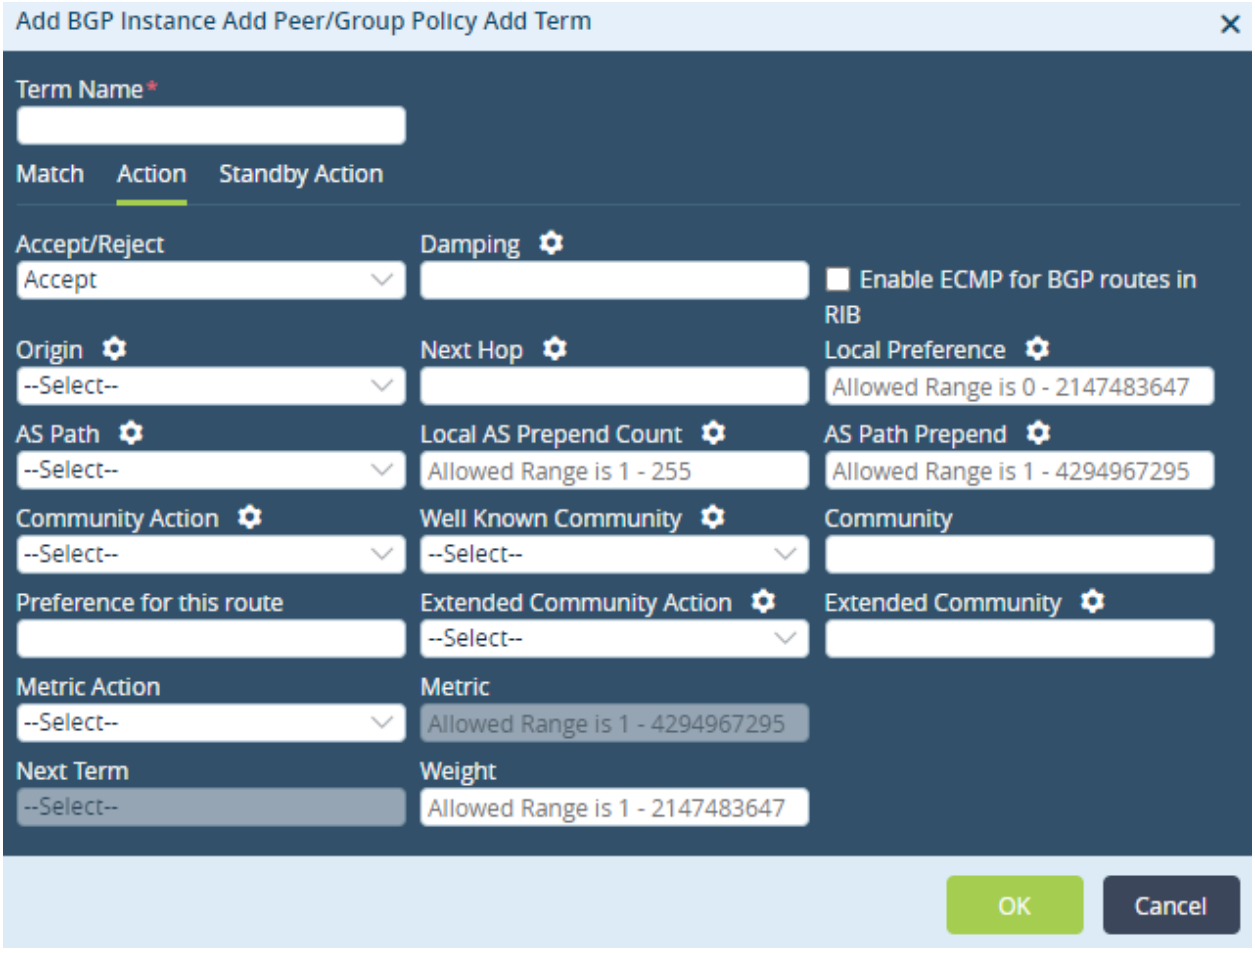

Field	Description
Accept/ Reject	Select whether to either accept or reject the route.

Field	Description
Damping	Enter a value for BGP route-flap damping.
Enable ECMP for BGP Routes in RIB	Select to perform ECMP for BGP paths in the route table. BGP performs equal-cost multipath load when two or more routes have the same administrative distance.
Origin	Select the source of the route: ◦ Local EGP ◦ Remote IGP ◦ Unknown Heritage
Next Hop	Enter the IP address of the next hop.
Local Preference	Enter local preference value to use to choose the outbound external BGP path. *Range:* 0 through 2147483647
AS Path	Select a regular expression to match the AS path for the route: ◦ No AS path action. ◦ Prepend the local AS path the number of times specified by local-as-prepend-count. ◦ Remove all AS numbers matched by match as-path. ◦ Remove all AS numbers matched by match as-path and prepend the local AS the number of times specified by the local-as-prepend-count.
Local AS Prepend Count	Enter the number of times to prepend the local AS number to the AS path. *Range: 1* through 255
AS Path Prepend	Select how to prepend the AS number to an AS path. *Range: 1* through 4294967295
Community Action	Select how to match the community list for a route: ◦ Community field is ignored. ◦ Remove all communities from the route.

Field	Description
	◦ Remove all communities with the value of set-community. ◦ Remove all communities that match set-extended-community. ◦ Append the value of set-community into the communities list.
Well-Known Community	(For Releases 21.2.1 and later.) Select a well-known community: ◦ no_advertise—A BGP speaker that receives a route containing this community value must not advertise the route to any external or internal peer. ◦ no_export—A BGP speaker that receives a route containing this community value must not advertise the route to its external BGP peers. However, the BGP speaker can advertise the route to its IBGP peers and to confederation peers in other member ASs within its local confederation. ◦ no_export_subconfed—A BGP speaker that receives a route containing this community value must not advertise the route to any external peer including peers in other member ASs within its confederation. If you select a well-known community, you cannot also configure a string in the Community field.
Community	Enter the BGP community string to match. A BGP community is a group of destinations with a common property. This path attribute in BGP update messages identifies community members and performs actions at a group level instead of to an individual level. BGP communities help identify and segregate BGP routes, enabling a smooth traffic flow. If you configure a community string, you cannot also select a community in the Well-Known Community field.
Preference for this Route	Enter the preference value for the route. *Range:* 1 through 255
Extended Community Action	Select how to match the extended community list for a route: ◦ Community field is ignored. ◦ Remove all communities from the route. ◦ Remove all communities with the value of set-community. ◦ Remove all communities that match set-extended-community. ◦ Append the value of set-community into the communities list.

Field	Description
Extended Community	Enter the BGP extended community value.
Metric Action	Select a metric action to take: ◦ Add—Add constant to attribute. ◦ IGP—Track the IGP metric. ◦ Set value—Absolute value of the metric to set. ◦ Subtract—Subtract a constant value from the attribute.
Metric	Enter the metric value for the route. *Range:* 1 through 4294967295
Next Term	Select the name of the next term to evaluate.
Weight	(For Releases 21.2.1 and later.) Enter a weight value to use when receiving routes from all neighbors in the peer group. The weight value is a locally significant parameter. *Range:* 1 through 2147483647 *Default:* 0

6. Select the Standby Action tab to enter information for the action to take when the VOS device is an interchassis high availability (HA) standby. For most of these fields, click the Parameterize icon to parameterize the value for this field.

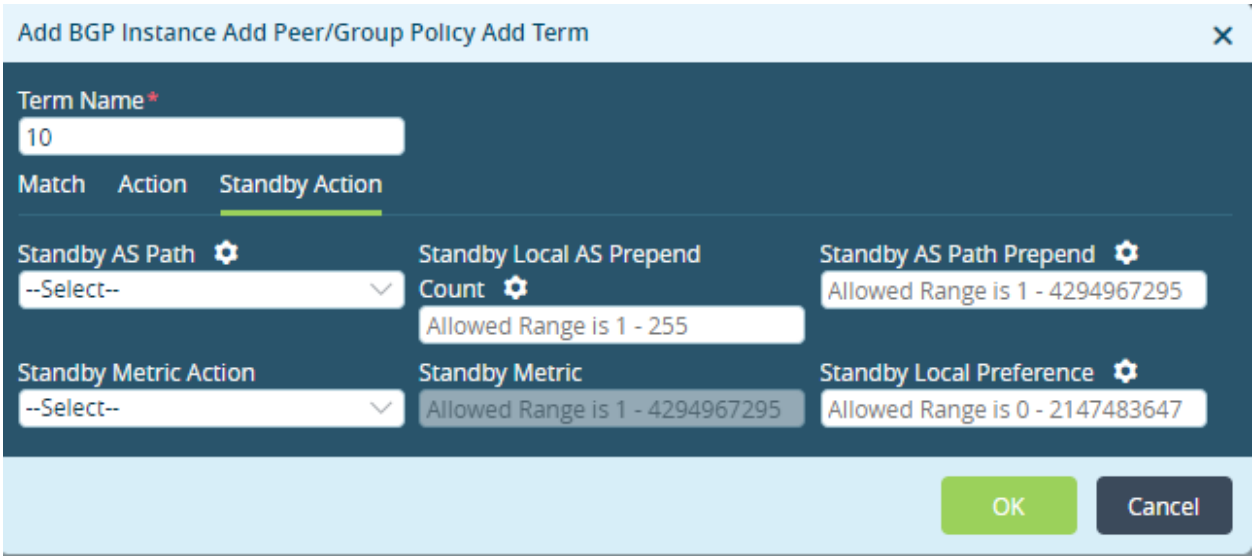

Field	Description
Standby AS Path	Select the AS path action to take when the Versa Operating System™ (VOS™) device is an interchassis HA standby device: ◦ No AS path action. ◦ Prepend the local AS the number of times specified by the local AS prepend count value. ◦ Remove all AS numbers that match the AS path. ◦ Remove all AS numbers that match the AS path and prepend the local AS the number of times specified by the local AS prepend count value.
Standby Local AS Prepend Count	Enter number of times local AS number is prepended to the AS path. *Range:* 1 through 255
Standby AS Path Prepend	Enter prepend specified AS numbers to an AS path while the device is a inter-chassis HA standby. *Range:* 1 through 4294967295
Standby Metric Action	Enter the metric value to be set when the VOS device is a interchassis HA standby device: ◦ Add—Add constant to attribute. ◦ IGP—Track the IGP metric.

Field	Description
	◦ Set value—Absolute value of the metric to set. ◦ Subtract—Subtract a constant value from the attribute.
Standby Metric	Enter the metric value. *Range:* 1 through 4294967295
Standby Local Preference	Enter the local preference for the route while the device is a inter-chassis HA standby. *Range:* 1 through 2147483647

7. Click OK. The Add BGP Instance popup window shows the configured peer and group policies.

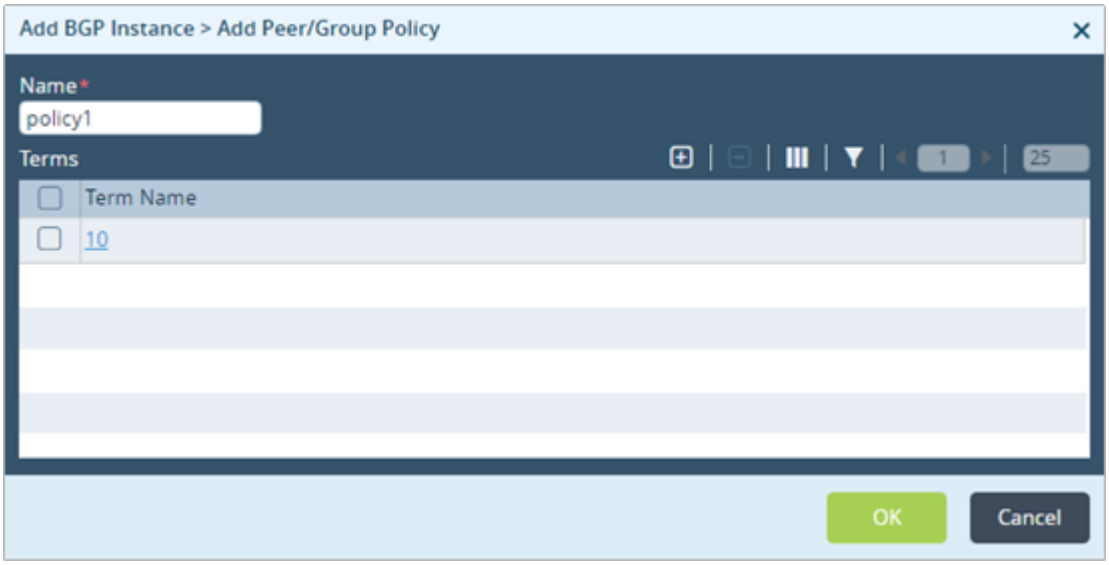

Configure a BGP Peer Group

1. In the Configure Virtual Router > BGP window, select the BGP Peer Group tab.
2. Click the Add icon. In the Add Peer Group popup window, enter information for the following fields.

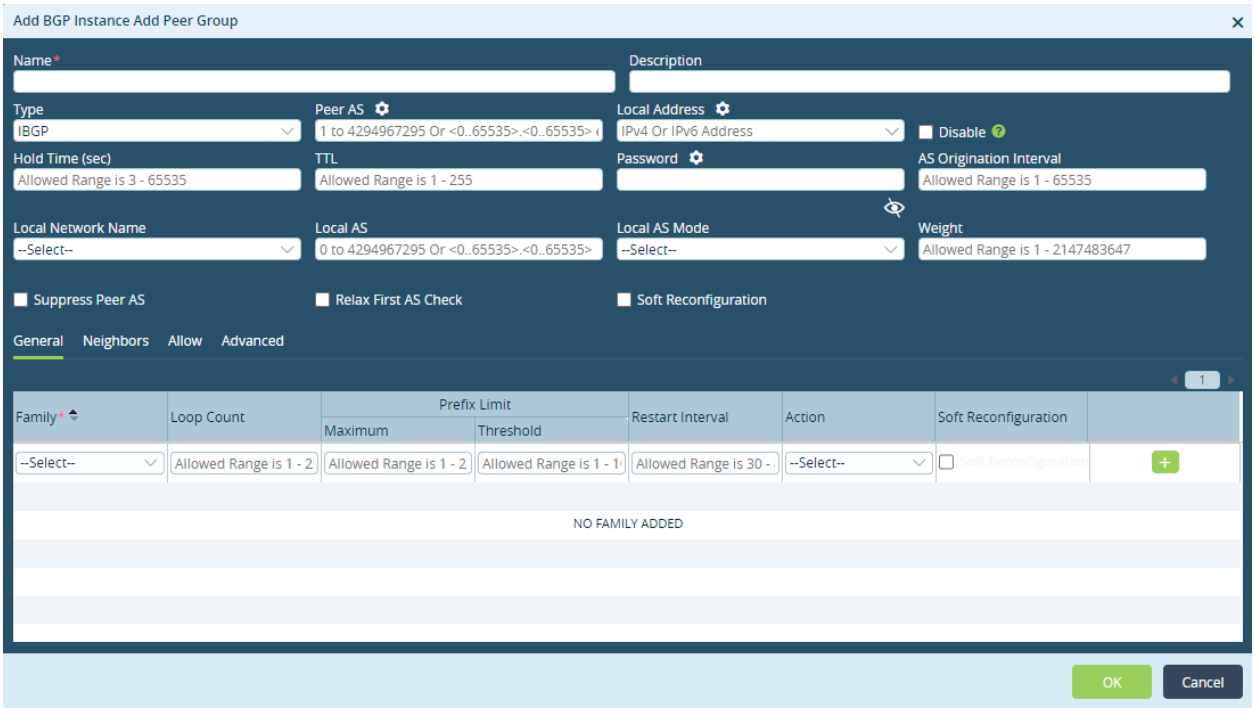

Field	Description
Name	Enter a name for the peer group.
Description	Enter a text description for the peer group.
Type	Select the type of peer group: ◦ EBGP ◦ IBGP
Peer AS	Enter the peer's AS number.
Local Address	Enter the IP address of the local end of a BGP session.
Disable	Click to disable all static neighbors in the BGP group.
Hold Time	Enter the hold time used when negotiating with a peer. *Range:* 3 through 65535
TTL	Enter the time-to-live value, which is the number of hops a packet can travel in a network before the packet expires.

Field	Description
	Range: 1 through 255 *Default for EBGP:* 64 (Note that you do not need to enable EBGP multihop.) *Default for IBGP:* 64
Password	Enter the MD5 password used by this peer group.
AS Origination Interval	Enter the minimum time that must elapse between successive advertisements of Update messages that report changes within the advertising BGP speaker's own AS. *Range:* 1 through 65535
Local Network Name	Select the network to which the peer group belongs. Specify the network name or the local address of the peer group.
Local AS	Enter the local AS number.
Local AS Mode	Select the BGP AS mode to use on the local device: ◦ 1—Peering session is established with the local AS configured in the BGP instance or with a BGP group or neighbor. When importing routes, an AS number is not inserted in the AS path. When exporting routes, the selected local AS number is prepended to the AS path. ◦ 2—Peering session is established with local AS configured as a BGP group or neighbor. When importing routes, the local AS number of the group or neighbor is inserted in AS path. When exporting routes, the local AS number configured on the BGP group or neighbor and the local AS number configured for the BGP instance are prepended to the AS path. This is the default. ◦ 3—Peering session is established with the local AS configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS configured for the BGP group or neighbor and and the local AS number configured for the BGP instance are prepended to AS path. ◦ 4—Peering session is established with the local AS number configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS number configured for the BGP group or neighbor is prepended to the AS path. *Default:* 2
Weight	(For Releases 21.2.1 and later.) Enter a weight value to use when receiving routes from all neighbors in the peer group. The weight value is a locally significant parameter. *Range:* 1 through 2147483647

Field	Description
	Default: 0
Suppress Peer AS	(For Releases 21.2.1 and later.) Click to suppress and not advertise routes received from an EBGP neighbor to another neighbor that is in the same AS as originating neighbor.
Relax First AS Check	(For Releases 21.2.1 and later.) Click to relax the check that the first (left-most) AS number in an AS path receives from an EBGP neighbor is same as neighbor's AS.
Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.

3. Select the General tab, and enter information for the following fields.

Field	Description
Family	Select the protocol family for the peer group: ◦ IPv4 Unicast—Applicable to BGP. ◦ IPv4 Multicast—Applicable to BGP. ◦ IPv4 Versa Private—Applicable to SD-WAN. ◦ IPv4 Layer 3 VPN Unicast—Applicable to Layer 3 VPN. ◦ IPv4 VPN Multicast—Applicable to Layer 3 VPN. ◦ IPv6 Unicast—Applicable to BGP. ◦ IPv6 Multicast—Applicable to BGP. ◦ IPv6 VPN Unicast— Applicable for Layer 3 VPN. ◦ IPv6 VPN Multicast—Applicable for Layer 3 VPN. ◦ L2VPN EVPN—Applicable for Layer 2 VPN.
Loop Count	Enter the number of times the local AS is allowed in the received AS path. For example, if you set the loop value to 5, the local AS in received AS paths can appear five times.
Prefix Limit	Enter the prefixes that a BGP instance can receive per session from its peer.
◦ Maximum	Enter the maximum number of received prefixes. *Range:* 1 through 2147483647

Field	Description
◦ Threshold	Enter the prefix limit number to reach before taking the configured action. *Range:* 1 through 100
Restart Interval	Enter the time, in seconds, to re-establish a session that exceeds prefix limit.
Action	Select the action to take when the prefix limit exceeds: ◦ Drop ◦ Warn
Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.

4. Select the Neighbors tab and click the Add icon. In the Add Neighbor popup window, enter information for the following fields.

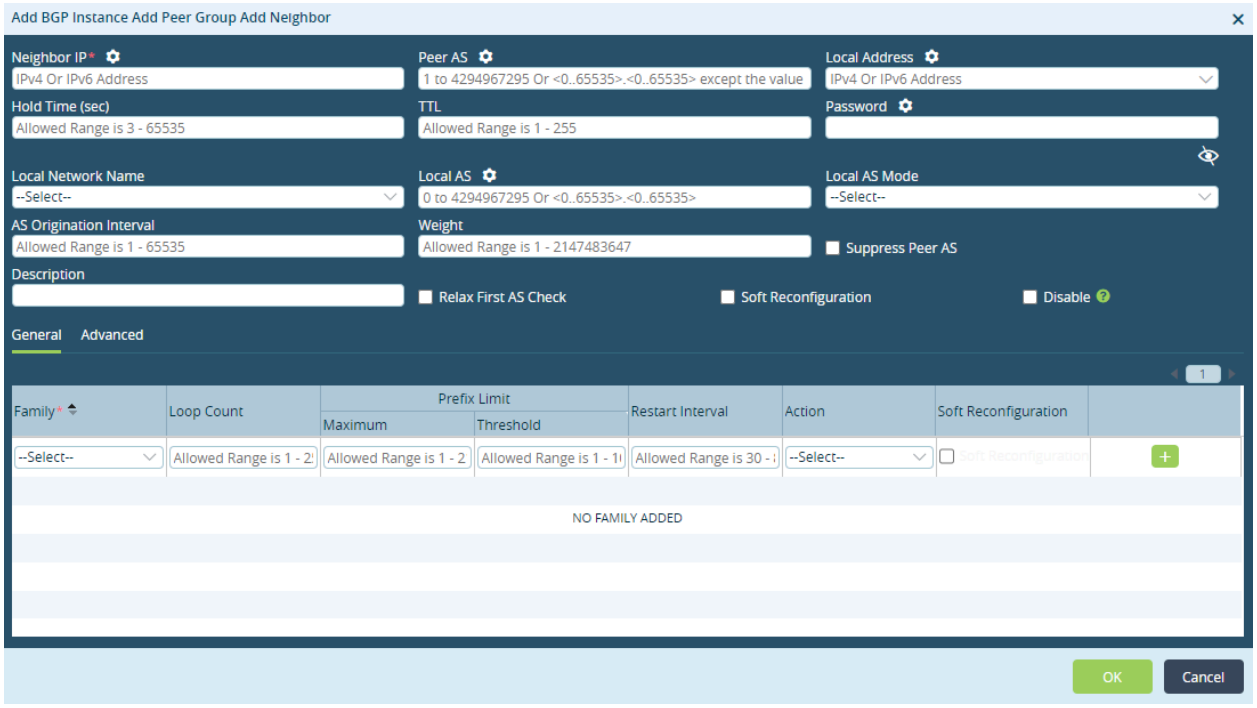

Field	Description
Neighbor IP	Enter the neighbor's IP address.
Peer AS	Enter the peer's AS number.
Local Address	Enter the IP address of the the local end of the BGP session.
Hold Time (sec)	Enter the hold time used when negotiating with a peer. *Range:* 3 through 65535
TTL	Enter the number of hops a packet can travel in a network before the packet expires. *Range:* 1 through 255 *Default for EBGP:* 64 (Note that you do not need to enable EBGP multihop.) *Default for IBGP:* 64
Password	Enter the MD5 password for the neighbor.
Local Network Name	Select the network to which the neighbor peer group belong. Specify the network name or the local address of the peer group.
Local AS	Enter the local AS number.
Local AS Mode	Select the BGP AS mode to use on the local device: ◦ 1—Peering session is established with the local AS configured in BGP instance or with a BGP group or neighbor. When importing routes, an AS number is not inserted in the AS path. When exporting routes, the selected local AS number is prepended to the AS path. ◦ 2—Peering session is established with local AS configured as a BGP group or neighbor. When importing routes, the local AS number of the group or neighbor is inserted in AS path. When exporting routes, the local AS number configured on the BGP group or neighbor and the local AS number configured for the BGP instance are prepended to the AS path. This is the default. ◦ 3—Peering session is established with the local AS configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS configured for the BGP group or neighbor and and the local AS number configured for the BGP instance are prepended to AS path. ◦ 4—Peering session is established with the local AS number configured for the BGP group or neighbor. When importing routes, no AS number is inserted in the AS path. When exporting routes, the local AS number configured for the BGP group or neighbor is prepended to the AS path.

Field	Description
	Default: 2
AS Origination Interval	Enter the minimum amount of time that must elapse between successive advertisements of Update messages that report changes within the advertising BGP speaker's own autonomous system. *Range:* 1 through 65535
Weight	(For Releases 21.2.1 and later.) Enter a weight value to use when receiving routes from all neighbors in the peer group. The weight value is a locally significant parameter. *Range:* 1 through 2147483647 *Default:* 0
Description	Enter a text description for the BGP peer.
Suppress Peer AS	(For Releases 21.2.1 and later.) Click to suppress and not advertise routes received from an EBGP neighbor to another neighbor that is in the same AS as originating neighbor.
Relax First AS Check	(For Releases 21.2.1 and later.) Click to relax the check that the first (left-most) AS number in an AS path receives from an EBGP neighbor is same as neighbor's AS.
Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.
Disable	Disable all static neighbors in the BGP group.
General (Tab)	
◦ Family	Select the protocol family of the neighbor peer group: ◦ IPv4 Unicast—Use for BGP. ◦ IPv4 Multicast—Use for BGP. ◦ IPv4 Versa Private—Use for SD-WAN. ◦ IPv4 Layer 3 VPN Unicast—Use for Layer 3 VPN. ◦ IPv4 VPN Multicast— Use for Layer 3 VPN. ◦ IPv6 Unicast—Use for BGP. ◦ IPv6 Multicast—Use for BGP.

Field	Description
	◦ IPv6 VPN Unicast— Use for Layer 3 VPN. ◦ IPv6 VPN Multicast—Use for Layer 3 VPN. ◦ L2VPN EVPN—Use for Layer 2 VPN.
◦ Loop Count	Enter the number of times the local AS is allowed in the received AS path.
◦ Prefix Limit	Enter the number of prefixes that a BGP instance can receive per session from its peer.
◦ Maximum	Enter the maximum prefix limit. *Range:* 1 through 2147483647
◦ Threshold	Enter the threshold prefix limit. *Range:* 1 through 100
◦ Restart Interval	Enter the time, in seconds, to re-establish a session that exceeds prefix limit.
◦ Action	Select the action to take when the prefix limit exceeds: ◦ Drop ◦ Warn
Soft Reconfiguration	(For Releases 21.2.1 and later.) Click to keep received routes even if the routing peer policy rejects the routes.
Advanced (Tab)	In the Advanced tab, enter information for the following fields.

Field	Description
	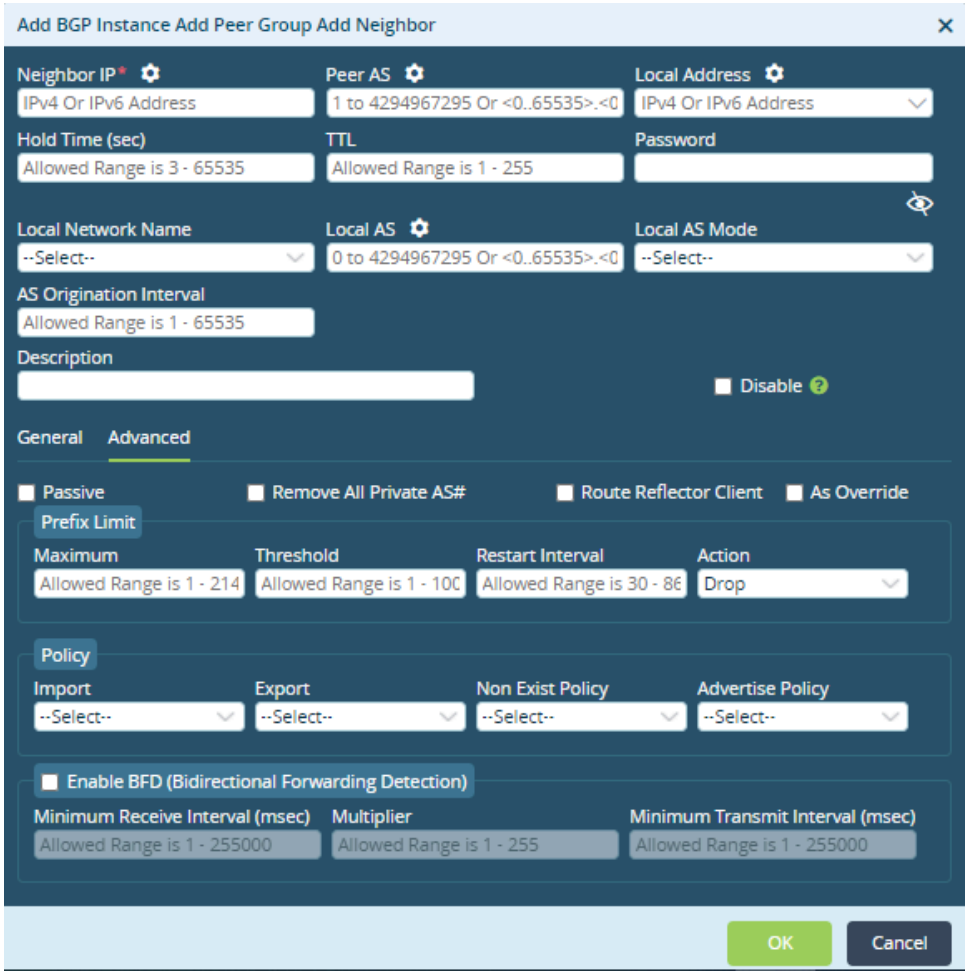

◦ Passive	Select to have BGP to accept traffic only and not to transmit any routes.
◦ Remove All Private AS	Select to remove all private AS numbers before advertising routes.
◦ Route Reflector Client	For IBGP, select to have the router function as a route reflector and to broadcast the routes of all other routers in the network
◦ AS Override	Select to replace neighbor AS numbers with the local AS numbers from the AS path.
Prefix Limit (Group of Fields)	

◦ Maximum	Enter the maximum number of received prefixes. *Range:* 1 through 2147483647
◦ Threshold	Enter the threshold prefix limit before taking the configured action. *Range:* 1 through 100
◦ Restart Interval	Enter the time, in seconds, to re-establish a session that exceeds prefix limit. *Range:* 30 through 86400
◦ Action	Select the action to take when the prefix limit exceeds: ◦ Drop ◦ Warn
Policy (Group of Fields)	
◦ Import	Select the peer group policy to apply to outgoing routing updates.
◦ Export	Select the peer group policy to apply to incoming routing updates.
◦ Non-Exit Policy	Select the policy that defines the route being withdrawn. Routes matching the advertised policy are advertised to peers only if no routes in the local route table match the non-exist policy.
◦ Advertise Policy	Select the policy that defines the route to be advertised to the BGP peer after the route in the non-exist policy is withdrawn.
Enable BFD (bidirectional forwarding detection)	Click to enable BFD on the interface, to allow BFD to report when a static route becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds.
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval.

◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds.

5. Click OK.
6. Select the Allow tab to configure the peer group routes to accept. Enter information for the following fields.

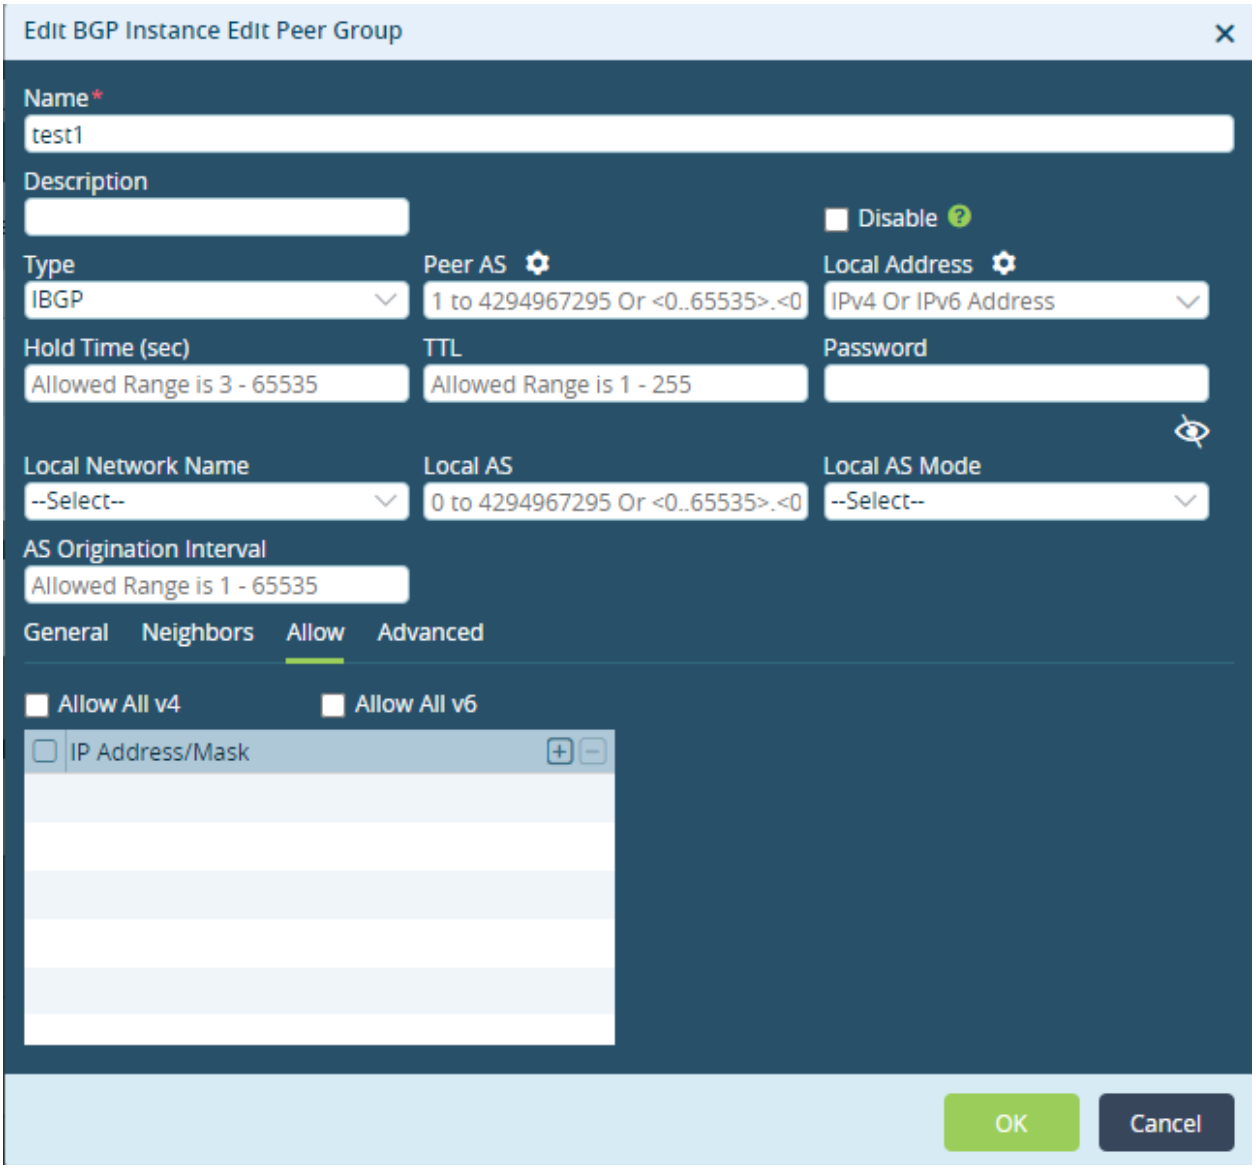

Field	Description
Allow All v4	Click to select all IP addresses as acceptable peer group routes.

Field	Description
Allow All v6	Click to select all IP addresses as acceptable peer group routes.
IP Address/Mask	Click the ⊞ Add icon, and then add the IP address of the route to be allowed.

7. Select the Advanced tab to configure the peer group routes to accept. Enter information for the following fields.

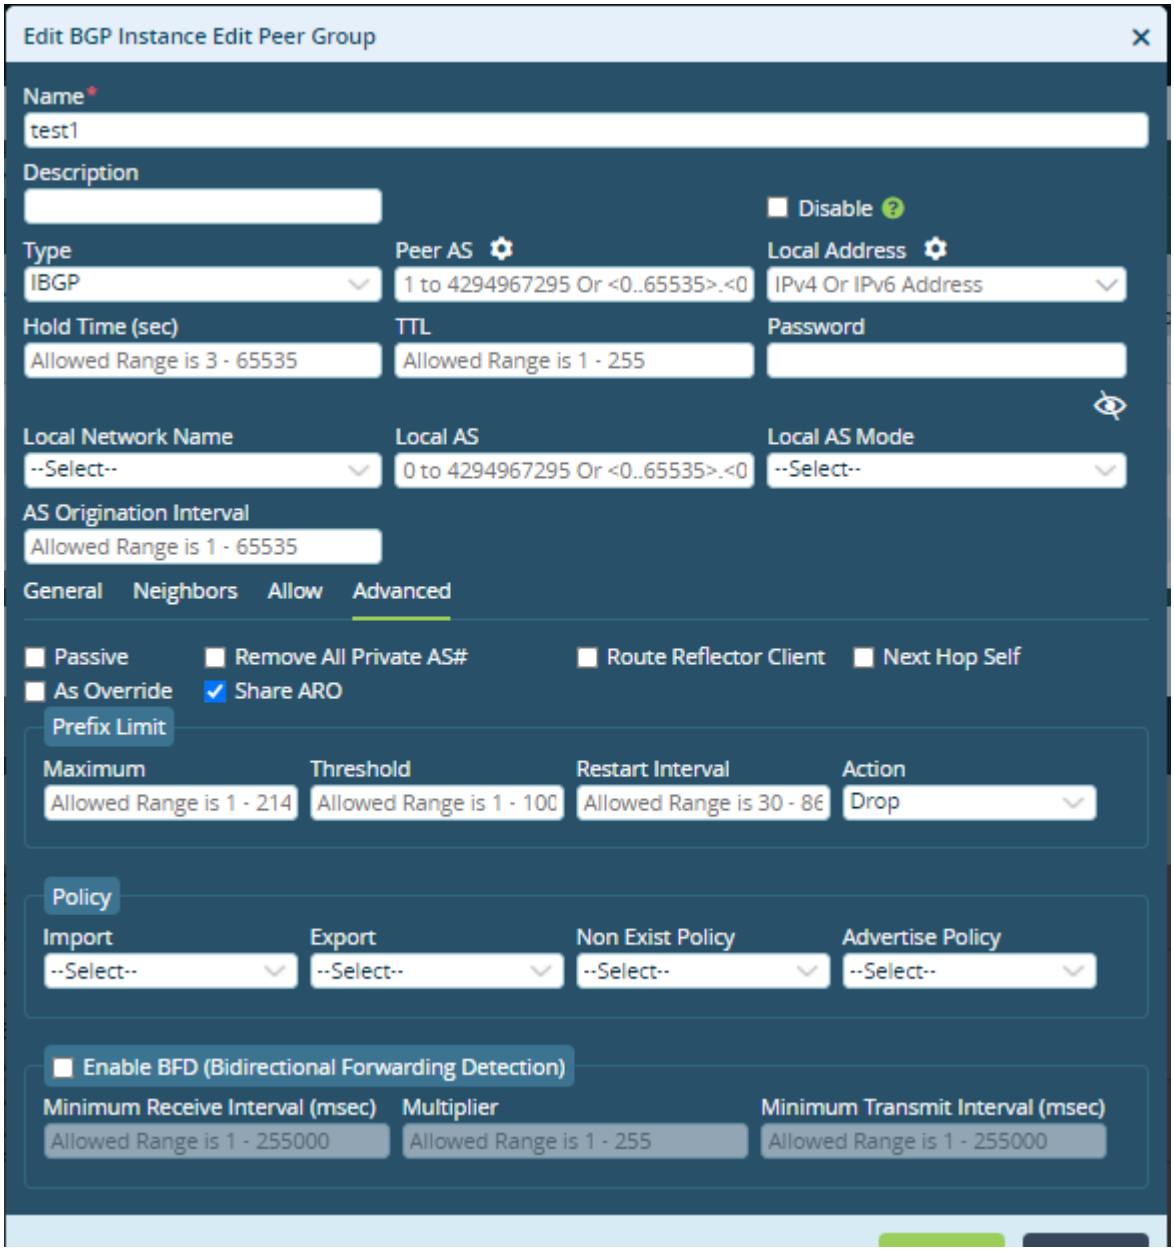

Field	Description
Passive	Click to have BGP to accept traffic only and not to transmit any routes.
Remove All Private AS#	Click to remove all private AS numbers before advertising routes.
Route Reflector Client	For IBGP, click to have the router function as a route reflector and to broadcast the routes of all other routers in the network
Next Hop Self	Click to use the IP address of the prefix list.
AS Override	Click to replace neighbor AS numbers with the local AS numbers from the AS path.
Share ARO	Click to share all the routes of Adj-RIB-Out (ARO) between all neighbors in an IBGP group to save internal memory.
Prefix Limit (Group of Fields)	
◦ Maximum	Enter the maximum number of prefixes that a BGP instance can receive per session from its peer. *Range:* 1 through 2147483647
◦ Threshold	Enter the prefix threshold limit that a BGP instance can receive per session from its peer. *Range:* 1 through 100
◦ Restart Interval	Enter the restart time. *Range:* 30 through 86400
◦ Action	Select the action to take when the prefix limit exceeds: ◦ Drop ◦ Warn

Field	Description
Policy (Group of Fields)	Routes matching this advertise-policy will be advertised to peers only if there are no routes in the local RIB matching the non-exist-policy.
◦ Import	Select the peer group policy to apply to outgoing routing updates.
◦ Export	Select the peer group policy to apply to incoming routing updates.
◦ Non-Exist Policy	(For Releases 16.1R2S11 and later.) Select the name of a policy to use when you are configuring conditional route advertisements. In the Non-Exist Policy field, select the policy that defines the route being withdrawn. For more information, see Configure BGP Conditional Route Advertisements.
◦ Advertise Policy	(For Releases 16.1R2S11 and later.) Select the name of a policy to use when you are configuring conditional route advertisements. In the Advertise Policy field, select the policy that defines the route to be advertised to the BGP peer after the route in the Non-Exist policy is withdrawn. For more information, see Configure BGP Conditional Route Advertisements.
Enable BFD	Click to enable Bidirectional Forwarding Detection on the interface, to allow BFD to report when a static route becomes unavailable.
◦ Minimum Receive Interval (msec)	Enter the minimum time interval to receive routes, in milliseconds.
◦ Multiplier	Enter the multiplier value used to calculate the final minimum receive interval and minimum transmit interval.
◦ Minimum Transmit Interval (msec)	Enter the time after which routes can be retransmitted, in milliseconds.

8. Click OK. The Add BGP Instance popup window displays the configured peer groups.

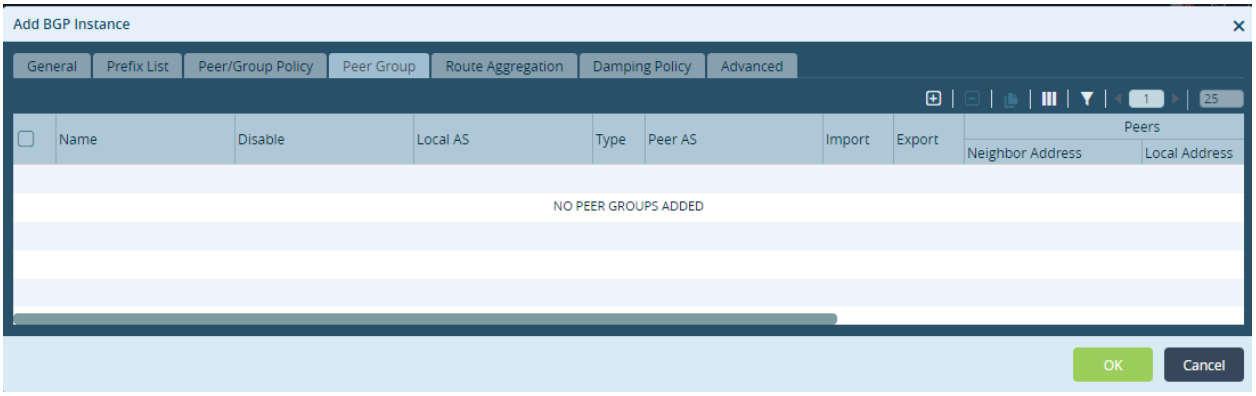

Configure Route Aggregation

Route aggregation is a method to minimize the number of routing tables required in an IP network.

To configure route aggregation:

1. In the Add BGP Instance tab, select the Route Aggregation tab.

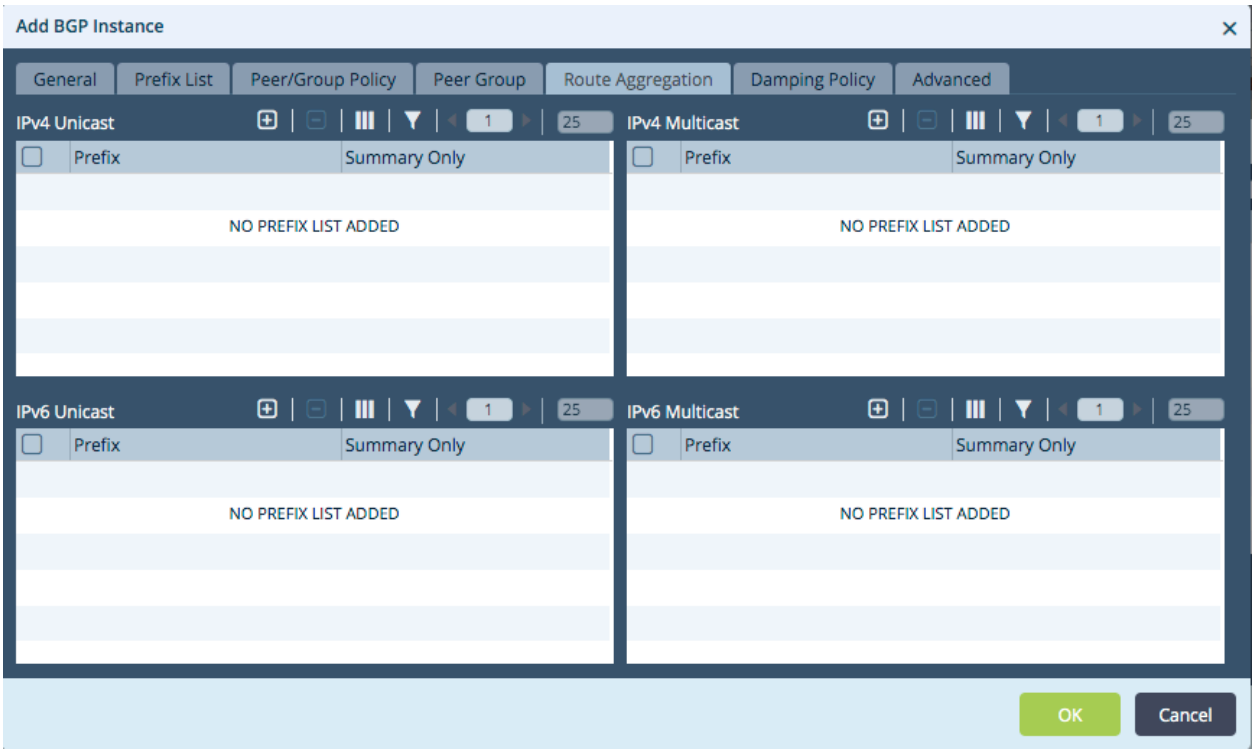

2. Click the Add icon in any of the IPv4/IPv6 Unicast/Multicast sections, and enter information for the following fields.

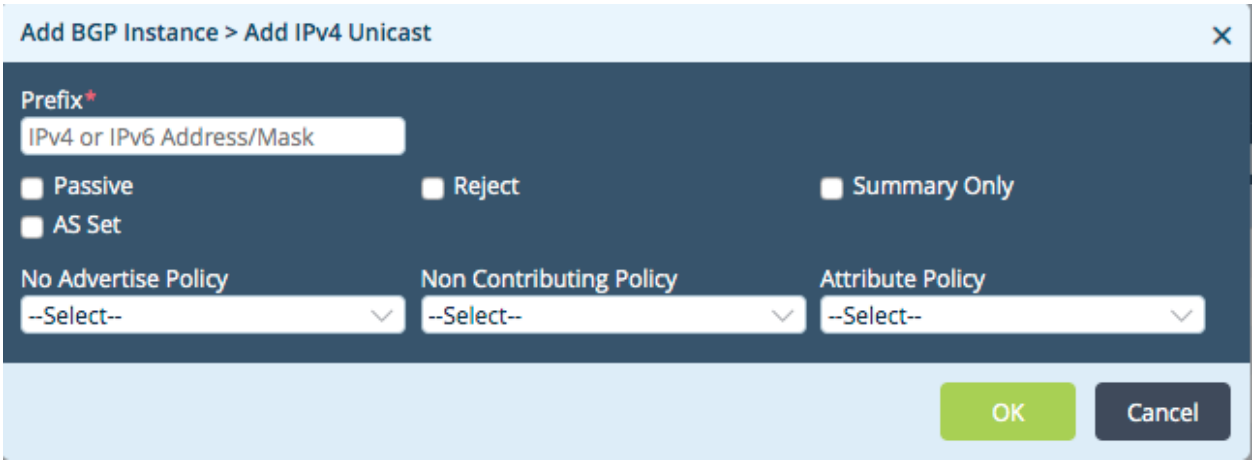

Field	Description
Prefix	Enter the IPv4 or IPv6 address.
Passive	Select this check box to indicate that the router is a passive listener that does not actively propagate messages. A passive router does not advertise itself.
Reject	Select to reject the traffic on the route ted.
Summary Only	Select to have an area border router advertise the area routers.
AS Set	Select to indicate that the AS number is set.
No Advertise Policy	Select the routing peer policy that defines the contributing routes that should not be advertised.
Non Contributing Policy	Select the routing peer policy that defines the set of routes that do not act as contributing routes.
Attribute Policy	Select the name of the policy that allows the setting of aggregated route attributes.

3. Click OK. The Add BGP Instance popup windows displays the configured aggregate routes.

Configure Damping Policy

1. In the Add BGP Instance tab, select the Damping Policy tab.
2. Click the Add icon, and enter information for the following fields.

Field	Description
Damping Name	Enter the name of the damping policy.
Suppress	Enter the cutoff threshold limit. Routes exceeding this level are suppressed.
Maximum Suppress Time (min)	Enter the maximum time a route can be suppressed (held), in minutes
Reuse	Enter the reuse threshold of a suppressed route.
Half Life Ok (min)	Enter the decay half life time, in minutes, to define the stability of the route while it is still reachable
Half Life Ng (min)	Enter the decay half life time, in minutes, to define the stability of the route while it is unreachable.
Maximum Time Ok (min)	Enter the maximum time, in minutes, any memory of a previous instability is retained for a reachable route.
Maximum Time Ng (min)	Enter the maximum time, in minutes, any memory of a previous instability is retained for an unreachable route.

3. Click OK. The Add BGP Instance popup window displays the configured damping policies.

Configure Advanced BGP Settings

1. In the Configure Virtual Router > BGP window, select the Advanced tab.
2. In the Add BGP Instance popup window, enter information for the following fields.

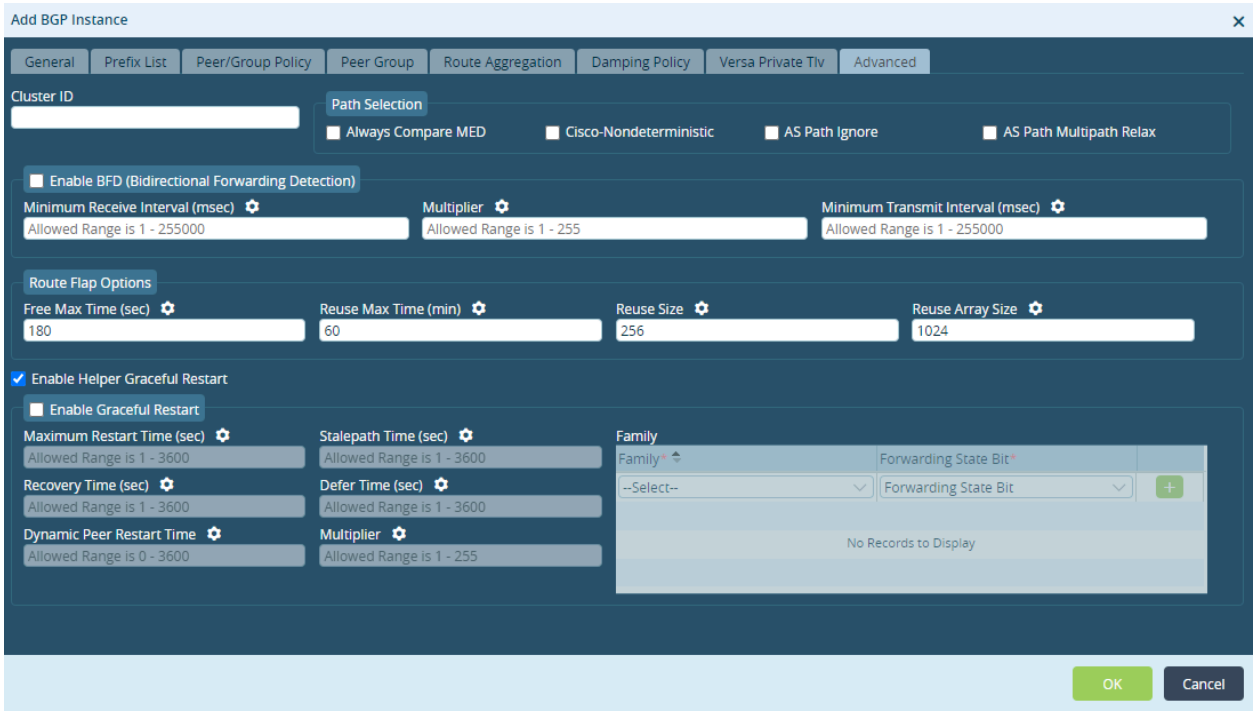

Field	Description
Cluster ID	Enter the cluster ID of the reflector clients.
Path Selection (Group of Fields)	
◦ Always Compare MED	Click to compare multiexit discriminators when sending routes to another router. A route with a lower MED is given priority.
◦ Cisco Nondeterministic	Click to use Cisco nondeterministic path selection. The active path is always first. All non-active, but eligible paths follow the active path and are maintained in the order in which they are received, with the most recent path first. Ineligible paths remain at the end of the list. When a new path is added to the routing table, path comparisons are made without removing from consideration those paths that should not be selected because those paths lose the MED tie-breaking rule.
◦ AS Path Ignore	(For Releases 21.2.1 and later.) Click to exclude the AS path from BGP best path computation.
◦ AS Path Multipath Relax	(For Releases 21.2.1 and later.) Click to allow different AS paths to be considered for multipath if the AS path length is equal.
Enable BFD	Click to mark the link as down whenever the Bidirectional Forwarding Detection (BFD) is down.
◦ Minimum Receive Interval	Enter the time interval, in milliseconds, to mark the link as down if the routing updates are not received.
◦ Multiplier	Enter value to use to compute the final minimum receive interval.
◦ Minimum Transmit Interval	Enter how often BGP instances communicate with each other, in milliseconds.
Route Flap Option (Group of Fields)	
◦ Free Maximum Time	Enter the maximum time to remember an assigned penalty to the router, in seconds. A penalty is assigned to a router when its routes go up and down.

◦ Reuse Maximum Time	Enter the time corresponding to the last reuse list, in minutes.
◦ Reuse Size	Enter the number of reuse lists.
◦ Reuse Array Size	Enter the size of the reuse index arrays.
Enable Graceful Restart (Group of Fields)	Click to enable BGP graceful restart.
◦ Maximum Restart Time	Enter the restart time, in seconds, that is advertised to a neighbor on which graceful restart is enabled. *Range:* 1 through 3600 seconds (1 hour) *Default:* 3600 seconds (for Controller BGP sessions)
◦ Stale Path Time	Enter the maximum time, in seconds, that BGP waits before removing stale routes from a neighbor after a graceful restart of the neighbor's session. Enter the maximum amount of time, in seconds, that a helper device waits for an End-of-RIB marker from a peer. When this time period expires, all stale path are deleted. *Range:* 1 through 3600 seconds (1 hour) *Default:* 3600 seconds (for Controller BGP sessions)
◦ Recovery Time	Enter the estimated recovery time, in seconds, after a restart. *Range:* 1 through 3600 seconds (1 hour) *Default:* 120 seconds
◦ Defer Time	Enter how long, in seconds, a restarting router defers its BGP best-path calculation after a restart occurs and while BGP peering sessions are establishing themselves. This time should should be long enough so that all peers have enough time to send all their routes to the restarting BGP router. *Range:* 1 through 3600 seconds (1 hour) *Default:* 30 seconds (for Controller BGP sessions)
◦ Dynamic Peer Restart Time	Enter the minimum time, in seconds, for a peer to dynamically reconnect after the BGP process restarts. *Range:* 1 through 3600 seconds (1 hour) *Default:* 120 seconds

◦ Multiplier	Enter a multiplier value that, when multiplied by the stale path time, provides a value that is used to hold the routes from a peer that are in RIB/RIB in the stale state after the BGP session with that peer goes down. Note that this multiplier value is not related to the BFD multiplier. *Range:* 1 through 255 *Default:* 8 (for Controller BGP sessions)
Family (Group of Fields)	Protocol family for NLRIs in updates.
◦ Family	Select the protocol: ◦ IPv4 Unicast—Use for BGP ◦ IPv4 Multicast—Use for BGP ◦ IPv4 Versa Private—Use for SD-WAN ◦ IPv4 Layer 3 VPN Unicast—Use for Layer 3 VPN ◦ IPv4 VPN Multicast—Use for Layer 3 VPN ◦ IPv6 Unicast—Use for BGP ◦ IPv6 Multicast—Use for BGP ◦ IPv6 VPN Unicast— Use for Layer 3 VPN ◦ IPv6 VPN Multicast—Use for Layer 3 VPN ◦ L2VPN EVPN—Use for Layer 2 VPN
◦ Forwarding State Bit	Select the forwarding state bit to preserve the forwarding state associated with the AFI/SAFI.

3. Click OK. The Configure Virtual Router window displays the BGP instance.

Configure BGP Conditional Route Advertisements

For Releases 16.1R2S11 and later.

To have BGP to advertise a route based on a specific situation, you configure a conditional route advertisement. To configure the conditional route, you define the prefixes and then you create policies and apply the policies to the BGP peer group.

To illustrate how to configure BGP conditional route advertisements, let's consider a situation where you want BGP to advertise a route only when another route is withdrawn and becomes unavailable. For this example, you configure the following:

- Two prefix lists, one for the unavailable prefix and one for the backup prefix to advertise.

- Two policies, one with a term that matches the unavailable prefix (which here is called *Non-exist*) and one with a term that matches the backup prefix to advertise (which here is called *Advertise*).
- A BGP peer group that refers to the two policies

In the example, the two routes are for the prefixes 192.168.0.0/24, which is the primary prefix that is used in the *non-exist* policy, and 192.168.1.0/24, which is the backup prefix used in the *advertise* policy.

To configure a BGP conditional route advertisement:

1. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
2. Select the Configuration tab in the top menu bar.
3. Select Networking > Virtual Routers in the left menu bar
4. Select an existing tenant, or click the Add icon to add a new tenant. The Configure Virtual Router popup window displays.
5. Select the BGP tab in the left menu bar. The main pane displays a list of the BGP instances that are already configured.
6. Click the Add icon. The Add BGP Instance popup window displays.
7. Select the Prefix List tab. The Add Prefix List popup window displays. Here, you create two prefix lists, *Non-Existent*, for the prefix 192.168.0.0/24, and *Advertise*, for 192.168.1.0/24.
8. To create the *Non-Exist* prefix list, click the Add icon and name the prefix list. Then, add a prefix list sequence.

   Click the Add icon, and then select the action Permit and enter the prefix 192.168.0.0./24. Make sure that you take into account the AFI/SAFI.

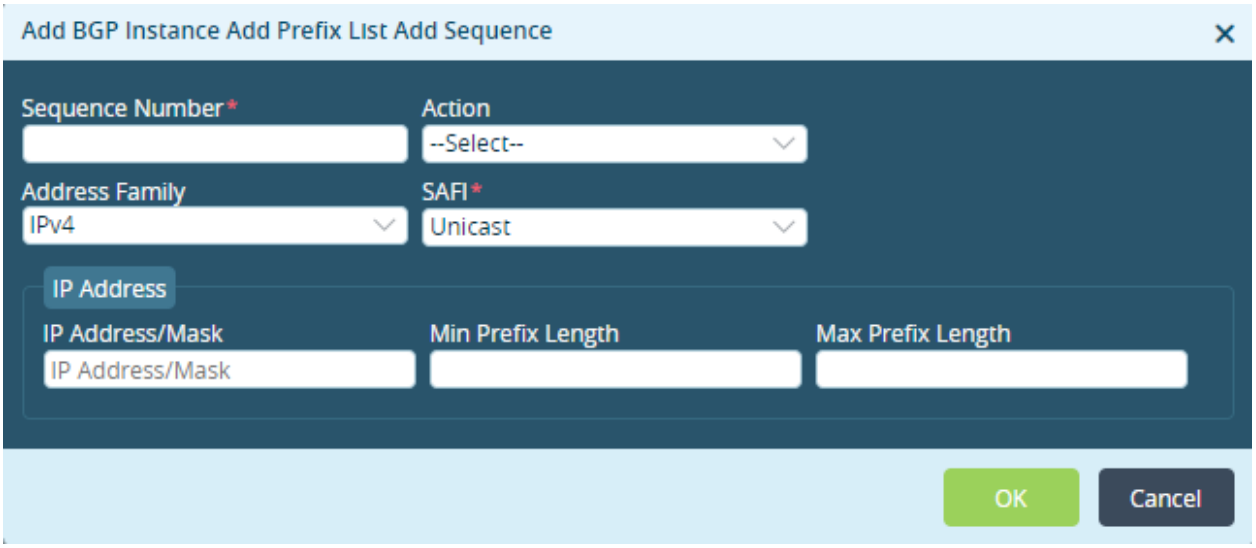

9. Click OK.

10. To create the *Advertise* prefix list, click the Add icon and name the prefix list. Then, add a prefix list sequence.

    Click the Add icon, and then select the action Permit and enter the prefix 192.168.1.0./24. Make sure that you take into account the AFI/SAFI.

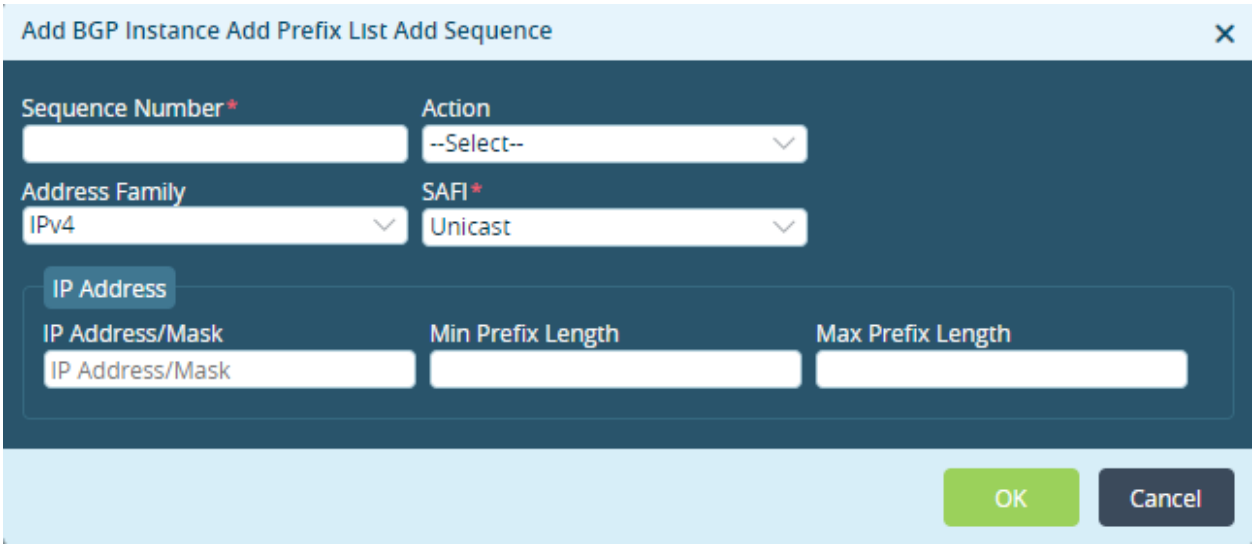

11. Click OK.
12. Click OK to return to the BGP Instance popup window.
13. Select the Peer/Group Policy tab. Here, you create two peer group policies, one for the *Non-Exist* prefix list and the second for the *Advertise* prefix list. You must create both policies.

14. To create the *Non-Exist* peer group policy, click the Add icon and name the policy. Then, create a term. Click the Add icon. Name the term, and in the Match tab, in the NLRI field, select the *Non-Existent* prefix list. In the Action tab, use the default action, which is Accept.

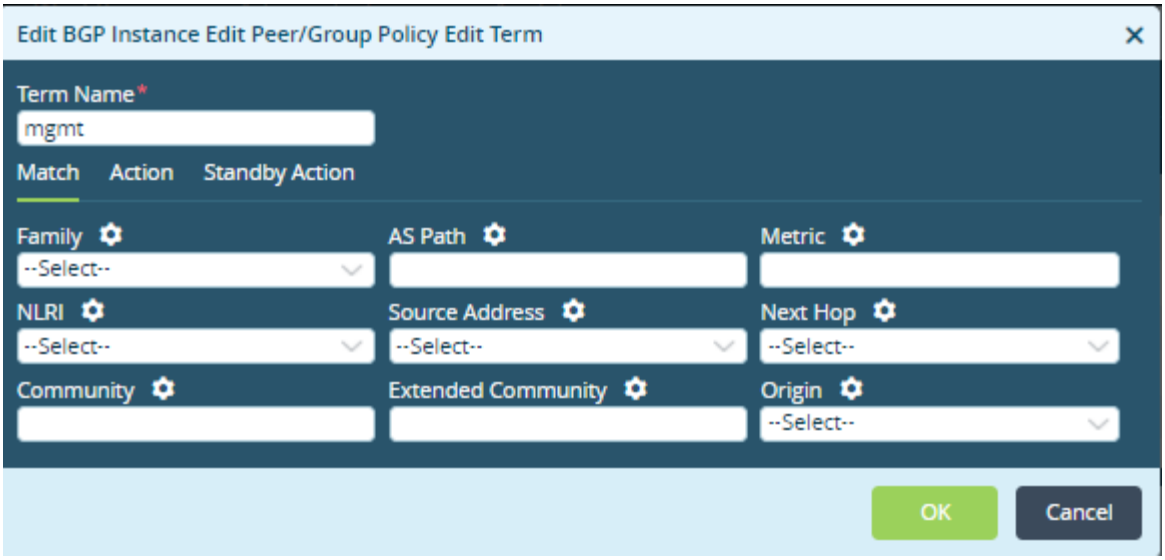

15. Click OK.

16. To create the *Advertise* peer group policy, click the Add icon and name the policy. Then, create a term. Click the Add icon, and in the Match tab, in the NLRI field, select the *Advertise* prefix list. In the Action tab, use the default action, which is Accept.

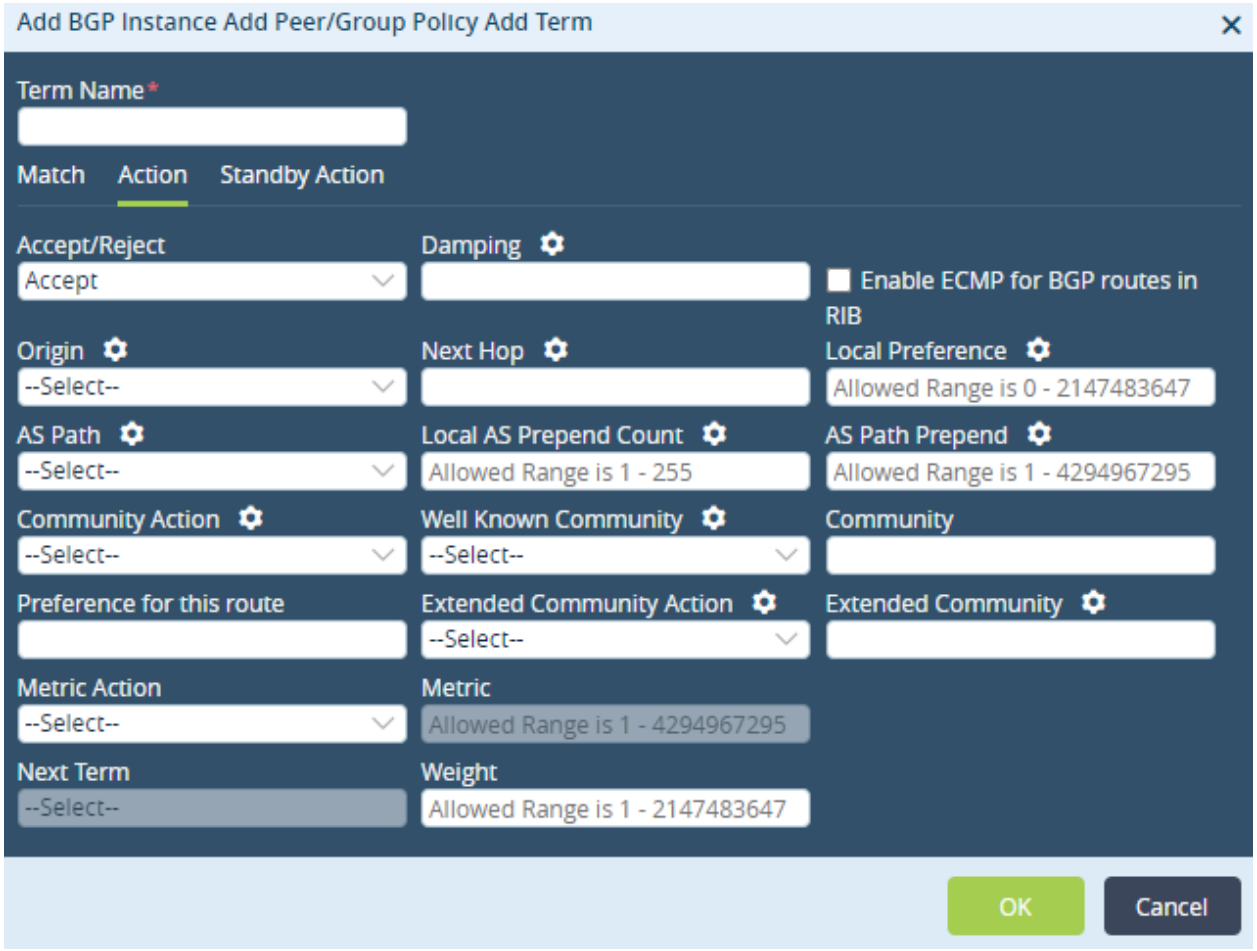

17. Click OK.
18. Click OK to return to the BGP Instance popup window.
19. Select the Peer Group tab to associate the policies with the BGP peer group. In this example, we want to advertise the prefix 192.168.1.0/24 to the peer only if the prefix 192.168.0.0/24 is withdrawn from the RIB. Select the peer group, and then select the Advanced tab. In the Policy group of fields, in the Non-Exist Policy. field select the *Non-Existent* policy, and in the Advertise Policy field select the *Advertise* policy. With this configuration, the BGP peer does not receive the route shown in the Advertise field until the route show in the Non-Exist field is withdrawn.

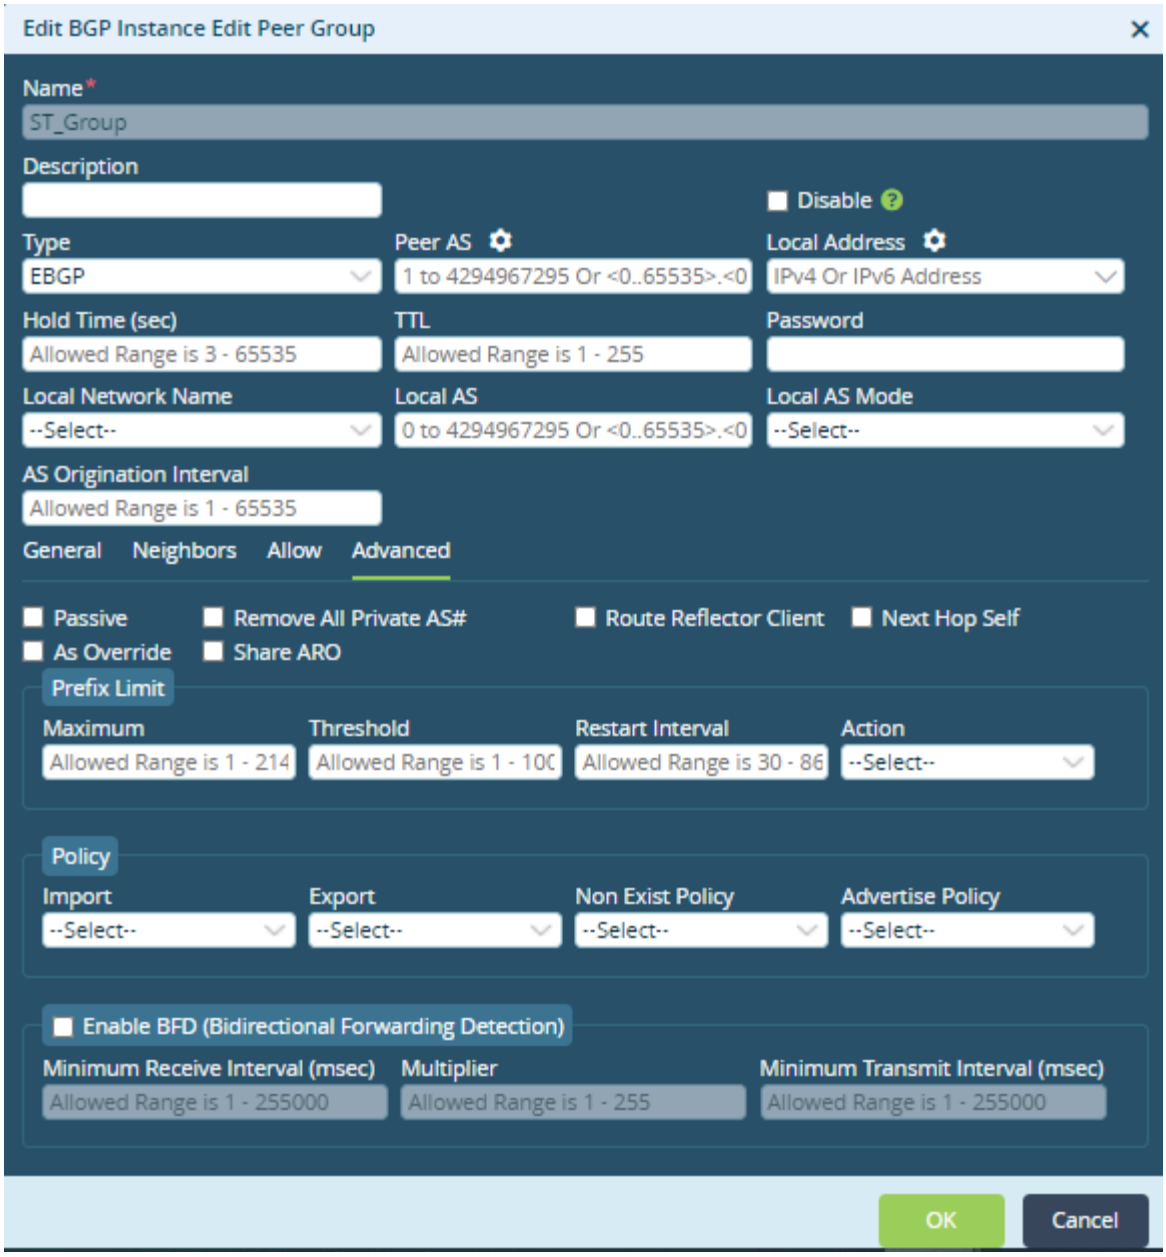

20. Click OK.

Configure PIM

For Releases 20.2 and later.

IP multicast is a method used to distribute application data from a single source host to many receiver hosts over an IP network in a bandwidth efficient manner. IP multicast sends data to hosts that have explicitly asked to receive the data. It is typically used by streaming video, online gaming, financial, and broadcasting applications.

PIM provides several different modes. Of these modes, Versa supports PIM-Sparse Mode (PIM-SM)—including PIM Anycast RP—and PIM-Source Specific Multicast (SSM):

- PIM-Sparse Mode (PIM-SM) sends multicast traffic only after it receives an IGMP join request from a downstream router. It does not assume that there are hosts on the network waiting to receive multicast traffic.
- PIM Anycast RP allows you to configure more than one RP in a multicast network in order to provide redundancy and load balancing.
- PIM-SSM builds trees that are rooted in a single source, offering a more secure and scalable model for a limited number of applications (primarily applications that broadcast content).

Note: By default, the global timers for a PIM instance apply to all PIM-enabled interfaces and networks. To override the global values, configure timers at the interface or network level.

To configure PIM on a virtual router:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking  > Virtual Routers in the left menu bar.
5. Click the  Add icon. The Configure Virtual Router popup window displays.
6. Select the PIM tab in the left menu bar.
7. Click the  Add icon. The Add PIM Instance screen displays.

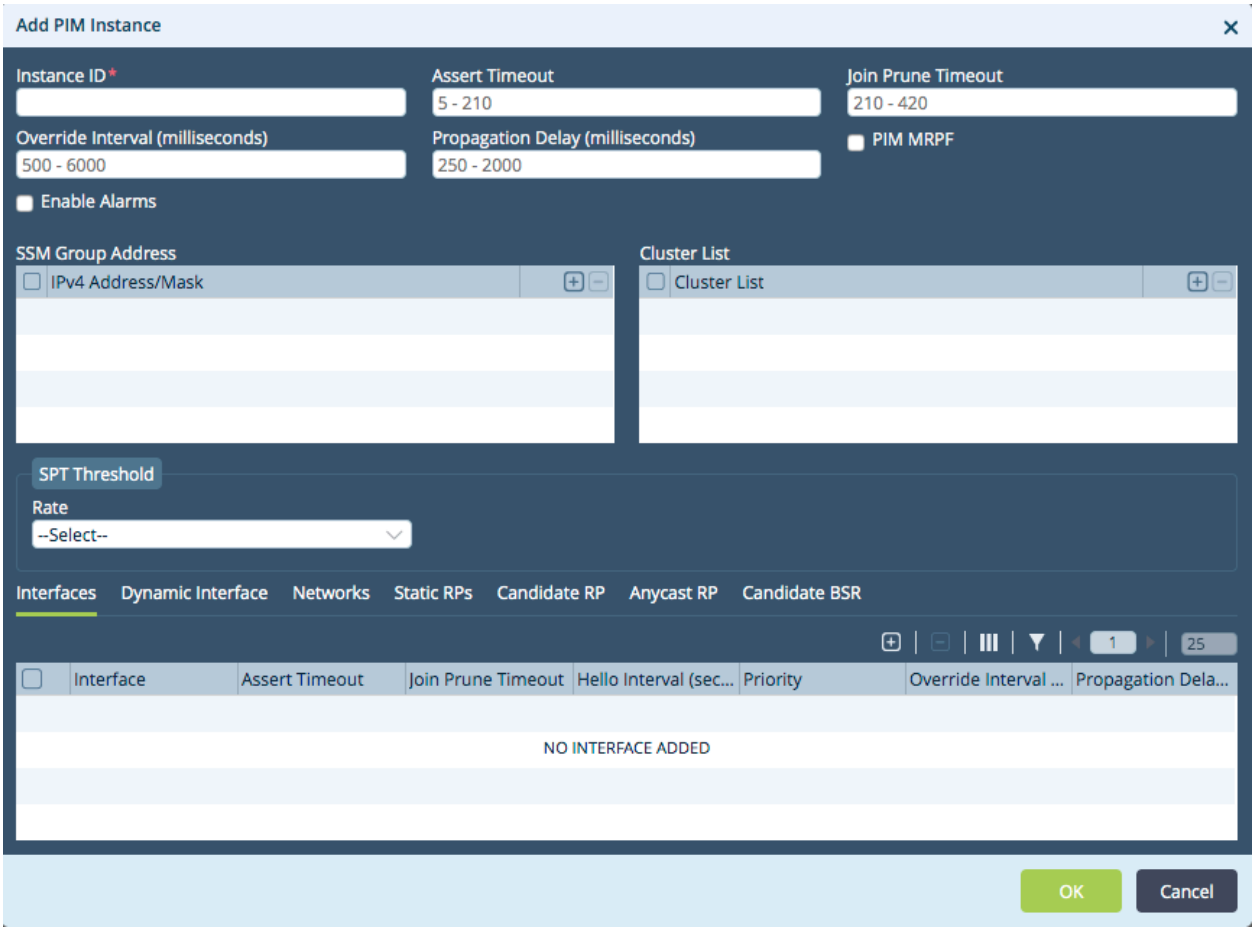

Field	Description
Instance ID	Enter the instance ID for PIM. Note that, for a given tenant, PIM and IGMP have the same instance ID.
Assert Timeout	Enter the time after which the multicast routing devices enter a PIM assert message cycle. *Range:* 5 through 210 seconds
Join Prune Timeout	Enter the time after which the join state is removed if the join state is not refreshed in that time period. The router sends periodic join or prune messages to refresh the state. *Range:* 210 through 420 seconds
Override Interval (milliseconds)	Enter the maximum time, in milliseconds, for which a router or switch waits before sending a join message to override another device's prune message. *Range:* 500 through 6000 milliseconds.
Propagation	Enter how long, in milliseconds, the routing device waits to detect whether a join message

Delay (milliseconds)	is being suppressed by another routing device.
PIM MRPF	Click to perform the RPF check based on the multicast route table. *Default:* RPF check is based on the unicast route table.
Enable Alarms	Enables alarm generation
SSM Group Address	Enter the group address for the source specific multicast (SSM). SSM identifies a set of multicast hosts both by group address and source. The default SSM range (232.0.0.0/8) is preconfigured on VOS devices.
Cluster List	Select the set of branches that form a cluster, which are branches that form PIM relationships among themselves. For multicast over SD-WAN, you must configure a cluster list. For native multicast routing over LANs or WANs, configuring a cluster list is optional.
SPT Threshold	
Rate	Select the threshold value to trigger a switchover from the rendezvous-point tree (RPT) to the shortest-path tree (SPT). You can specify the rate in pps or kbps. To disable switchover, select infinity. *Default:* 1 ppm

8. Select the Interfaces tab, and click Add icon. In the Add Interface popup window, enter information for the following fields.

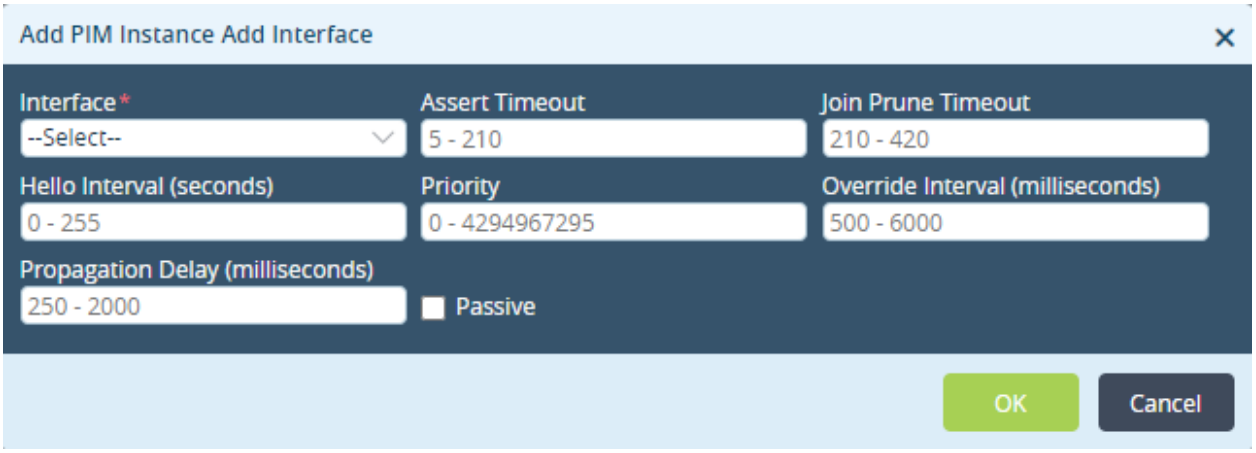

Field	Description
Interface	Select the interface on which to enable PIM.

Assert Timeout	Enter how long a multicast routing devices waits before enter a PIM assert message cycle. *Range:* 5 through 210 seconds
Join Prune Timeout	Enter how long to wait before the join state is removed, if the join state is not refreshed within that time. The router sends periodic join or prune messages to refresh the state. *Range:* 210 through 420 seconds
Hello Interval (seconds)	Enter the time, in seconds, between successive PIM hello messages. *Default:* 35 seconds
Priority	Enter the interface priority for PIM designated router selection.
Override Interval (milliseconds)	Enter the maximum time, in milliseconds, that a router or switch waits before it sends a join message to override another device's prune message. *Range:* 500 through 6000 milliseconds
Propagation Delay (milliseconds)	Enter how long, in milliseconds, the routing device waits to detect whether a join message is currently being suppressed by another routing device.
Passive	Select this check box to indicate that the router is a passive listener.

9. Select the Dynamic Interfaces tab, and enter information for the following fields.

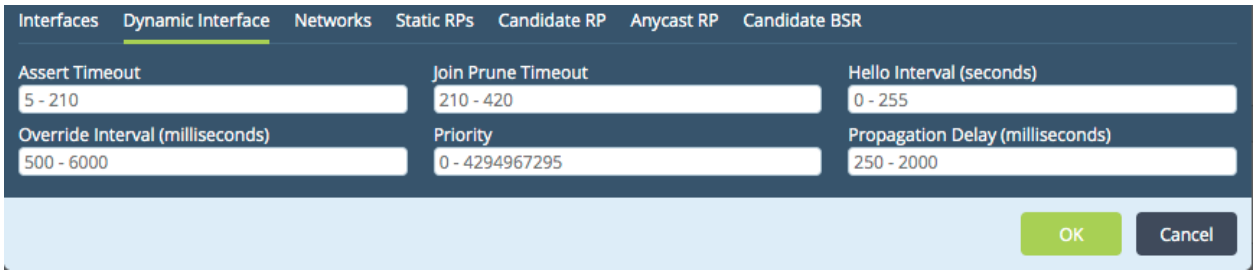

Field	Description
Assert Timeout	Enter how long a multicast routing devices waits before enter a PIM assert message cycle. *Range:* 5 through 210 seconds
Join Prune Timeout	Enter how long to wait before the join state is removed, if the join state is not refreshed within that time. The router sends periodic join or prune messages to refresh the state.

	Range: 210 through 420 seconds
Hello Interval (seconds)	Enter the time, in seconds, between successive PIM hello messages. *Default:* 35 seconds
Override Interval (milliseconds)	Enter the maximum time, in milliseconds, that a router or switch waits before it sends a join message to override another device's prune message. *Range:* 500 through 6000 milliseconds
Priority	Enter the interface priority for PIM designated router selection. *Range:* 0 through 4294967295
Propagation Delay (milliseconds)	Enter how long, in milliseconds, the routing device waits to detect whether a join message is currently being suppressed by another routing device. *Range:* 250 through 2000

10. Select the Networks tab, and click the Add icon. In the Add Networks popup window, enter information for the following fields.

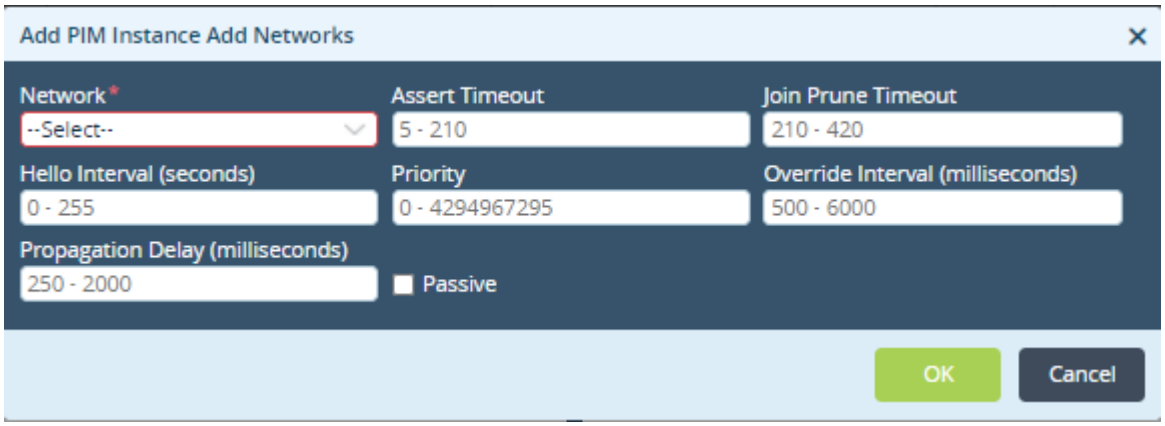

Field	Description
Network	Select the network for PIM.
Assert Timeout	Enter how long a multicast routing devices waits before enter a PIM assert message cycle.

	Range: 5 through 210 seconds
Join Prune Timeout	Enter how long to wait before the join state is removed, if the join state is not refreshed within that time. The router sends periodic join or prune messages to refresh the state. *Range:* 210 through 420 seconds
Hello Interval (seconds)	Enter the time, in seconds, between successive PIM hello messages. *Default:* 35 seconds
Priority	Enter the interface priority for PIM designated router selection. *Range:* 0 through 4294967295
Override Interval (milliseconds)	Enter the maximum time, in milliseconds, that a router or switch waits before it sends a join message to override another device's prune message. Range: 500 through 6000 milliseconds
Propagation Delay (milliseconds)	Enter how long, in milliseconds, the routing device waits to detect whether a join message is currently being suppressed by another routing device. *Range:* 500 through 2000 milliseconds
Passive	Select this check box to indicate that the router is a passive listener.

11. Select the Static RPs tab, and click the Add icon. In the Add Static RPs popup window, enter information for the following fields.

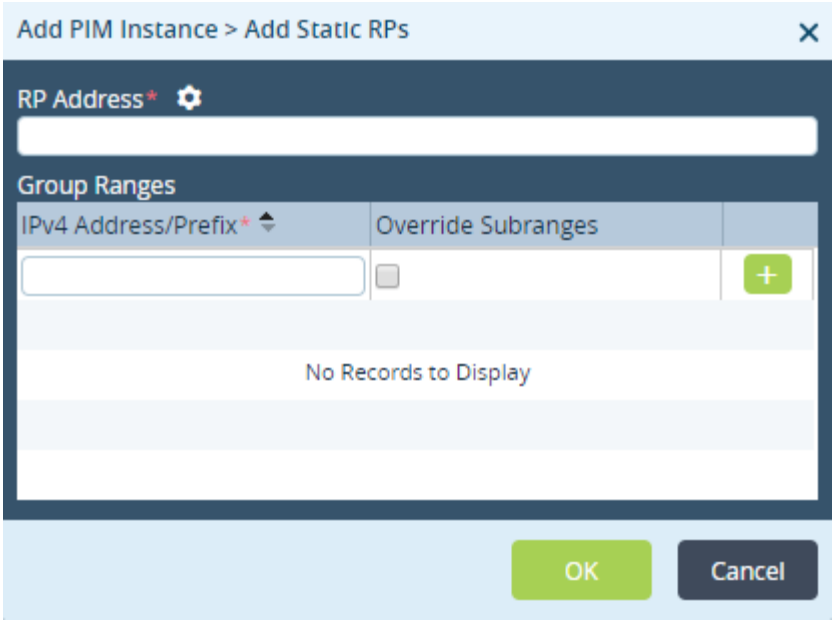

Field	Description
RP Address	Enter the PIM static RP address for a multicast group range.
Group Ranges	Enter the group address range for the rendezvous point (RP).
IPv4 Address/Prefix	Enter the IPv4 multicast group range.
Override Subranges	Select to override the subranges of the IPv4 addresses.

12. Select the Candidate RP tab, and click the Add icon. In the Add Candidate RP popup window, enter information for the following fields.

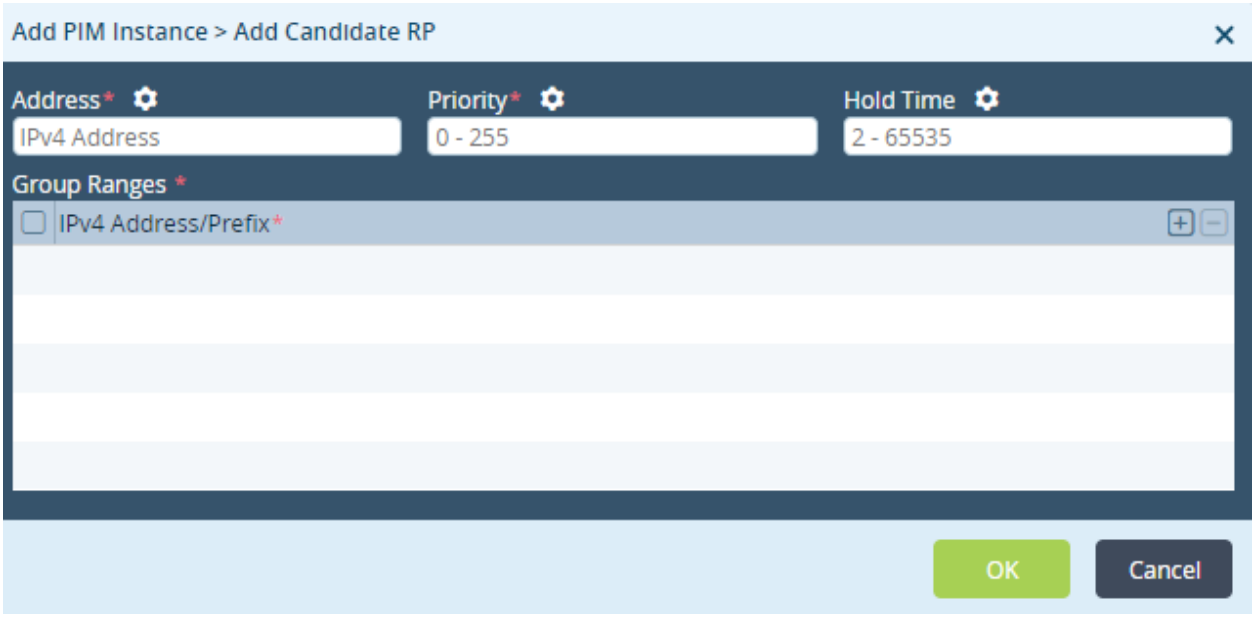

Field	Description
Address	Enter the IPv4 address for candidate RP.
Priority	Enter the candidate RP priority. *Range:* 0 through 255
Hold Time	Enter how long a candidate RP advertisement is valid. *Range: 3* through 65535
Group Ranges	Enter the group address range for the RP.
IPv4 Address/Prefix	Enter the IPv4 multicast group range.

13. Select Anycast RP tab, and click the Add icon. In the Add Anycast RP popup window, enter information for the following fields.

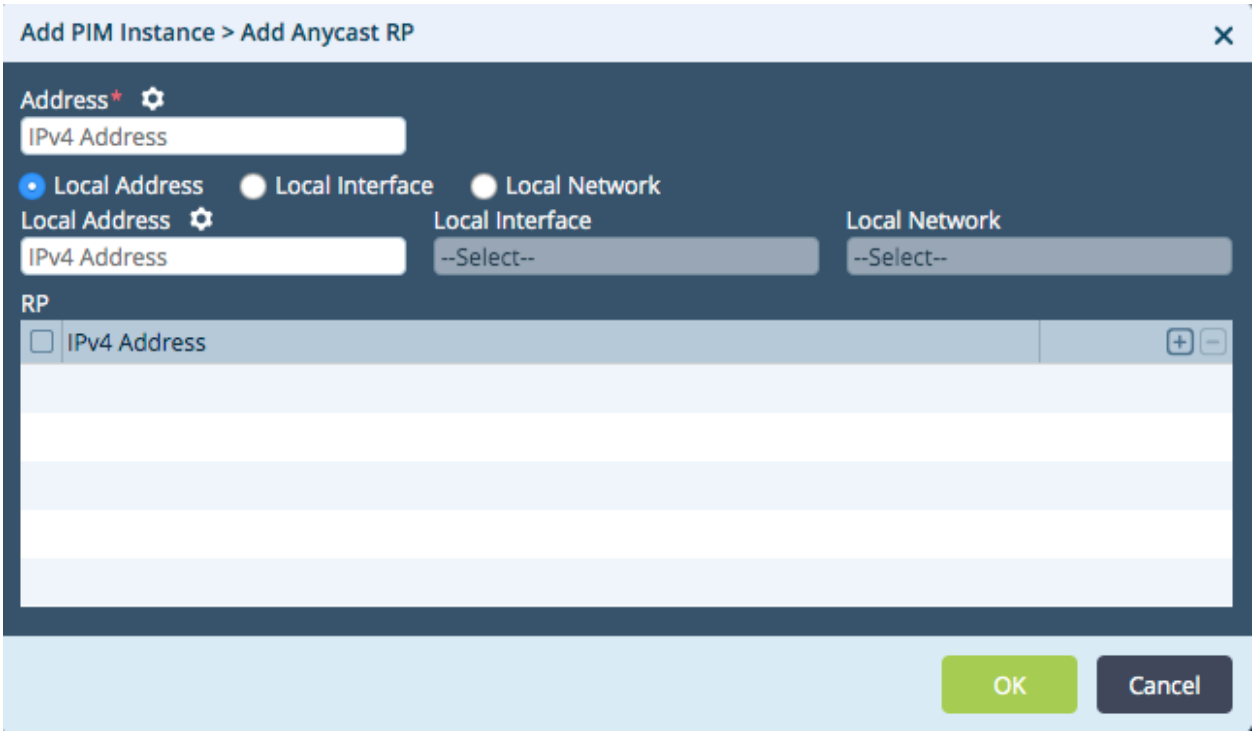

Field	Description
Address	Enter the IPv4 address for Anycast RP.
Local Address	Click the button and enter the local IPv4 address
Local Interface	Click the button and enter the local interface name.
Local Network	Click the button and enter the network name.
RP	Click the Add icon, then enter the IPv4 address of the other RP in the anycast RP set. For a set of multiple RPs, repeat this step to add the IPv4 address for each of the RPs in the set.

14. Select the Candidate BSR tab, and enter information for the following fields.

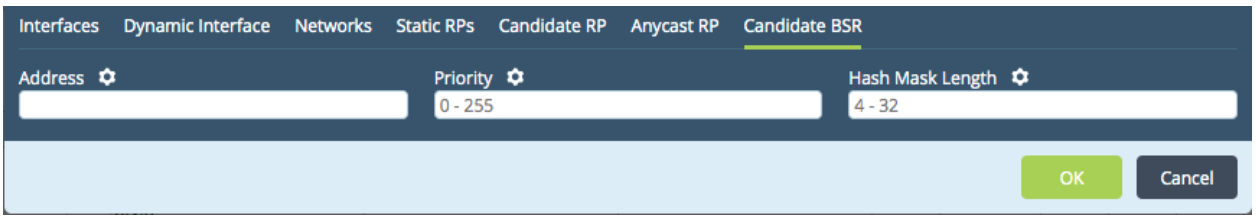

Field	Description
Address	Enter the address for the candidate bootstrap router (BSR).
Priority	Enter the candidate BSR priority. Higher priority is preferred for candidate BSR. *Range:* 0 through 255
Hash Mask Length	Enter the length, in bits, of the mask to use in the hash function. *Range: 4* through 32

15. Click OK. The Configure Virtual Router popup window displays the configured PIM instance for a tenant.

Configure IGMP

For Release 20.2 and later.

Internet group management protocol (IGMP) is a multicast protocol that hosts and adjacent routers on IP networks use to establish and manage multicast group memberships. You enable IGMP between a host device and a local multicast router. A host requests membership to a multicast group by sending join requests to any local multicast router.

VOS devices support IGMP Version 1, Version 2, and Version 3. IGMPv3 with the Include Filter (only) is used for PIM-SSM. VOS devices also support static IGMP groups.

By default, IGMPv2 is enabled on all the PIM-enabled interfaces or networks. To enable IGMPv3 on an interface or network, or to override default properties such as timers, immediate leave, and IGMP static groups, you must create an IGMP instance.

To configure IGMP:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking  > Virtual Routers in the left menu bar.

5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the IGMP tab in the left menu bar.
7. Click the Add icon. In the Add IGMP Instance popup window, enter information for the following fields.

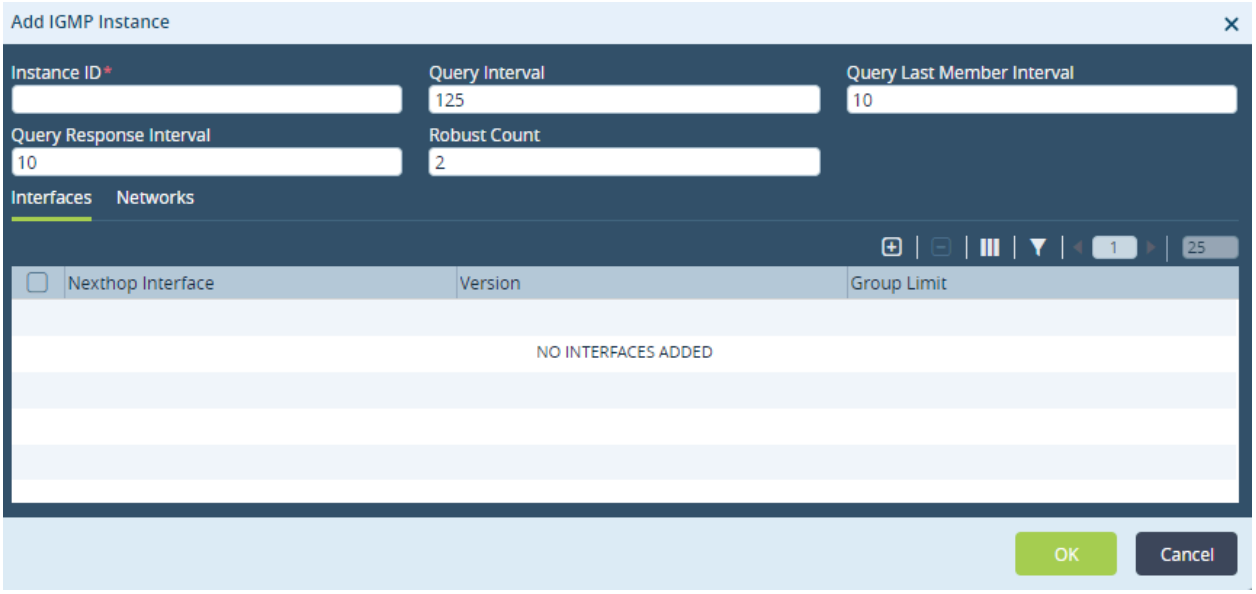

Field	Description
Instance ID	Enter the instance ID assigned to IGMP. Note that, for a given tenant, PIM and IGMP have the same instance ID.
Query Interval	Enter the interval after which the IGMP router periodically sends general host-query messages on each attached network to inquire about the membership information. *Default:* 125 seconds
Query Last Member Interval	Enter the maximum amount of time between group-specific query messages. *Default:* 10 seconds
Query Response Interval	Enter the maximum amount of time between the router sending a host-query message and the router receiving a response from a host. *Default:* 10 seconds
Robust Count	Enter the number of IGMP group query messages sent by the router when the same router receives an IGMP leave message on a shared network.

5. Select the Interfaces tab, and click the Add icon. In the Add Interfaces popup window, enter information for the

following fields.

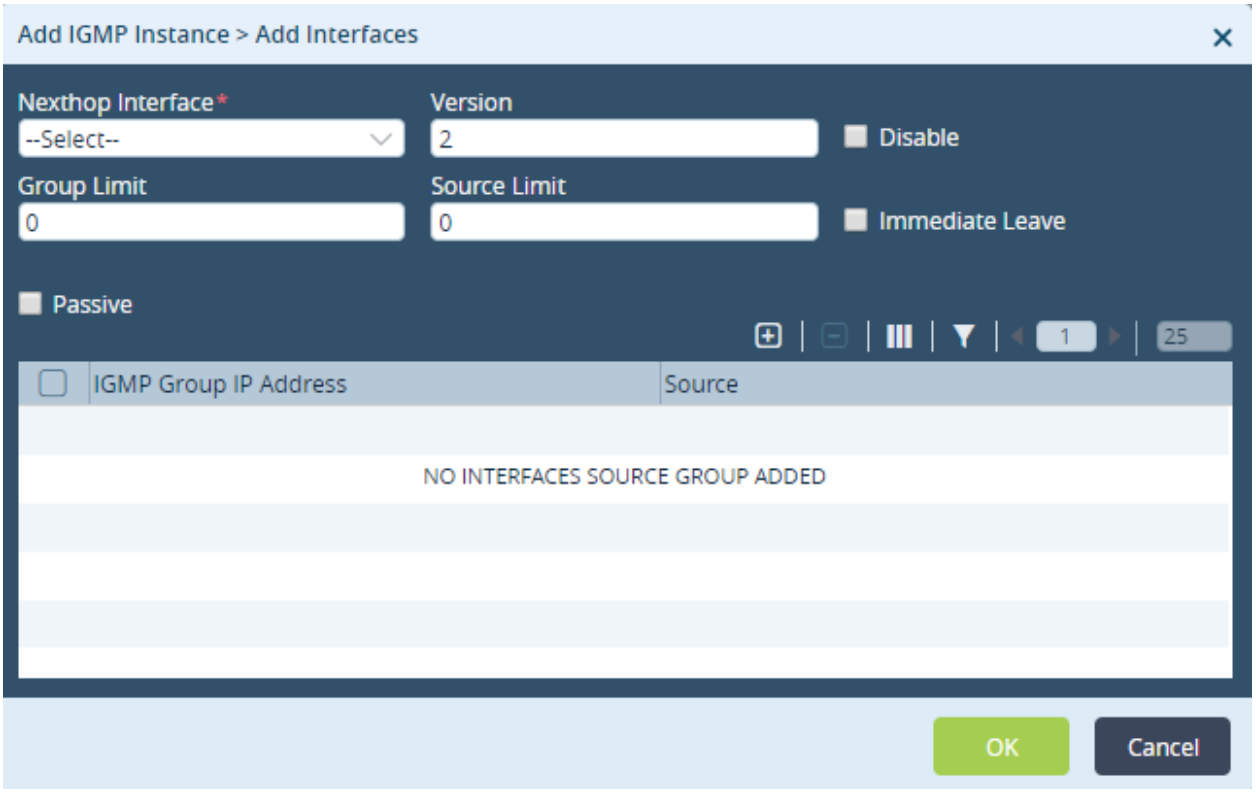

Field	Description
Next-Hop Interface	Select an interface to enable IGMP and to use for the next-hop interface.
Version	Enter the IGMP version. *Default:* 2
Disable	Click to disable IGMP on the interface.
Group Limit	Enter the maximum number of multicast group joins for logical interfaces.
Source Limit	Enter the maximum number of multicast source joins for logical interfaces.
Immediate Leave	Click to flush the IGMP group cache entries immediately after receiving an IGMP leave message on an interface.
Passive	Click to mark the router as a passive listener that does not advertise itself.
IGMP Group IP Address	Enter the IGMP static group IP address.

Source	Enter the source IP address corresponding to a static IGMPv3 group.

6. Select the Networks tab, and click the Add icon. In the Add Networks popup window, enter information for the following fields.

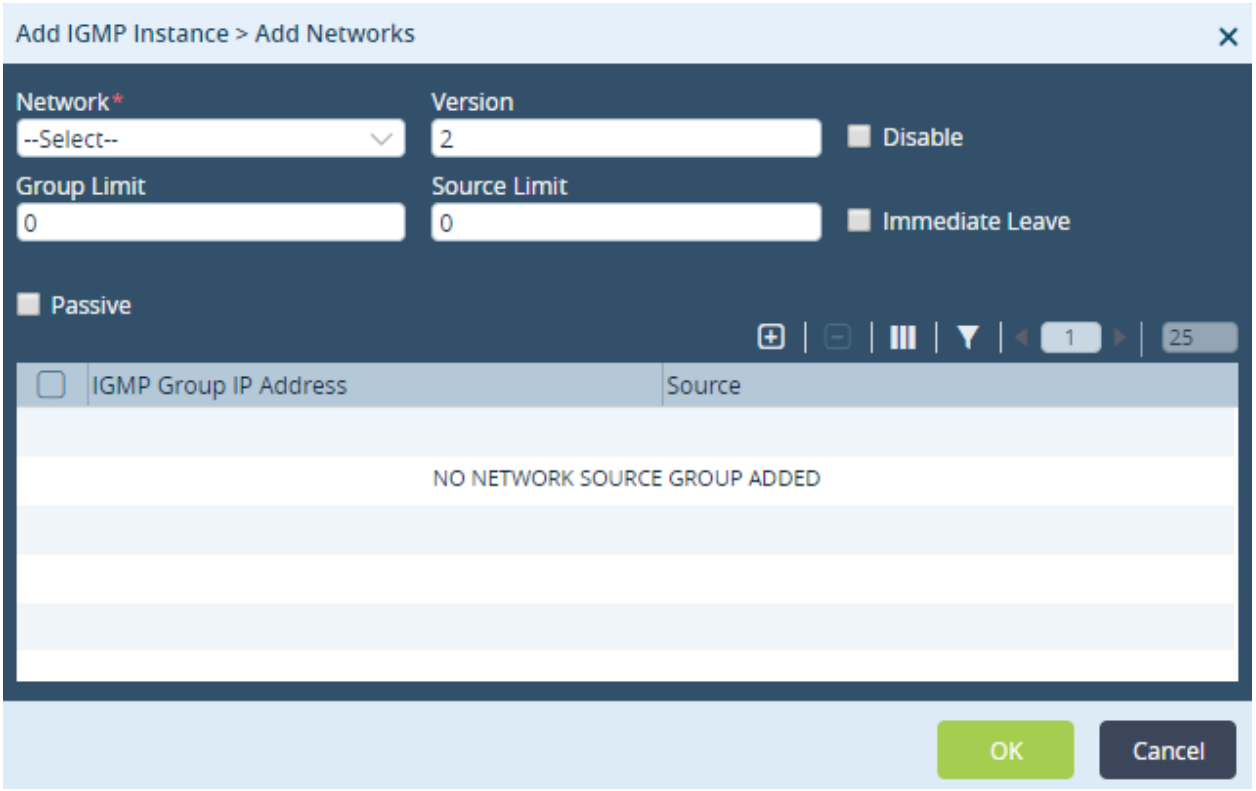

Field	Description
Network	Select the network to enable IGMP.
Version	Enter the IGMP version. *Default:* 2
Disable	Click to disable IGMP on the network.
Group Limit	Enter the maximum number of multicast group joins for logical interfaces.
Source Limit	Enter the maximum number of multicast source joins for logical interfaces.
Immediate Leave	Click to flush IGMP group cache entries immediately after receiving an IGMP leave message on a given network..

Passive	Click to mark the router as a passive listener that does not advertise itself.
IGMP Group IP Address	Enter the static group IP address for IGMP.
Source	Enter the source IP address corresponding to a static IGMPv3 group.

7. Click OK. The Configure Virtual Router popup window displays the configured IGMP instances.

Configure Router Advertisements

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the Router Advertisement tab in the left menu bar.
7. Click the Add icon. In the Add Router Advertisement popup window, enter information for the following fields.

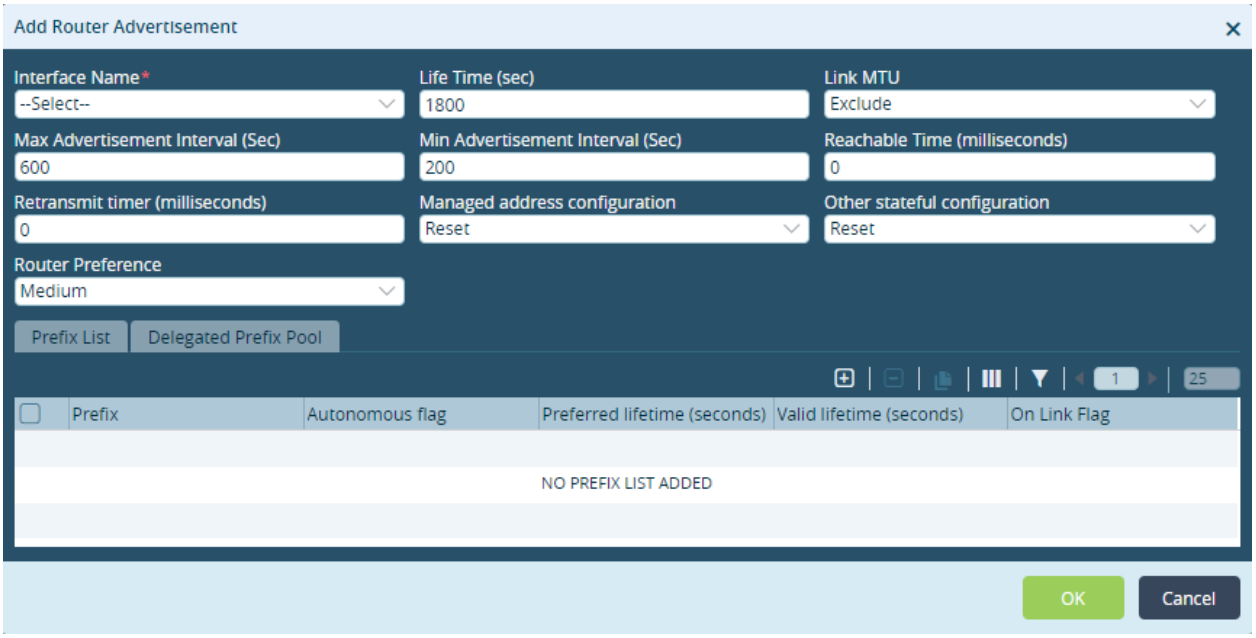

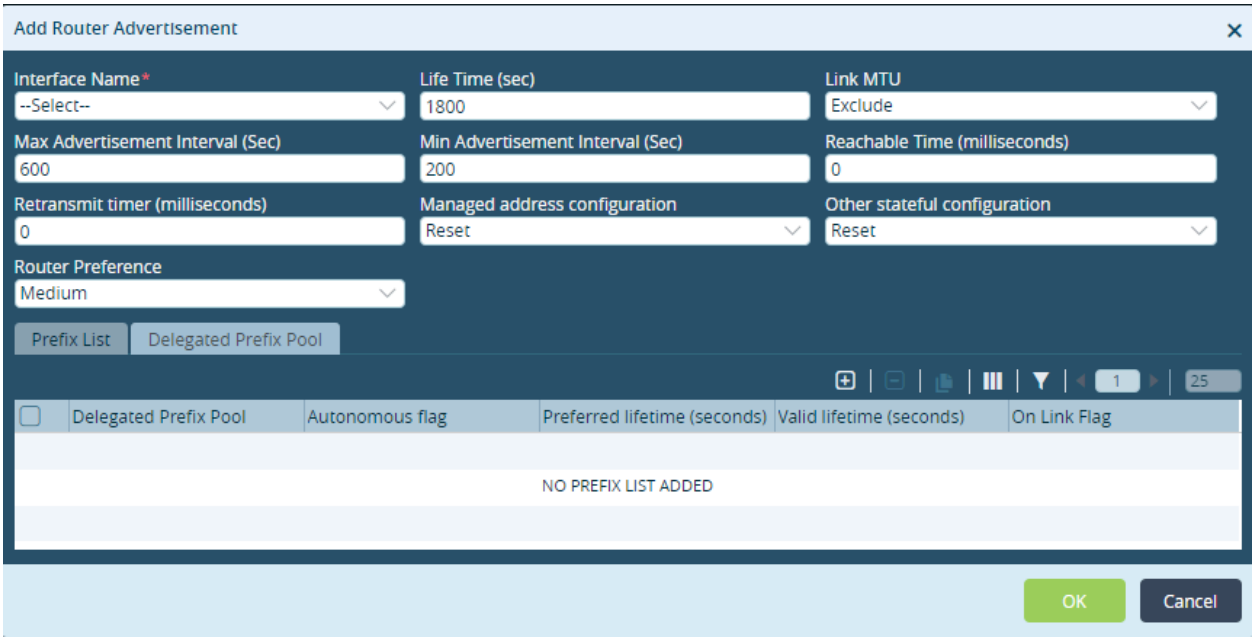

Field	Description
Interface Name	Select the interface to use.
Life Time (sec)	Enter the default router lifetime
Link MTU	Include/Exclude link MTU in RA
Max Advertisement Interval (Sec)	Enter the maximum interval between each router advertisement message, in seconds.
Min Advertisement Interval (Sec)	Enter the minimum interval between each router advertisement message, in seconds.
Reachable Time (milliseconds)	Enter how long the host or router considers a neighbor as reachable until another reachability confirmation is received from that neighbor, in milliseconds.
Retransmit timer (milliseconds)	Enter how often to retransmit neighbor solicitation messages, in milliseconds.
Managed Address Configuration	Select whether to have the host to use a stateful autoconfiguration protocol for address autoconfiguration, in addition to any already configured stateless autoconfiguration.
Other Stateful Configuration	Select whether to enable autoconfiguration of other non-address-related information. ◦ Reset ◦ Set

Field	Description
Router Preference	Select the advertise router preference in router advertisement (RA). ◦ High ◦ Low ◦ Medium When an IPv6 host receives a router advertisement message, it can use the router preference setting to select a default router.

8. Select the Prefix List tab, and click the Add icon. In the Add Prefix List popup window, enter information for the following fields.

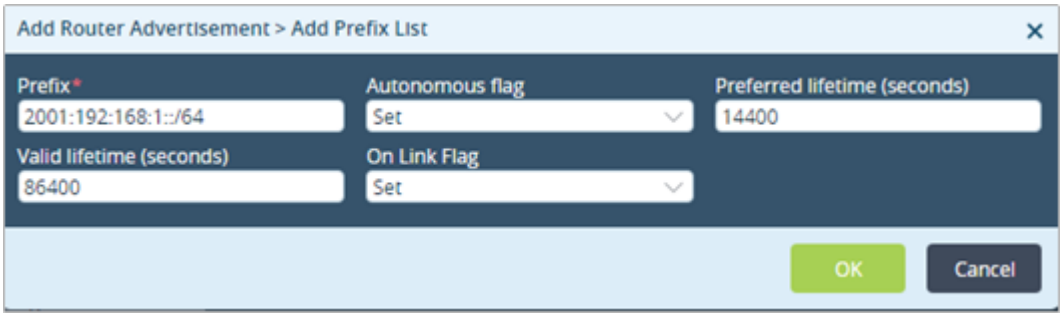

Field	Description
Prefix	Enter the IP prefix.
Autonomous Flag	Select whether prefixes in the router advertisement messages are used for stateless address autoconfiguration: ◦ Reset ◦ Set
Preferred Lifetime (seconds)	Enter how long to prefer the autoconfigured prefix, in seconds.
Valid Lifetime (seconds)	Enter how long the prefix remains valid, in seconds.
On Link Flag	Select whether the prefix advertised in a router-advertisement message is an on-link prefix:

Field	Description
	◦ Reset ◦ Set

9. Click OK.
10. Select the Delegated Prefix Pool tab.

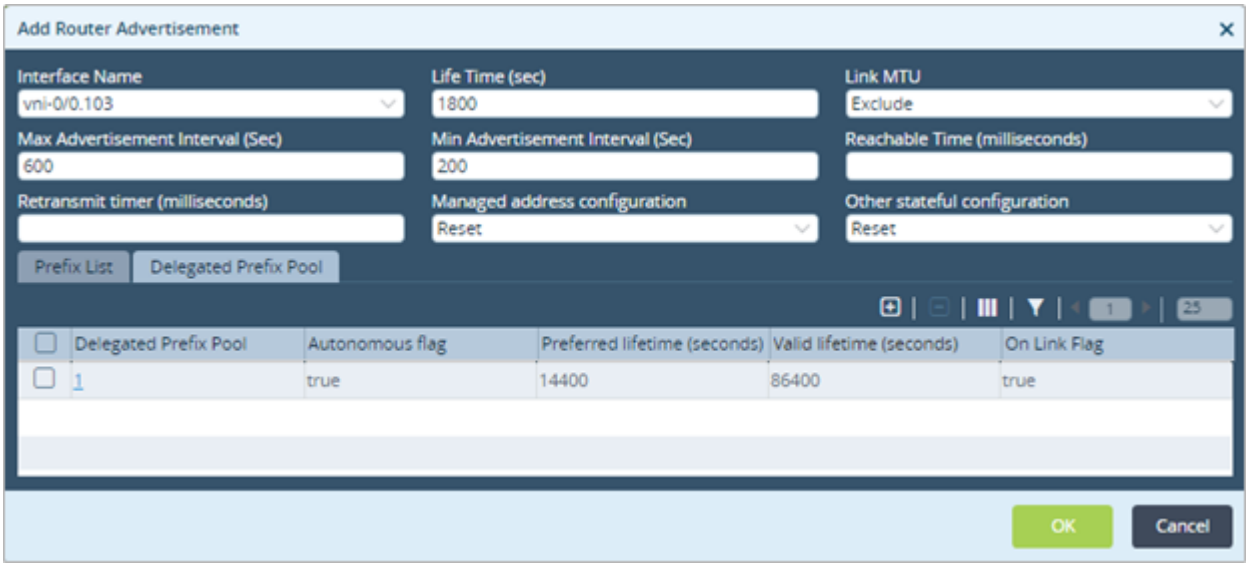

11. Click the ⊞ Add icon. In the Add Prefix List popup window, enter information for the following fields.

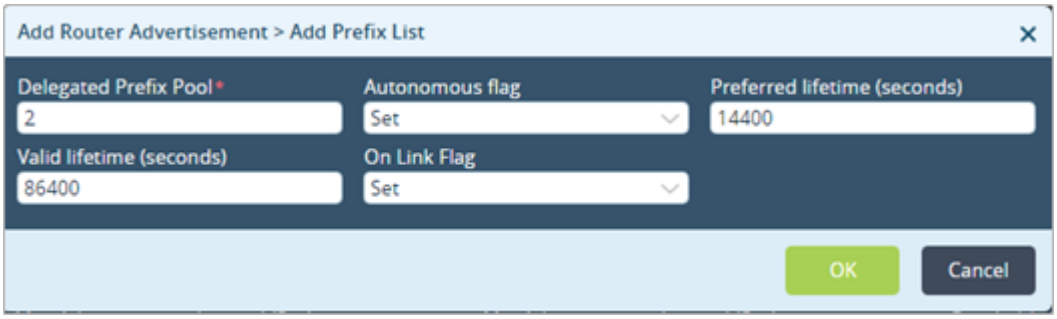

Field	Description
Delegated Prefix Pool	Enter the number of delegated prefixes to include in router advertisement messages.

Field	Description
Autonomous Flag	Select whether prefixes in the router advertisement messages are used for stateless address autoconfiguration: ◦ Reset ◦ Set
Preferred Lifetime (seconds)	Enter how long to prefer the autoconfigured prefix, in seconds.
Valid Lifetime (seconds)	Enter how long the prefix remains valid, in seconds.
On Link Flag	Select whether the prefix advertised in a router-advertisement message is an on-link prefix: ◦ Set ◦ Reset

10. Click OK. The Add Router Advertisement popup window displays the configured router advertisements.

Configure Prefix Lists

To configure prefix list:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click Add icon. The Configure Virtual Router popup window displays.
6. Select Prefix Lists tab in the left menu bar.

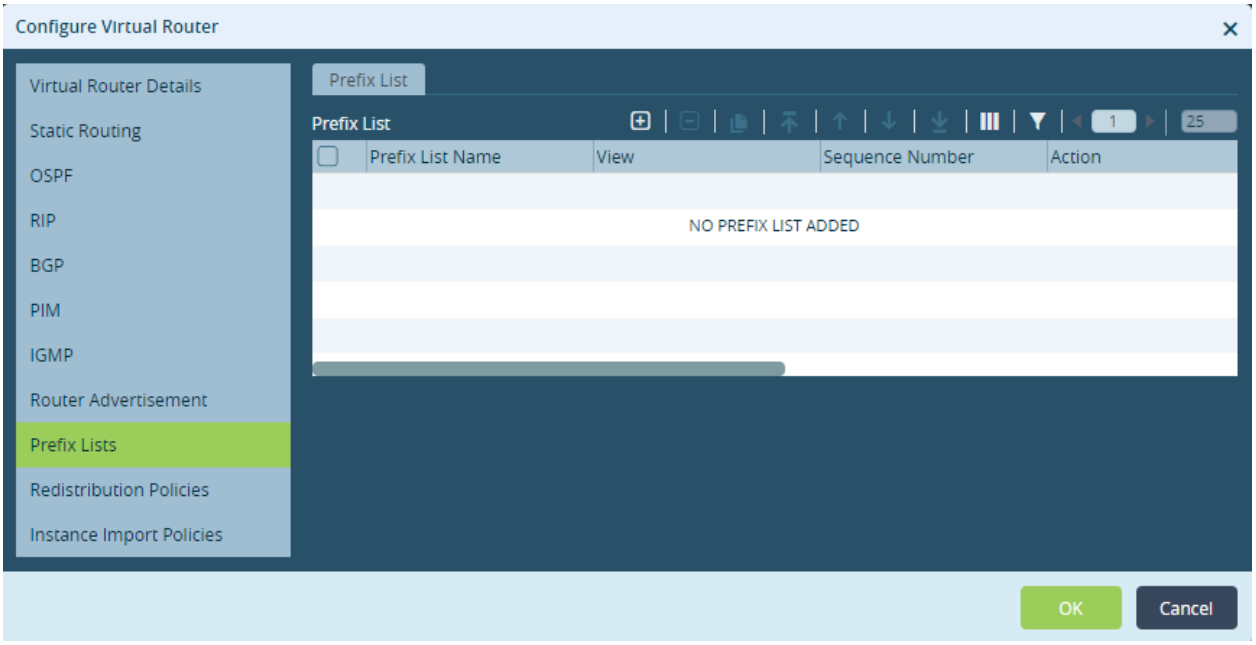

7. Click ⊞ Add icon. The Add Prefix List window displays.

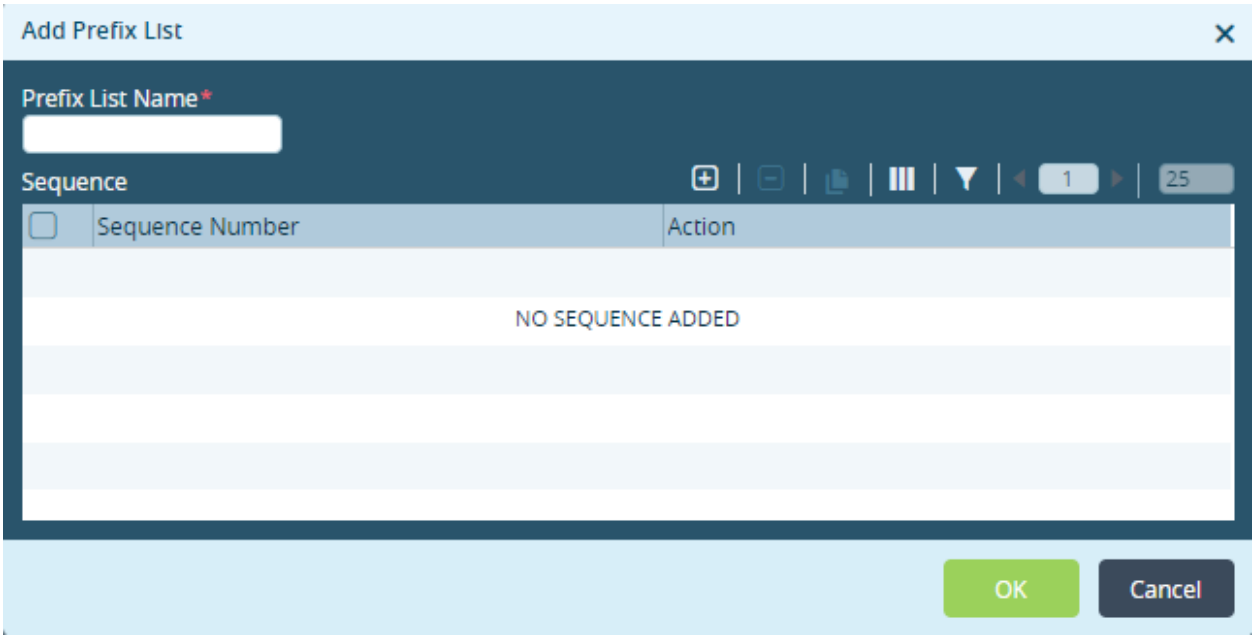

8. In the Prefix List Name field, enter a name for the prefix list.
9. In the Sequence field, click ⊞ Add icon to add a sequence. Enter information for the following fields.

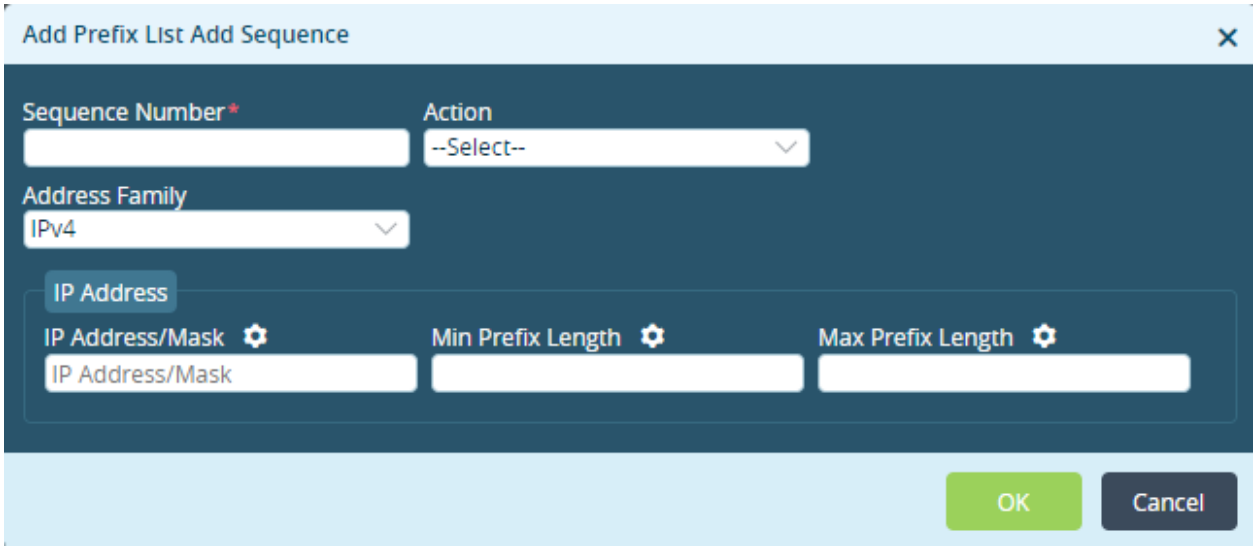

Field	Description
Sequence Number	Enter a sequence number for the prefix list.
Action	Select the action to take on the routes: ◦ Deny—Select to deny routes on this prefix list. ◦ Permit—Select to allow routes on this prefix list.
Address Family	Select the broadcast address family protocol of the route: ◦ IPv4 ◦ IPv6
IP Address (Group of Fields)	
◦ IP Address/Mask	Enter the IP prefix of the routes grouped in this prefix list.
◦ Minimum Prefix Length	Enter the minimum number of prefix length to match. *Range:* 0 through 32
◦ Maximum Prefix Length	Enter the maximum number of prefix length to match. *Range:* 0 through 32

10. Click OK. The Add Prefix List popup window displays the configured sequence.

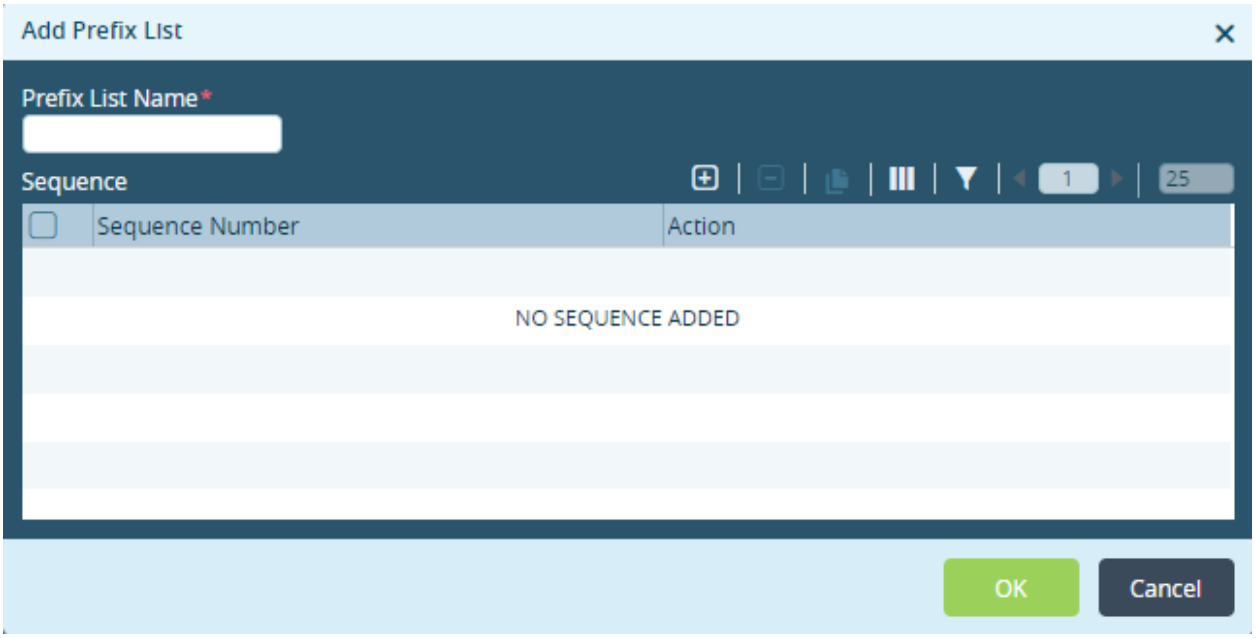

11. Click OK. The Configure Virtual Router popup window displays the configured prefix lists.

Configure Redistribution Policies

You configure redistribution policies to forward routes from one routing protocol to another protocol. You can redistribute routes among static, OSPF, and BGP For example, to send static routes to an OSPF route, you need a redistribution policy.

To configure redistribution policies:

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the Router Advertisement tab in the left menu bar.

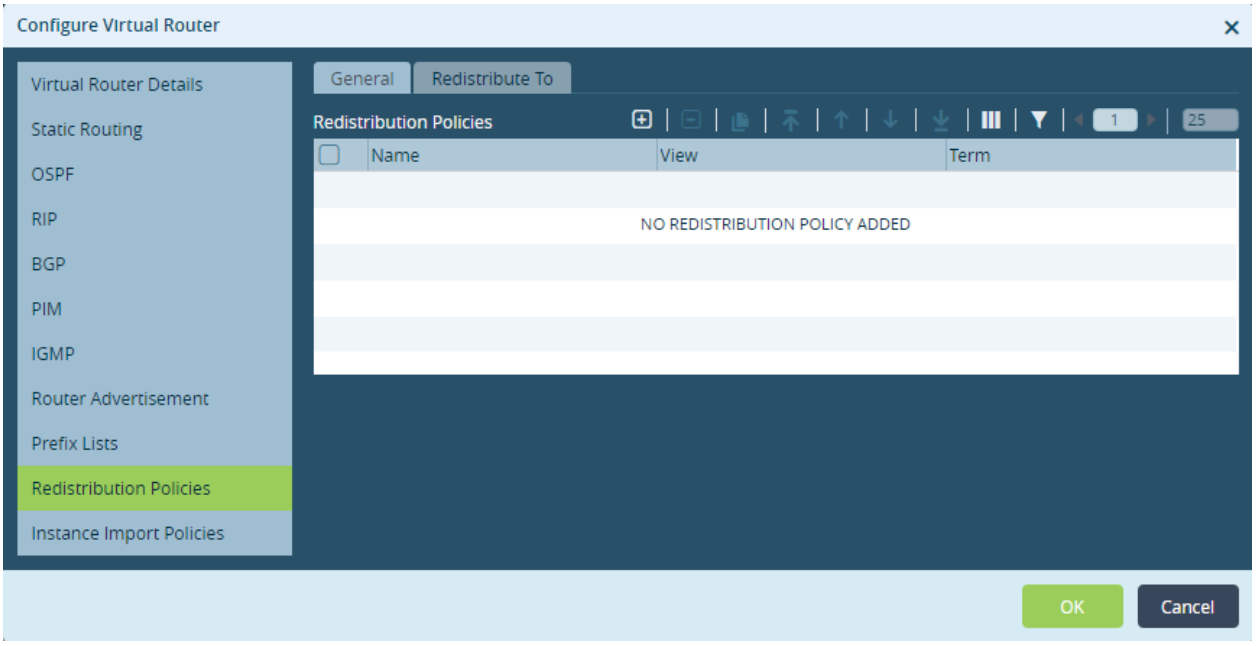

7. Click the ⊞ Add icon.
8. In the Add Redistribution Policy popup window, click the ⊞ Add icon to add the name of the policy.

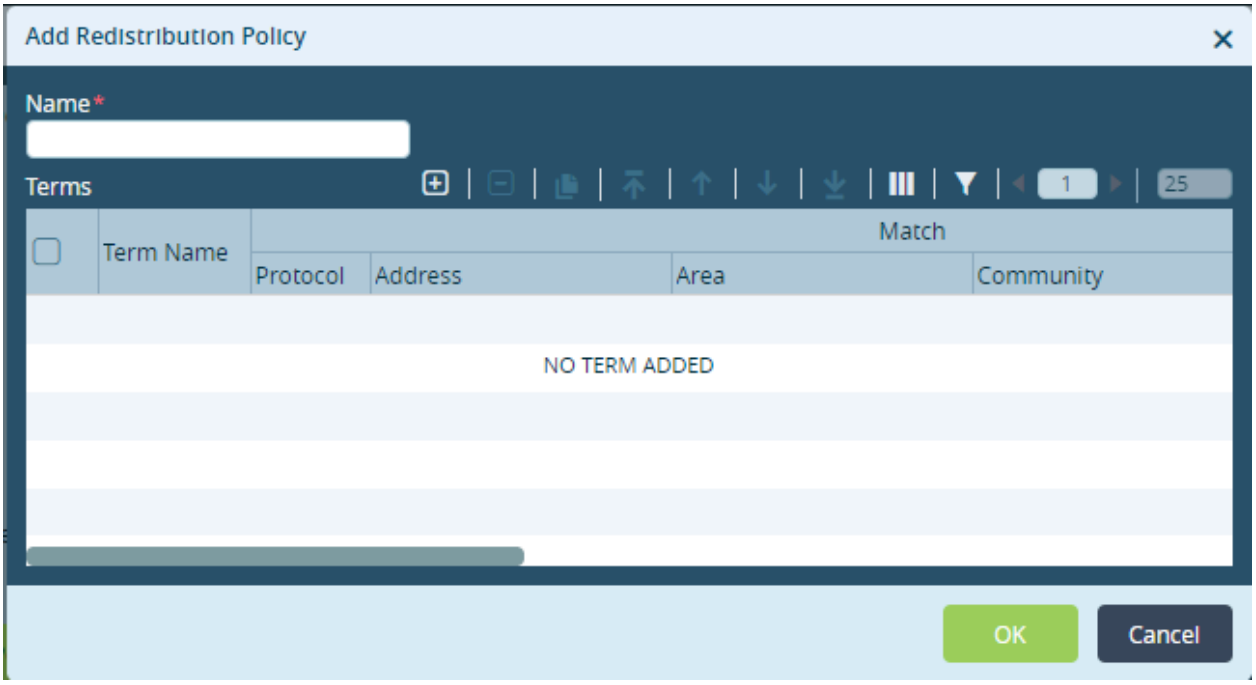

9. Click the ⊞ Add icon to add a policy term. Enter information for the following fields.

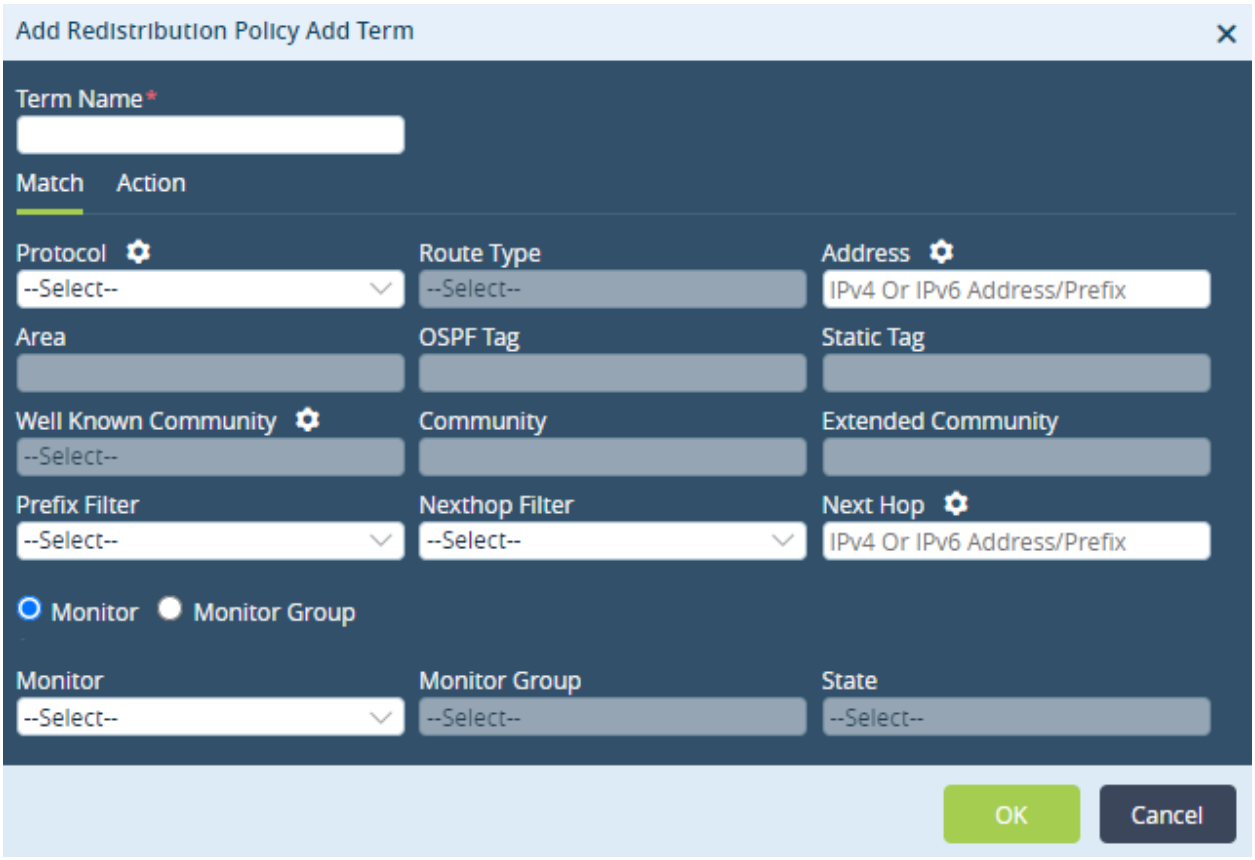

Field	Description
Term Name	Enter a name for the term in the redistribution policy. The first instance created is evaluated first by the policy rule, and the remaining terms are evaluated in the order they are listed in the term name table.

10. Select the Match tab to define redistribution policy match conditions. Enter information for the following fields.

Field	Description
Protocol	Select the protocol to match for redistribution: ◦ BGP ◦ DHCP ◦ Direct ◦ OSPF ◦ RIP ◦ SD-WAN ◦ Static

Field	Description
Route Type	Select the route type for the protocol.
Address	Enter the IPv4 or IPv6 address of the route to match.
Area	Enter the OSPF area to match.
OSPF Tag	Enter the OSPF tag to match.
Static Tag	Enter the OSPF static tag to match.
Well Known Community	(For Releases 21.2.1 and later.) Select a well-known community: ◦ no_advertise—A BGP speaker that receives a route containing this community value must not advertise the route to any external or internal peer. ◦ no_export—A BGP speaker that receives a route containing this community value must not advertise the route to its external BGP peers. However, the BGP speaker can advertise the route to its IBGP peers and to confederation peers in other member ASs within its local confederation. ◦ no_export_subconfed—A BGP speaker that receives a route containing this community value must not advertise the route to any external peer including peers in other member ASs within its confederation. If you select a well-known community, you cannot also configure a string in the Community field.
Community	Enter the BGP community string to match. A BGP community is a group of destinations with a common property. This path attribute in BGP update messages identifies community members and performs actions at a group level instead of to an individual level. BGP communities help identify and segregate BGP routes, enabling a smooth traffic flow. If you configure a community string, you cannot also select a community in the Well-Known Community field.
Extended Community	Enter the extended BGP community identifier.
Prefix Filter	Name of the prefix list that defines the terms for the route prefixes to be advertised.
Nexthop Filter	Name of the prefix list that defines the terms for matching the next hop of the route.
Next Hop	Enter the next-hop address for the route.

Field	Description
Monitor	Name of the monitor used for liveness detection, the state of which should be matched.
Monitor Group	Name of the monitor-group used for liveness detection, the state of which should be matched.
State	State of the monitor or monitor group to match.

11. Select the Action tab to define redistribution policy action conditions. Enter information for the following fields.

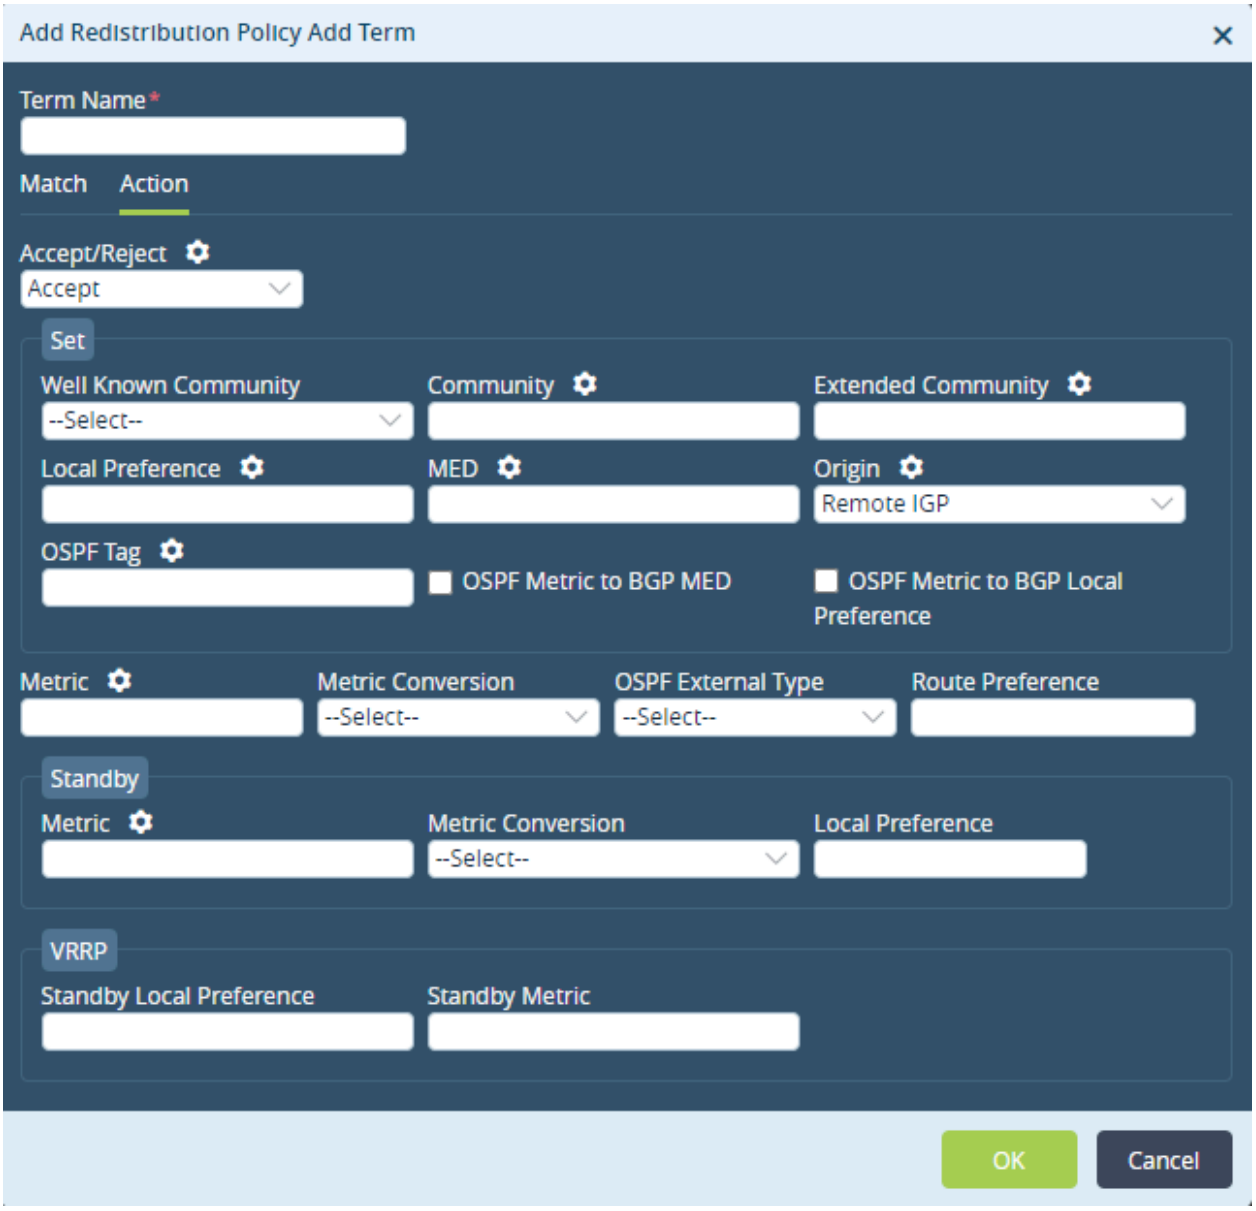

Field	Description
Accept/Reject	◦ Select the action to take for the route: ◦ Accept—Accept all the traffic for the route. ◦ Reject—Rejects all the traffic for the route.
Set (Group of Fields)	
Well-Known Community	(For Releases 21.2.1 and later.) Select a well-known community: ◦ no_advertise—A BGP speaker that receives a route containing this community value must not advertise the route to any external or internal peer. ◦ no_export—A BGP speaker that receives a route containing this community value must not advertise the route to its external BGP peers. However, the BGP speaker can advertise the route to its IBGP peers and to confederation peers in other member ASs within its local confederation. ◦ no_export_subconfed—A BGP speaker that receives a route containing this community value must not advertise the route to any external peer including peers in other member ASs within its confederation. If you select a well-known community, you cannot also configure a string in the Community field.
Community	Enter the BGP community string to match. A BGP community is a group of destinations with a common property. This path attribute in BGP update messages identifies community members and performs actions at a group level instead of to an individual level. BGP communities help identify and segregate BGP routes, enabling a smooth traffic flow. If you configure a community string, you cannot also select a community in the Well-Known Community field.
Extended Community	Enter the BGP extended community identifier to add to the route.
Local Preference	Enter the local BGP preference to add to the route.
MED	Enter the multiexit BGP discriminator to add to the route.
Origin	For BGP, select the source of the BGP route: ◦ Local EGP

Field	Description
	◦ Remote IGP ◦ Unknown Heritage
OSPF Tag	For OSPF, enter the OSPF tag to add to the route.
OSPF Metric to BGP MED	For routes being redistributed from OSPF to BGP, click to set the MED of BGP route to the same value as the OSPF metric.
OSPF Metric to BGP Local Preference	For routes being redistributed from OSPF to BGP, click to set the local preference of BGP route to a mapped value from OSPF metric (294967295 minus the OSPF metric).
Metric	Enter a value to determine how one route should be chosen over another: ◦ Bandwidth ◦ Communication cost ◦ Hop count ◦ Load ◦ MTU ◦ Path cost ◦ Path length
Metric Conversion	Select the conversion factor for the metric value: ◦ Inverse ◦ Scale Down ◦ Scale Up ◦ Set ◦ Truncate
OSPF External Type	Select the OSPF external type to use when distributing a route to OSPF: ◦ E1 ◦ E2
Route Preference	Enter a value for the route preference.
Standby (Group of Fields)	Configure interchassis HA standby metrics.

Field	Description
◦ Metric	Enter the metric value for the interchassis HA standby.
◦ Metric Conversion	Select how to convert the for metric value from a route during redistribution: ◦ Inverse ◦ Scale Down ◦ Scale Up ◦ Set ◦ Truncate
◦ Local Preference	Enter the local preference to use during the route redistribution.
VRRP (Group of Fields)	
◦ Standby Preference	Enter the standby preference to use when exporting the prefix corresponding to the subnet of the interface in VRRP backup state.
◦ Standby Metric	Enter the metric value to set while the device is in VRRP backup state.

12. Click OK. The Add Redistribution Policy popup window displays the configured policies.

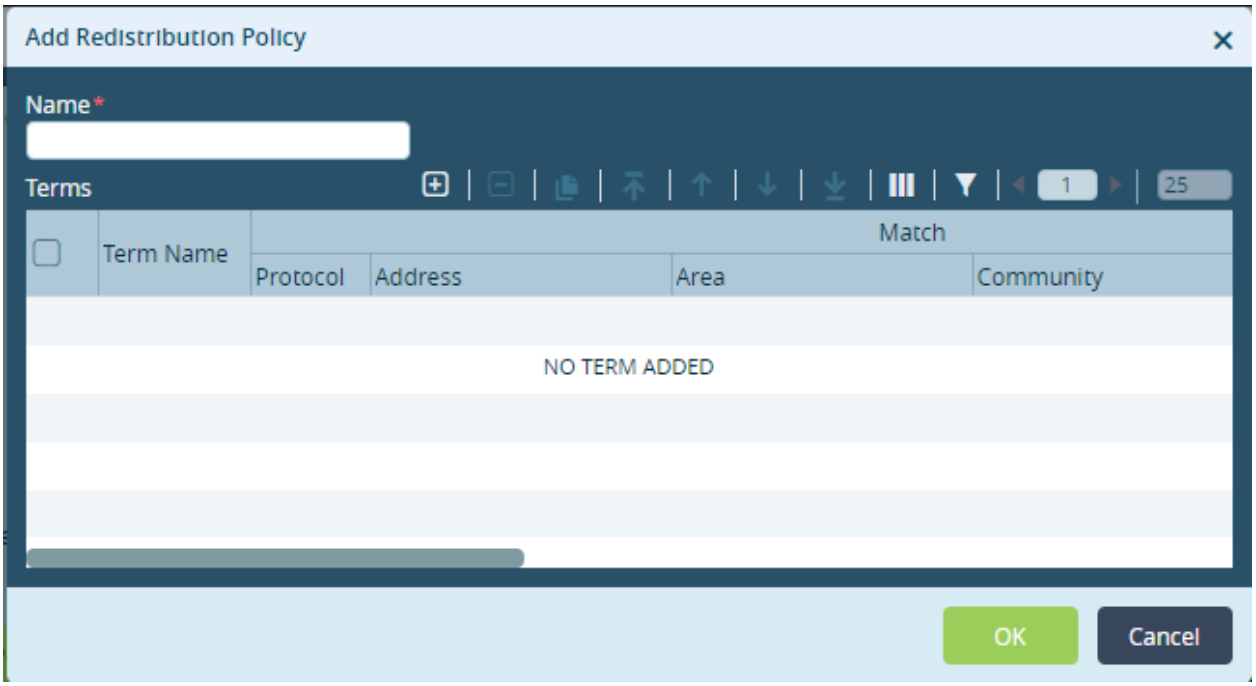

13. Select the Redistribute To tab, and click the Add icon. Enter information for the following fields.

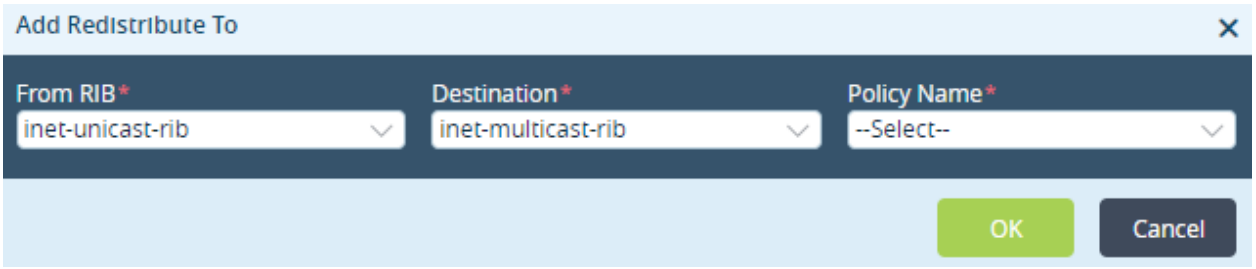

Field	Description
From RIB	Select the instance from which the redistribution occurs. ◦ inet-unicast-rib ◦ inet6-unicast-rib ◦ inet-multicast-rib ◦ inet6-multicast-rib
Destination	Select the destination for the redistribution. ◦ inet-multicast-rib ◦ inet6-multicast-rib

Field	Description
	◦ ospf ◦ rip ◦ bgp
Policy Name	Select the name of the redistribution policy to use.

14. Click OK. The Configure Virtual Router popup window displays the configured redistribution policies.

Configure Instance Import Policies

1. If you are continuing from the previous section, skip to Step 6.
2. In Director view:
   a. Select the Configuration tab in the top menu bar.
   b. Select Templates > Device Templates in the horizontal menu bar.
   c. Select an organization in the left menu bar.
   d. Select a post-staging template in the main pane. The view changes to Appliance view.
3. Select the Configuration tab in the top menu bar.
4. Select Networking > Virtual Routers in the left menu bar.
5. Click the Add icon. The Configure Virtual Router popup window displays.
6. Select the Instance Import Policies tab in the left menu bar.

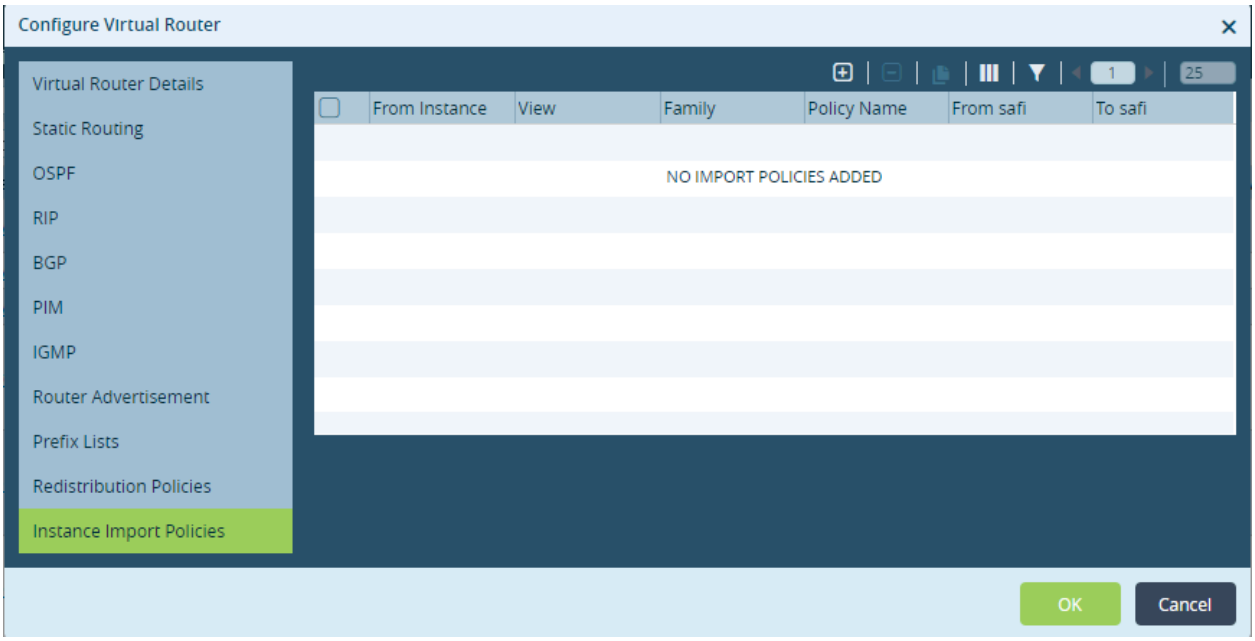

7. Click the Add icon and enter information for the following fields.

Field	Description
From Instance	Select the instance to use for the import policies.
Family	Select the address family of the routes.
Policy Name	Select the name of the redistribution policy.
From Safi	Enter the subsequent address family identifier from which to import the policy. ◦ multicast ◦ unicast
To Safi	Enter the subsequent address family identifier to which to export the policy. ◦ multicast ◦ unicast

8. Click OK. The Configure Virtual Router popup window displays the configured import policies.

Default Routing Preferences

The following tables lists the default values for route preferences, also referred to as administrative distances. The route with the lowest preference is the most likely to become the active route.

Route Source	Default Preference
Connected	0
Static	1
EBGP	20
OSPF internal	30
OSPF external	110
RIP	120
IBGP	200

Software Release Information

Releases 16.1R2 and later support all content described in this article, except:

- Releases 16.1R2S11 and later support configuring BGP conditional advertisements.
- Releases 20.2 and later add support for multicast protocols, specifically, IGMP versions 1, 2, and 3, PIM sparse mode, SSM, PIM static RP, and PIM bootstrap router.
- Releases 21.2.1 and later add support for the following new BGP configuration fields: AS path ignore, AS path multipath relax, community 4 byte, relax first AS check, soft reconfiguration, well-known community, weight, and suppress peer AS.

Additional Information

Configure Interchassis HA
Configure IP Multicast
Configure IP SLA Monitor Objects
Troubleshoot Routing Protocols