



Border Gateway Protocol

- [Feature history for BGP, on page 1](#)
- [BGP for Cisco SD-WAN overlay networks, on page 2](#)
- [Configure BGP, on page 3](#)
- [Verify BGP redistribute route in OMP, on page 31](#)
- [Redistribute BGP routes and AS path information, on page 32](#)

Feature history for BGP

This table describes the developments of this feature, by release.

Table 1: Feature History

Feature Name	Release Information	Description
MPLS-BGP Support on the Service Side	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature allows you to enable support on Multiprotocol Label Switching (MPLS). Multiple Service VPNs use inter autonomous system (AS) BGP labelled path to forward the traffic, which in turn helps scaling the service side VPNs with less control plane signaling. Label distribution for a given VPN routing and forwarding (VRF) instance on a given device can be handled by Border Gateway Protocol (BGP).
BGP Community Propagation	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature enables propagation of BGP communities between routing protocols during route redistribution. On one node, the OMP redistributes routes from BGP and on the other node, the BGP redistributes routes from OMP. In addition to configurable AS path attribute propagation, there is an option to propagate BGP communities. The BGP community propagation helps in propagating BGP communities between Cisco Catalyst SD-WAN sites, across VPNs using OMP redistribution. To propagate the BGP communities during route redistribution from OMP, use the propagate-community command.

Feature Name	Release Information	Description
Ability to Match and Set Communities during BGP to OMP Redistribution	Cisco IOS XE Catalyst SD-WAN Release 17.4.1a Cisco vManage Release 20.4.1	This feature enhances the implementation of match and set clauses for redistribution of routes from BGP to OMP and vice versa on Cisco IOS XE Catalyst SD-WAN devices. You can redistribute the routes from a BGP into an OMP routing to allow route traffic to help increase the accessibility within the network. The <code>route-maps</code> are defined locally on each device to filter the routes from the source routing protocol. You can manipulate OMP communities to propagate BGP routes. The following commands are updated: <pre>route-map advertise bgp route-map bgp-to-omp redistribute omp route-map omp-to-bgp</pre>

BGP for Cisco SD-WAN overlay networks

BGP is the routing protocol that directs traffic across the internet by exchanging routing information between different networks, known as autonomous systems (AS). It determines the best paths for data packets to travel between these large networks to ensure efficient and reliable delivery.

Cisco Catalyst SD-WAN overlay networks support BGP unicast routing protocols. These protocols can be configured on Cisco IOS XE Catalyst SD-WAN devices within any Virtual Routing and Forwarding (VRF) instance, excluding transport and management VRFs. This configuration enables reachability to local site networks. Cisco IOS XE Catalyst SD-WAN devices can also redistribute route information learned from BGP into the Overlay Management Protocol (OMP), allowing OMP to make more informed path selections within the overlay network.

BGP topologies

- **Direct Connection to Layer 3 VPN MPLS WAN Cloud:** When a local site connects directly to a Layer 3 VPN (L3VPN) MPLS WAN cloud, the Cisco IOS XE Catalyst SD-WAN devices function as MPLS Customer Edge (CE) devices. They establish a BGP peering session with the Provider Edge (PE) router in the L3VPN MPLS cloud.
- **Indirect Connection to WAN Cloud:** If devices at a local site are one or more hops away from the WAN cloud and connect indirectly through a non-Cisco IOS XE Catalyst SD-WAN device, standard routing must be enabled on the devices' DTLS connections to reach the WAN. In such scenarios, either OSPF or BGP can serve as the routing protocol.

In both of these topologies, BGP sessions operate over a Datagram Transport Layer Security (DTLS) connection. This DTLS connection is established on the loopback interface within VRF 0, which is the dedicated transport VRF for carrying control traffic in the overlay network. The Cisco SD-WAN Validator learns about this DTLS connection via the loopback interface and relays this information to the Cisco SD-WAN Controller for tracking TLOC-related data. Although VRF 0 also hosts the physical interface connecting the Cisco IOS XE Catalyst SD-WAN device to its neighbor (e.g., PE router in MPLS or hub/next-hop router), a DTLS tunnel connection is not established on this physical interface.

BGP Community Propagation

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.3.1a, the BGP community propagation feature is supported. Previously, BGP communities were not sent to BGP neighbors even if attached. With this feature, Cisco IOS XE Catalyst SD-WAN devices can propagate communities attached to BGP entries to their neighbors. This is particularly useful when migrating a BGP overlay to a Cisco Catalyst SD-WAN overlay, as it ensures BGP route attributes are propagated between Cisco Catalyst SD-WAN sites across VPNs. Further details can be found using the `propagate-community` command.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.4.1a, administrators gain the ability to manipulate communities during propagation between BGP and OMP, and vice versa, using the `route-map` command. This command defines conditions for redistributing routes between routing protocols. Each `route-map` command includes `match` commands, which specify the conditions (e.g., matching communities) under which redistribution is permitted, and `set` commands, which define specific redistribution actions to be performed if the `match` criteria are met. For more information on the commands, refer [Command Reference Guide](#).

Configure BGP

The BGP can be used for service-side routing to provide reachability to networks at the local site, and it can be used for transport-side routing to enable communication between Cisco IOS XE Catalyst SD-WAN devices when a device is not directly connected to the WAN cloud. Create separate BGP templates for the two BGP routing types.



Note Cisco IOS XE Catalyst SD-WAN devices use VRFs in place of VPNs. However, the following steps still apply to configure Cisco IOS XE Catalyst SD-WAN devices through Cisco SD-WAN Manager. When you complete the configuration, the system automatically maps the VPN configurations to VRF configurations.

To configure the BGP routing protocol using Cisco SD-WAN Manager templates:

Procedure

- Step 1** Create a BGP feature template to configure BGP parameters.
- Step 2** Create a VPN feature template to configure VPN parameters for either service-side BGP routing (in any VPN other than VPN 0 or VPN 512) or transport-side BGP routing (in VPN 0).

Create a BGP template

Procedure

- Step 1** From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
- Step 2** Click **Device Templates**.
- Step 3** Click **Create Template**.

- Step 4** From the **Create Template** drop-down list, choose **From Feature Template**.
- Step 5** From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
- Step 6** To create a template for **VPN 0** or **VPN 512**:
- Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
 - Under **Additional VPN 0 Templates**, click **BGP**.
 - From the **BGP** drop-down list, click **Create Template**. The BGP template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining BGP parameters.
- Step 7** To create a template for VPNs **1** through **511**, and **513** through **65530**:
- Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
 - Click the **Service VPN** drop-down list.
 - Under **Additional VPN Templates**, click **BGP**.
 - From the **BGP** drop-down list, click **Create Template**. The BGP template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining BGP parameters.
- Step 8** In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
- Step 9** In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Configure basic BGP parameters

Procedure

- Step 1** To configure Border Gateway Protocol (BGP), click **Basic Configuration** and configure the following parameters. Parameters marked with an asterisk are required to configure BGP.

Parameter Name	Description
Shutdown*	Click No to enable BGP for the VPN.
AS number*	Enter the local AS number.
Router ID	Enter the BGP router ID in decimal four-part dotted notation.
Propagate AS Path	Click On to carry BGP AS path information into OMP.
Internal Routes Distance	Enter a value to apply as the BGP route administrative distance for routes coming from one AS into another. Range: 0 through 255 Default: 200

Parameter Name	Description
Local Routes Distance	Specify the BGP route administrative distance for routes within the local AS. By default, a route received locally from BGP is preferred over a route received from OMP. Range: 0 through 255 Default: 200
External Routes Distance	Specify the BGP route administrative distance for routes learned from other sites in the overlay network. Range: 0 through 255 Default: 20

Step 2 Click **Save** to save the feature template.

For service-side BGP, configure Overlay Management Protocol (OMP) to advertise to the Cisco SD-WAN Controller any BGP routes that the device learns. By default, Cisco IOS XE Catalyst SD-WAN devices advertise to OMP both the connected routes on the device and the static routes that are configured on the device, but it does not advertise BGP external routes learned by the device. You configure this route advertisement in the OMP template for devices.

For transport-side BGP, you must also configure a physical interface and a loopback interface in VPN 0. In addition, you should create a policy for BGP to advertise the loopback interface address to its neighbors, and apply the policy in the BGP instance or to a specific neighbor.

Configure BGP address family

Procedure

Step 1 To configure global BGP address family information, click **Unicast Address Family** and configure the following parameters:

Parameter	Option	Sub-Option	Description
IPv4 / IPv6			Click IPv4 to configure an IPv4 Unicast Address Family. Click IPv6 to configure an IPv6 Unicast Address Family.
Maximum Paths			Specify the maximum number of parallel IBGP paths that can be installed into a route table to enable IBGP multipath load sharing. Range: 0 to 32
Mark as Optional Row			Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.

Parameter	Option	Sub-Option	Description	
Redistribute	Click Redistribute > New Redistribute .			
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.		
	Protocol	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are:		
		static	Redistribute static routes into BGP.	
		connected	Redistribute connected routes into BGP.	
		ospf	Redistribute Open Shortest Path First routes into BGP.	
		omp	Redistribute Overlay Management Protocol routes into BGP.	
		nat	Redistribute Network Address Translation routes into BGP.	
		natpool-outside	Redistribute outside NAT routes into BGP.	
	At a minimum, choose the following:			
<ul style="list-style-type: none"> • For service-side BGP routing, choose OMP. By default, OMP routes are not redistributed into BGP. • For transport-side BGP routing, choose Connected, and then under Route Policy, specify a route policy that has BGP advertise the loopback interface address to its neighbors. 				
Route Policy	Enter the name of the route policy to apply to redistributed routes.			
Click Add to save the redistribution information.				
Network	Click Network > New Network .			
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.		
	Network Prefix	Enter a network prefix, in the format <i>prefix/length</i> to be advertised by BGP.		
	Click Add to save the network prefix.			

Parameter	Option	Sub-Option	Description
Aggregate Address	Click Aggregate Address > New Aggregate Address .		
	Mark as Optional Row		Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
	Aggregate Prefix IPv6 Aggregate Prefix		Enter the prefix of the addresses to aggregate for all BGP sessions in the format <i>prefix/length</i> .
	AS Set Path		Click On to generate the set path information for aggregated prefixes.
	Summary Only		Click On to filter out specific routes from the BGP updates.
	Click Add to save the aggregate address.		

Step 2 Click **Save** to save the feature template.

Step 3 To change the AS number, remove the BGP configuration and wait for few seconds.

Step 4 Configure the BGP again with changed global-as and the local-as configuration.

Configure BGP neighbors

Before you begin

For BGP to function, you must configure at least one neighbor.

Procedure

Step 1 To configure a neighbor, click **Neighbor > New Neighbor**, and configure the following parameters:

Parameter Name	Options	Sub-Options	Description
IPv4 / IPv6	Click IPv4 to configure IPv4 neighbors. Click IPv6 to configure IPv6 neighbors.		
Address/IPv6 Address	Specify the IP address of the BGP neighbor.		
Description	Enter a description of the BGP neighbor.		
Remote AS	Enter the AS number of the remote BGP peer.		

Parameter Name	Options	Sub-Options	Description	
Address Family	Click On and select the address family. Enter the address family information. The software supports only the BGP IPv4 unicast address family.			
	Address Family	Select the address family. The software supports only the BGP IPv4 unicast address family.		
	Maximum Number of Prefixes	Specify the maximum number of prefixes that can be received from the neighbor. Range: 1 through 4294967295 Default: 0		
		Threshold	Specify the threshold at which to generate a warning message or restart the BGP connection. The threshold is a percentage of the maximum number of prefixes. You can specify either a restart interval or a warning only.	
		Restart Interval	Specify the duration to wait for restarting the BGP connection. <i>Range:</i> 1 through 65535 minutes	
		Warning Only	Click On to display a warning message without restarting the BGP connection.	
		Route Policy In	Click On and specify the name of a route policy that will have the prefixes from the neighbour.	
Route Policy Out	Click On and specify the name of a route policy that will have the prefixes sent to the neighbour.			
Shutdown	Click On to enable the connection to the BGP neighbor.			

Step 2 Click **Save**.

Configure MPLS interface

The Cisco IOS XE Catalyst SD-WAN devices support Multiprotocol Label Switching (MPLS) to enable multiple protocol environment. MPLS offers an extremely scalable, protocol agnostic, data-carrying mechanism that transfers data packets with assigned labels across the network through virtual links. Extensions of the BGP protocol can be used to manage an MPLS path. The Cisco IOS XE Catalyst SD-WAN devices also have the capability of BGP MPLS VPN Option B.

The multiple service VPNs use inter autonomous system (AS) BGP labeled path to forward the traffic, that in turn helps scale the service side VPNs with less control plane signaling. MPLS interface is supported only in global VRF.

To configure an MPLS interface, do the following:

Procedure

Step 1 Click **MPLS Interface**.

Step 2 Enter the interface name in the **Interface Name** field.

Step 3 You can click on + to add more interfaces and save the configuration.

Configure label range

Cisco SD-WAN Manager automatically programs the label space for BGP MPLS. The labels are allocated per VPN. To view the configuration, use the command, **show sdwan running-config**.

Sample configuration:

```
show sdwan running-config
mpls label range 100000 1048575 static 16 999
mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf
mpls label mode all-vrfs protocol bgp-vpnv6 per-vrf
```

Configure route targets

You can configure route targets on the Cisco IOS XE Catalyst SD-WAN devices. Route targets configuration is supported only on eBGP and IPv4 peer devices. All the supported protocols can be redistributed to BGP.

Procedure

Step 1 To configure route targets, click **Route Targets** and configure the following parameters:

Parameter	Option	Sub-Option	Description
IPv4 / IPv6			Click IPv4 to configure a route target for IPv4 interfaces. Click IPv6 to configure a route target for IPv6 interfaces.
Add VPN			Click Add VPN to add VPNs.
VPN ID for IPv4			Specify the VPN ID for IPv4 interface.
Import			Imports routing information from the target VPN extended community.
Export			Exports routing information to the target VPN extended community.

Step 2 Click **Save** to save the feature template.

Initially, the devices have default route targets, then you can add additional entries as required.

Configure advanced neighbor parameter

Procedure

Step 1 To configure advanced parameters for the neighbor, click **Neighbor > Advanced Options**.

Parameter Name	Description
Next-Hop Self	Click On to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Click On to send the local router's BGP community attribute to the BGP neighbor.
Send Extended Community	Click On to send the local router's BGP extended community attribute to the BGP neighbor.
Negotiate Capability	Click On to allow the BGP session to learn about the BGP extensions that are supported by the neighbor.
Source Interface Address	Enter the IP address of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor.
Source Interface Name	Enter the name of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor, in the format <i>ge port/slot</i> .
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 0 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered available. Specify the keepalive time for the neighbor to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Hold Time	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive timer)
Connection Retry Time	Specify the number of seconds between retries to establish a connection to a configured BGP neighbor peer that has gone down. Range: 0 through 65535 seconds Default: 30 seconds
Advertisement Interval	For the BGP neighbor, set the minimum route advertisement interval (MRAI) between when BGP routing update packets are sent to that neighbor. Range: 0 through 600 seconds Default: 5 seconds for IBGP route advertisements; 30 seconds for EBGP route advertisements

Step 2 To save the feature template, click **Save**.

Change the scope of a parameter value

Before you begin

When you first open a feature template, for each parameter that has a default value, the scope is set to Default, and the default setting or value is shown.

Procedure

To change the default or to enter a value, click the scope drop-down list to the left of the parameter field and select one of the following:

Parameter Name	Description
Device Specific	<p>Use a device-specific value for the parameter. For device-specific parameters, you cannot enter a value in the feature template. You enter the value when you attach a device to a device template.</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a device to a device template.</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure advanced BGP parameters

Procedure

- Step 1** From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
- Step 2** Click **Device Templates**.
- Step 3** Click **Create Template**.
- Step 4** From the **Create Template** drop-down list, choose **From Feature Template**.

Step 5 From the **Device Model** drop-down list, choose the type of device for which you are creating the template.

Step 6 To create a template for **VPN 0** or **VPN 512**:

- a. Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
- b. Under **Additional VPN 0 Templates**, click **BGP**.
- c. From the **BGP** drop-down list, click **Create Template**. The BGP template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining BGP parameters.

Step 7 To configure advanced parameters for BGP, click **Advanced** and configure the following parameters:

Parameter Name	Description
Hold Time	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local device then terminates the BGP session to that peer. This hold time is the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive timer)
Keepalive	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local device is still active and should be considered available. This keepalive time is the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Compare MED	Click On to always compare MEDs regardless of whether the peer ASs of the compared routes are the same.
Deterministic MED	Click On to compare multiple exit discriminators (MEDs) from all routes received from the same AS, regardless of when the route was received.
Missing MED as Worst	Click On to consider a path as the worst path if the path is missing a MED attribute.
Compare Router ID	Click On to compare the device IDs among BGP paths to determine the active path.
Multipath Relax	Click On to have the BGP best-path process select from routes in different in ASs. By default, when you are using BGP multipath, the BGP best path process selects from routes in the same AS to load-balance across multiple paths.

Step 8 Click **Save**.

Configure BGP routing in a service profile using configuration group

Before you begin

On the **Configuration > Configuration Groups** page, choose **SD-WAN** as the solution type.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration > Configuration Groups**.

Step 2 Create and configure a BGP Routing feature in a Service profile.

- a. Configure Basic Configuration fields.

Table 2: Basic Configuration

Field	Description
AS Number	Enter the local AS number.
Router ID	Enter the BGP router ID, in decimal four-part dotted notation.
Propagate AS Path	Enable this option to carry BGP AS path information into OMP.
Propagate Community	Enable this option to propagate BGP communities between Cisco Catalyst SD-WAN sites, across VPNs using OMP redistribution.
External Routes Distance	Specify the BGP route administrative distance for routes learned from other sites in the overlay network. Range: 1 through 255 Default: 20
Internal Routes Distance	Enter a value to apply as the BGP route administrative distance for routes coming from one AS into another. Range: 1 through 255 Default: 200
Local Routes Distance	Specify the BGP route administrative distance for routes within the local AS. By default, a route received locally from BGP is preferred over a route received from OMP. Range: 1 through 255 Default: 20

- b. Configure Unicast Address Family fields.

Table 3: Unicast Address Family

Field	Description
IPv4 Settings	
Maximum Paths	Specify the maximum number of parallel internal BGP paths that can be installed into a route table to enable internal BGP multipath load sharing. Range: 0 to 32

Field	Description
Originate	Enable this option to allow the default route to be artificially generated and injected into the BGP Route Information Base (RIB), regardless of whether it is present in the routing table. The newly injected default is advertised to all the BGP peers.
Redistribute	
Protocol*	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are static , connected , ospf , omp , eigrp , and nat . At a minimum, choose omp . By default, OMP routes are not redistributed into BGP.
Route Policy	Enter the name of the route policy to apply to redistributed routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Network	
Network Prefix*	Enter a network prefix to be advertised by BGP. The network prefix is composed of the IPv4 subnet and the mask. For example, 192.0.2.0 and 255.255.255.0.
Aggregate Address	
Aggregate Prefix*	Enter the prefix of the addresses to aggregate for all BGP sessions. The aggregate prefix is composed of the IPv4 subnet and the mask. For example, 192.0.2.0 and 255.255.255.0.
AS Set Path	Enable this option to generate set path information for the aggregated prefixes.
Summary Only	Enable this option to filter out more specific routes from BGP updates.
Table Map	
Policy Name	Enter the route map that controls the downloading of routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Filter	When you enable this option, the route map specified in the Policy Name field controls whether a BGP route is to be downloaded to the Route Information Base (RIB). A BGP route is not downloaded to the RIB if it is denied by the route map. When you disable this option, the route map specified in the Policy Name field is used to set certain properties, such as the traffic index, of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.
IPv6 Settings	
Maximum Paths	Specify the maximum number of parallel internal BGP paths that can be installed into a route table to enable internal BGP multipath load sharing. Range: 0 to 32

Field	Description
Originate	Enable this option to allow the default route to be artificially generated and injected into the BGP RIB, regardless of whether it is present in the routing table. The newly injected default is advertised to all the BGP peers.
Redistribute	
Protocol*	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are static , connected , ospf , omp , and eigrp . At a minimum, choose omp . By default, OMP routes are not redistributed into BGP.
Route Policy	Enter the name of the route policy to apply to redistributed routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Network	
Network Prefix*	Enter a network prefix to be advertised by BGP. The IPv6 network prefix is composed of the IPv6 address and the prefix length (1-128). For example, the IPv6 subnet is 2001:DB8:0000:0000:: and the prefix length is 64.
Aggregate Address	
Aggregate Prefix*	Enter the prefix of the addresses to aggregate for all BGP sessions. The IPv6 aggregate prefix is composed of the IPv6 address and the prefix length (1-128). For example, the IPv6 subnet is 2001:DB8:0000:0000:: and the prefix length is 64.
AS Set Path	Enable this option to generate set path information for the aggregated prefixes.
Summary Only	Enable this option to filter out more specific routes from BGP updates.
Table Map	
Policy Name*	Enter the route map that controls the downloading of routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Filter	When you enable this option, the route map specified in the Policy Name field controls whether a BGP route is to be downloaded to the Route Information Base (RIB). A BGP route is not downloaded to the RIB if it is denied by the route map. When you disable this option, the route map specified in the Policy Name field is used to set certain properties, such as the traffic index, of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.

- c. Configure Neighbor fields.

Table 4: Neighbor

Field	Description
IPv4 Settings	
Address*	Specify the IP address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Interface Name	Enter the interface name. This interface is used as the source of the TCP session when establishing neighborship. We recommend that you use a loopback interface.
Allowas in Number	Enter the number of times to allow the advertisement of the autonomous system number (ASN) of a provider edge (PE) device. The range is 1 to 10. If no number is specified, the default value of three times is used.
AS Override	Enable this option to replace the AS number of the originating router with the AS number of the sending BGP router.
Shutdown	Disable this option to enable BGP for the VPN.
Advanced Options	
Next-Hop Self	Enable this option to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Enable this option to send the BGP community attribute of the local router to the BGP neighbor.
Send Extended Community	Enable this option to send the BGP extended community attribute of the local router to the BGP neighbor.
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 1 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time (seconds)	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered to be available. Specify the keepalive time for the neighbor, to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)

Field	Description
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Send Label	Enable this option to allow the routers advertise to each other so that they can send MPLS labels with the routes. If the routers successfully negotiate their ability to send MPLS labels, the routers add MPLS labels to all the outgoing BGP updates.
Add Neighbor Address Family	
Family Type*	Choose the BGP IPv4 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Maximum Prefix Reach Policy*	Choose one of the following options: <ul style="list-style-type: none"> • Policy Off: Policy is off. • Policy On - Restart: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: <ul style="list-style-type: none"> • Maximum Number of Prefixes*: Enter the maximum prefix limit. Range: 1 to 4294967295 • Threshold (percentage): Enter the threshold value: Range: 1 to 100 Default: 75 • Restart Interval (minutes)*: Enter the time interval. Range: 1 to 65535 minutes • Policy On - Warning message: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • Policy On - Disable Peer Neighbor: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

Field	Description
IPv6 Settings	
Address*	Specify the IP address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Interface Name	Enter the interface name. This interface is used as the source of the TCP session when establishing neighborhood. We recommend that you use a loopback interface.
Allowas in Number	Enter the number of times to allow the advertisement of the autonomous system number (ASN) of a provider edge (PE) device. The range is 1 to 10. If no number is specified, the default value of three times is used.
AS Override	Enable this option to replace the AS number of the originating router with the AS number of the sending BGP router.
Shutdown	Disable this option to enable BGP for the VPN.
Advanced Options	
Next-Hop Self	Enable this option to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Enable this option to send the BGP community attribute of the local router to the BGP neighbor.
Send Extended Community	Enable this option to send the BGP extended community attribute of the local router to the BGP neighbor.
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 1 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time (seconds)	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered to be available. Specify the keepalive time for the neighbor, to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)

Field	Description
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Add IPv6 Neighbor Address Family	
Family Type*	Choose the BGP IPv6 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Maximum Prefix Reach Policy*	Choose one of the following options: <ul style="list-style-type: none"> • Policy Off: Policy is off. • Policy On - Restart: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: <ul style="list-style-type: none"> • Maximum Number of Prefixes*: Enter the maximum prefix limit. Range: 1 to 4294967295 • Threshold (percentage): Enter the threshold value: Range: 1 to 100 Default: 75 • Restart Interval (minutes)*: Enter the time interval. Range: 1 to 65535 minutes • Policy On - Warning message: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • Policy On - Disable Peer Neighbor: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

Configure BGP routing in a transport profile using a configuration group

Before you begin

On the **Configuration > Configuration Groups** page, choose **SD-WAN** as the solution type.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration > Configuration Groups**.

Step 2 Create and configure BGP Routing in Transport and Management Profile.

- a. Configure Basic Configuration fields.

Table 5: Basic Configuration

Field	Description
AS Number	Enter the local AS number.
Router ID	Enter the BGP router ID, in decimal four-part dotted notation.
Propagate AS Path	Enable this option to carry BGP AS path information into OMP.
Propagate Community	Enable this option to propagate BGP communities between Cisco Catalyst SD-WAN sites, across VPNs using OMP redistribution.
External Routes Distance	Specify the BGP route administrative distance for routes learned from other sites in the overlay network. Range: 1 through 255 Default: 20
Internal Routes Distance	Enter a value to apply as the BGP route administrative distance for routes coming from one AS into another. Range: 1 through 255 Default: 200
Local Routes Distance	Specify the BGP route administrative distance for routes within the local AS. By default, a route received locally from BGP is preferred over a route received from OMP. Range: 1 through 255 Default: 20

- b. Configure Unicast Address Family.

Table 6: Unicast Address Family

Field	Description
IPv4 Settings	

Field	Description
Maximum Paths	Specify the maximum number of parallel internal BGP paths that can be installed into a route table to enable internal BGP multipath load sharing. Range: 0 to 32
Originate	Enable this option to allow the default route to be artificially generated and injected into the BGP Route Information Base (RIB), regardless of whether it is present in the routing table. The newly injected default is advertised to all the BGP peers.
Redistribute	
Protocol*	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are static , connected , ospf , omp , eigrp , and nat . At a minimum, choose connected , and then under Route Policy , specify a route policy that has BGP advertise the loopback interface address to its neighbors. Route policy is not supported in Cisco vManage Release 20.9.1.
Route Policy	Enter the name of the route policy to apply to redistributed routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Network	
Network Prefix*	Enter a network prefix to be advertised by BGP. The network prefix is composed of the IPv4 subnet and the mask. For example, 192.0.2.0 and 255.255.255.0.
Aggregate Address	
Aggregate Prefix*	Enter the prefix of the addresses to aggregate for all BGP sessions. The aggregate prefix is composed of the IPv4 subnet and the mask. For example, 192.0.2.0 and 255.255.255.0.
AS Set Path	Enable this option to generate set path information for the aggregated prefixes.
Summary Only	Enable this option to filter out more specific routes from BGP updates.
Table Map	
Policy Name	Enter the route map that controls the downloading of routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Filter	When you enable this option, the route map specified in the Policy Name field controls whether a BGP route is to be downloaded to the Route Information Base (RIB). A BGP route is not downloaded to the RIB if it is denied by the route map. When you disable this option, the route map specified in the Policy Name field is used to set certain properties, such as the traffic index, of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.

Field	Description
IPv6 Settings	
Maximum Paths	Specify the maximum number of parallel internal BGP paths that can be installed into a route table to enable internal BGP multipath load sharing. Range: 0 to 32
Originate	Enable this option to allow the default route to be artificially generated and injected into the BGP Route Information Base (RIB), regardless of whether it is present in the routing table. The newly injected default is advertised to all the BGP peers.
Redistribute	
Protocol*	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are static , connected , ospf , omp , and eigrp . At a minimum, choose connected , and then under Route Policy , specify a route policy that has BGP advertise the loopback interface address to its neighbors. Route policy is not supported in Cisco vManage Release 20.9.1.
Route Policy	Enter the name of the route policy to apply to redistributed routes. Route policy is not supported in Cisco vManage Release 20.9.1.
Network	
Network Prefix*	Enter a network prefix to be advertised by BGP. The IPv6 network prefix is composed of the IPv6 address and the prefix length (1-128). For example, the IPv6 subnet is 2001:DB8:0000:0000:: and the prefix length is 64.
Aggregate Address	
Aggregate Prefix*	Enter the prefix of the addresses to aggregate for all BGP sessions. The IPv6 aggregate prefix is composed of the IPv6 address and the prefix length (1-128). For example, the IPv6 subnet is 2001:DB8:0000:0000:: and the prefix length is 64.
AS Set Path	Enable this option to generate set path information for the aggregated prefixes.
Summary Only	Enable this option to filter out more specific routes from BGP updates.
Table Map	
Policy Name	Enter the route map that controls the downloading of routes. Route policy is not supported in Cisco vManage Release 20.9.1.

Field	Description
Filter	<p>When you enable this option, the route map specified in the Policy Name field controls whether a BGP route is to be downloaded to the Route Information Base (RIB). A BGP route is not downloaded to the RIB if it is denied by the route map.</p> <p>When you disable this option, the route map specified in the Policy Name field is used to set certain properties, such as the traffic index, of the routes for installation into the RIB. The route is always downloaded, regardless of whether it is permitted or denied by the route map.</p>

- c. Configure MPLS Interface.

Table 7: MPLS Interface

Field	Description
Interface Name*	Enter a name for the MPLS interface.

- d. Configure Neighbor.

Table 8: Neighbor

Field	Description
IPv4 Settings	
Address*	Specify the IP address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Interface Name	Enter the interface name. This interface is used as the source of the TCP session when establishing neighborhood. We recommend that you use a loopback interface.
Allows in Number	Enter the number of times to allow the advertisement of the autonomous system number (ASN) of a provider edge (PE) device. The range is 1 to 10. If no number is specified, the default value of three times is used.
AS Override	Enable this option to replace the AS number of the originating router with the AS number of the sending BGP router.
Shutdown	Disable this option to enable BGP for the VPN.
Advanced Options	
Next-Hop Self	Enable this option to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Enable this option to send the BGP community attribute of the local router to the BGP neighbor.

Field	Description
Send Extended Community	Enable this option to send the BGP extended community attribute of the local router to the BGP neighbor.
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 1 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time (seconds)	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered to be available. Specify the keepalive time for the neighbor, to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Send Label	Enable this option to allow the routers advertise to each other so that they can send MPLS labels with the routes. If the routers successfully negotiate their ability to send MPLS labels, the routers add MPLS labels to all the outgoing BGP updates.
Add Neighbor Address Family	
Family Type*	Choose the BGP IPv4 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.

Field	Description
Maximum Prefix Reach Policy*	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Policy Off: Policy is off. • Policy On - Restart: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: <ul style="list-style-type: none"> • Maximum Number of Prefixes*: Enter the maximum prefix limit. Range: 1 to 4294967295 • Threshold (percentage): Enter the threshold value: Range: 1 to 100 Default: 75 • Restart Interval (minutes)*: Enter the time interval. Range: 1 to 65535 minutes • Policy On - Warning message: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • Policy On - Disable Peer Neighbor: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.
IPv6 Settings	
Address*	Specify the IP address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Interface Name	Enter the interface name. This interface is used as the source of the TCP session when establishing neighborhood. We recommend that you use a loopback interface.
Allowas in Number	Enter the number of times to allow the advertisement of the autonomous system number (ASN) of a provider edge (PE) device. The range is 1 to 10. If no number is specified, the default value of three times is used.
AS Override	Enable this option to replace the AS number of the originating router with the AS number of the sending BGP router.
Shutdown	Disable this option to enable BGP for the VPN.
Advanced Options	
Next-Hop Self	Enable this option to configure the router to be the next hop for routes advertised to the BGP neighbor.

Field	Description
Send Community	Enable this option to send the BGP community attribute of the local router to the BGP neighbor.
Send Extended Community	Enable this option to send the BGP extended community attribute of the local router to the BGP neighbor.
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 1 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time (seconds)	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered to be available. Specify the keepalive time for the neighbor, to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Add IPv6 Neighbor Address Family	
Family Type*	Choose the BGP IPv6 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.

Field	Description
Maximum Prefix Reach Policy*	<p>Choose one of the following options:</p> <ul style="list-style-type: none"> • Policy Off: Policy is off. • Policy On - Restart: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: <ul style="list-style-type: none"> • Maximum Number of Prefixes*: Enter the maximum prefix limit. Range: 1 to 4294967295 • Threshold (percentage): Enter the threshold value: Range: 1 to 100 Default: 75 • Restart Interval (minutes)*: Enter the time interval. Range: 1 to 65535 minutes • Policy On - Warning message: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • Policy On - Disable Peer Neighbor: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

e. Configure Advanced fields.

Table 9: Advanced

Field	Description
Keepalive (seconds)	<p>Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered to be available. This keepalive time is the global keepalive time.</p> <p>Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)</p>
Hold Time (seconds)	<p>Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. This hold time is the global hold time.</p> <p>Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)</p>
Compare MED	<p>Enable this option to compare the router IDs among BGP paths to determine the active path.</p>

Field	Description
Deterministic MED	Enable this option to compare MEDs from all routes received from the same AS regardless of when the route was received.
Missing MED as Worst	Enable this option to consider a path as the worst path if the path is missing a MED attribute.
Compare Router ID	Enable this option to always compare MEDs regardless of whether the peer ASs of the compared routes are the same.
Multipath Relax	Enable this option to have the BGP best-path process select from routes in different ASs. By default, when you are using BGP multipath, the BGP best-path process selects from routes in the same AS to load-balance across multiple paths.

Configure BGP using CLI

This is an example of a BGP configuration on a Cisco IOS XE Catalyst SD-WAN device for releases prior to Cisco IOS XE Catalyst SD-WAN Release 17.14.1a.

```

router bgp 100
  bgp log-neighbor-changes
  distance bgp 20 200 20
  !
  address-family ipv4 vrf 100
    bgp router-id 10.0.0.0
    redistribute omp
    neighbor 10.0.0.1 remote-as 200
    neighbor 10.0.0.1 activate
    neighbor 10.0.0.1 send-community both
    neighbor 10.0.0.1 route-map OMP_BGP-POLICY in
    neighbor 10.0.0.1 maximum-prefix 2147483647 100

route-map OMP_BGP-POLICY permit 1
  match ip address prefix-list OMP-BGP-TEST-PREFIX-LIST
  set omp-tag 10000
route-map OMP_BGP-POLICY permit 65535

ip prefix-list OMP-BGP-TEST-PREFIX-LIST seq 5 permit 10.0.0.0/8

```



Note Starting from Cisco IOS XE Catalyst SD-WAN Release 17.14.1a, the following changes apply to BGP configuration under non-VRF address-family:

- The keyword **remote-as** is not supported under the non-VRF **address-family** command. For non-VRF address-family, the remote-as ASN must be configured under router bgp mode.
- BGP distance configuration is not supported under router bgp mode. BGP distance must be configured under the specified non-VRF address-family.

You must update the device CLI template or the CLI Add-on feature template manually to modify the configuration to incorporate the changes introduced.

Following is the sample BGP configuration for Cisco IOS XE Catalyst SD-WAN Release 17.14.1a and later:

```
router bgp 100
  neighbor 10.10.10.10 remote-as
  address-family ipv4
    distance bgp 20 200 200
    neighbor 10.10.10.10 activate
  address-family ipv4 unicast vrf RED
    distance bgp 30 300 300
    neighbor 10.11.11.11 remote-as
    neighbor 10.11.11.11 activate
```

Example of configuring Service-Side routing using CLI

To set up routing on the Cisco IOS XE Catalyst SD-WAN device, you provision one VPN or multiple VPNs if segmentation is required. Within each VPN, you configure the interfaces that participate in that VPN and the routing protocols that operate in that VPN.

vpn-ID can be any service-side VPN, which is a VPN other than VPN 0 and VPN 512. VPN 0 is the transport VPN and carries only control traffic, and VPN 512 is the management VPN.

```
Device(config)# vrf definition vpn-id
Device(config-vrf)# address-family ipv4
Device(config-ipv4)# exit
Device(config-vrf)# address-family ipv6
Device(config-ipv6)# exit
Device(config-vrf)# exit
Device(config)#
```

Configure the local AS number: You can specify the AS number in 2-byte ASDOT notation (1 through 65535) or in 4-byte ASDOT notation (1.0 through 65535.65535).

```
Device(config)# router bgp local-as-number
Device(config-router)# address-family ipv4 unicast vrf vpn-id
```

Configure the BGP peer, specifying its address and AS number (the remote AS number), and enable the connection to the peer:

```
Device(config-router-af)# neighbor neighbor-ip-address remote-as remote-as-number
```

Configure a system IP address for the Cisco IOS XE Catalyst SD-WAN device:

```
Device(config)# system system-ipaddress
```

Example of BGP Configuration on a Cisco IOS XE Catalyst SD-WAN device

```

Device# show running-config system
system
  system-ip 10.1.2.3
!
Device# show running-config vpn 1
router bgp 2
  bgp log-neighbor-changes
  timers bgp 1 111
  neighbor 10.20.25.16 remote-as 1

!
address-family ipv4 unicast
  neighbor 10.20.25.16 activate
exit-address-family
!
address-family vpv4 unicast
  neighbor 10.20.25.16 activate
  neighbor 10.20.25.16 send-community extended
exit-address-family
!
address-family vpv6 unicast
  neighbor 10.20.25.16 activate
  neighbor 10.20.25.16 send-community extended
exit-address-family
!
address-family ipv4 unicast vrf 1
  redistribute connected
  redistribute static
exit-address-family
!
address-family ipv6 unicast vrf 1
  redistribute connected
  redistribute omp

exit-address-family
!
address-family ipv4 unicast vrf 2
  redistribute connected

exit-address-family

```

Example of configuring route targets:

```

vrf config

vrf definition 1
  rd 1:1

!
address-family ipv4

  route-target export 200:1

  route-target import 100:1

exit-address-family
!
address-family ipv6
  route-target export 101:1
  route-target import 201:1
exit-address-family

```

Verify BGP redistribute route in OMP

To verify BGP redistribute route in OMP:

```
Device# show sdwan omp routes 10.0.0.0/8
-----
omp route entries for vpn 100 route 10.0.0.0/8
-----
RECEIVED FROM:
peer          172.16.0.0
path-id       470777
label         1002
status        C,I,R
loss-reason   not set
lost-to-peer  not set
lost-to-path-id not set
Attributes:
  originator   10.0.0.1
  type         installed
  tloc         172.16.0.1, mpls, ipsec
  ultimate-tloc not set
  domain-id    not set
  overlay-id   1
  site-id      1
  preference   not set
  tag          10000 <=====
  origin-proto eBGP
  as-path      not set
  unknown-attr-len not set
```

The following example shows the propagation of BGP community on Cisco IOS XE Catalyst SD-WAN devices:

```
Device# show sdwan omp routes 192.168.0.0/16 detail
-----
omp route entries for vpn 1 route
192.168.0.0/16-----
RECEIVED FROM:
peer          10.0.0.0
path-id       70
label         1007
status        C,Red,R
loss-reason   not set
lost-to-peer  not set
lost-to-path-id not set
Attributes:
  originator   192.168.0.0
  type         installed
  tloc         192.168.0.1, lte, ipsec
  ultimate-tloc not set
  domain-id    not set
  overlay-id   1
  site-id      500
  preference   not set
  tag          not set
  origin-proto iBGP
  origin-metric 0
  as-path      not set
  community    100:1 100:2 100:3
  unknown-attr-len not set
  ADVERTISED TO:
```

```
peer 192.168.0.1
```

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.18.1a and Cisco IOS XE Catalyst SD-WAN Release 17.15.2, the **show sdwan omp routes** command requires specifying both the **tenant** *tenant-id* and **vpn** *vpn-id* when used with a prefix address.

```
Device# show sdwan omp routes tenant 0 vpn 1 10.2.2.0
```

```
Generating output, this might take time, please wait ...
```

```
Code:
```

```
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA -> On-demand inactive
U -> TLOC unresolved
BR-R -> Border-Router reoriginated
TGW-R -> Transport-Gateway reoriginated
R-TGW-R -> Reoriginated Transport-Gateway reoriginated
DERIVED
AFFINITY AFFINITY
PATH PSEUDO
GROUP GROUP
FROM PEER ID LABEL STATUS KEY TLOC IP COLOR
ENCAP PREFERENCE NUMBER NUMBER
```

```
-----
12.16.255.20 1 1003 C,I,R 1 10.2.1.25 publicinternet
ipsec - None None
12.16.255.20 2 1003 C,I,R 1 10.2.1.25 lte
ipsec - None None
12.16.255.20 3 1003 C,I,R 1 10.2.1.25 gold
ipsec - None None
12.16.255.20 4 1003 C,I,R 1 10.2.1.25 silver
ipsec - None None
12.16.255.20 5 1003 C,I,R 1 10.2.1.25
```

Redistribute BGP routes and AS path information

By default, routes from other routing protocols are not redistributed into BGP. It can be useful for BGP to learn OMP routes, because OMP learns routes to destinations throughout the overlay network. BGP on the Cisco IOS XE Catalyst SD-WAN devices, then advertises the OMP routes to all the BGP routers in the service-side of the network.

```
config-transaction
router bgp 2
address-family ipv4 unicast
redistribute omp route-map route_map
```

To redistribute OMP routes into BGP so that these routes are advertised to all BGP routers in the service side of the network, configure redistribution in any VRF except transport VRF or Mgmt VRF:

For Cisco IOS XE Catalyst SD-WAN device, under router BGP configuration, **redistribute omp route-map set/match** is used instead of **redistribute omp metric 0** setting as the **redistribute omp metric** is disabled in all the branches.

```
Device(config)# router bgp 100
Device(config-router)# address-family ipv4 vrf 100
Device(config-router-af)# redistribute omp [route-map policy-name]
```

```
config-transaction
router bgp 100
address-family ipv4 vrf 100
redistribute omp route_map route_map
```

You can also redistribute routes learned from other protocols such as OSPF, rip into BGP, and apply policy as shown in the example above:

You can control redistribution of routes on a per-neighbor basis:

```
config-transaction
router bgp 100
address-family ipv4
neighbor 10.0.100.1 route-map route_map (in | out)
```

You can configure the SD-WAN Controller to advertise BGP routes that it has learned, through OMP, from the SD-WAN Controller. Doing so allows the SD-WAN Controller to advertise these routes to other Cisco IOS XE Catalyst SD-WAN devices in the overlay network. You can advertise BGP routes either globally or for a specific VRF:

```
config-transaction
sdwan
omp
address-family ipv4 vrf 100
advertise bgp
exit
```

