

Unicast Overlay Routing

The overlay network is controlled by the Cisco Catalyst SD-WAN Overlay Management Protocol (OMP), which is at the heart of Cisco Catalyst SD-WAN overlay routing. This solution allows the building of scalable, dynamic, on-demand, and secure VPNs. The Cisco Catalyst SD-WAN solution uses a centralized controller for easy orchestration, with full policy control that includes granular access control and a scalable secure data plane between all edge nodes.

The Cisco Catalyst SD-WAN solution allows edge nodes to communicate directly over any type of transport network, whether public WAN, internet, metro Ethernet, MPLS, or anything else.

- Supported Protocols, on page 1
- Configure Unicast Overlay Routing, on page 15

Supported Protocols

This section explains the protocols supported for unicast routing.

OMP Routing Protocol

The Cisco Catalyst SD-WAN Overlay Management Protocol (OMP) is the protocol responsible for establishing and maintaining the Cisco Catalyst SD-WAN control plane. It provides the following services:

- Orchestration of overlay network communication, including connectivity among network sites, service chaining, and VPN or VRF topologies
- Distribution of service-level routing information and related location mappings
- Distribution of data plane security parameters
- Central control and distribution of routing policy

OMP is the control protocol that is used to exchange routing, policy, and management information between Cisco Catalyst SD-WAN Controllers and Cisco vEdge devices in the overlay network. These devices automatically initiate OMP peering sessions between themselves, and the two IP end points of the OMP session are the system IP addresses of the two devices.

OMP is an all-encompassing information management and distribution protocol that enables the overlay network by separating services from transport. Services provided in a typical VPN setting are usually located within a VPN domain, and they are protected so that they are not visible outside the VPN. In such a traditional architecture, it is a challenge to extend VPN domains and service connectivity.

OMP addresses these scalability challenges by providing an efficient way to manage service traffic based on the location of logical transport end points. This method extends the data plane and control plane separation concept from within routers to across the network. OMP distributes control plane information along with related policies. A central Cisco Catalyst SD-WAN Controller makes all decisions related to routing and access policies for the overlay routing domain. OMP is then used to propagate routing, security, services, and policies that are used by edge devices for data plane connectivity and transport.

OMP Route Advertisements

On Cisco Catalyst SD-WAN Controllers and Cisco vEdge devices, OMP advertises to its peers the routes and services that it has learned from its local site, along with their corresponding transport location mappings, which are called TLOCs. These routes are called OMP routes or vRoutes to distinguish them from standard IP routes. The routes advertised are actually a tuple consisting of the route and the TLOC associated with that route. It is through OMP routes that the Cisco Catalyst SD-WAN Controllers learn the topology of the overlay network and the services available in the network.

OMP interacts with traditional routing at local sites in the overlay network. It imports information from traditional routing protocols, such as OSPF and BGP, and this routing information provides reachability within the local site. The importing of routing information from traditional routing protocols is subject to user-defined policies.

Because OMP operates in an overlay networking environment, the notion of routing peers is different from a traditional network environment. From a logical point of view, the overlay environment consists of a centralized controller and a number of edge devices. Each edge device advertises its imported routes to the centralized controller and based on policy decisions, this controller distributes the overlay routing information to other edge devices in the network. Edge devices never advertise routing information to each other, either using OMP or any other method. The OMP peering sessions between the centralized controller and the edge devices are used exclusively to exchange control plane traffic; they are never, in any situation, used for data traffic.

Registered edge devices automatically collect routes from directly connected networks as well as static routes and routes learned from IGP protocols. The edge devices can also be configured to collect routes learned from BGP.

Route map AS path and community configuration, for example, AS path prepend, are not supported when route-maps are configured for protocol redistribution. The AS path for redistributed OMP routes can be configured and applied by using a route map on the BGP neighbor outbound policy.

OMP performs path selection, loop avoidance, and policy implementation on each local device to decide which routes are installed in the local routing table of any edge device.



Note Route advertisements to OMP are done by either applying the configuration at the global level or at the specific VPN level. To configure route advertisements to OMP at the global level, use the OMP feature template. On the other hand, to configure route advertisements to OMP at the specific VPN level, use the VPN feature template. For more information about configuring route advertisements to OMP, see Configure OMP using templates, on page 59.



Note Any recursive lookup for service side routes over OMP protocol is not supported on Cisco Catalyst SD-WAN. Starting from Cisco IOS XE SD-WAN Release 17.12.1a, the recursive route lookup on service side routes over OMP protocol on Cisco IOS XE Catalyst SD-WAN is not supported.

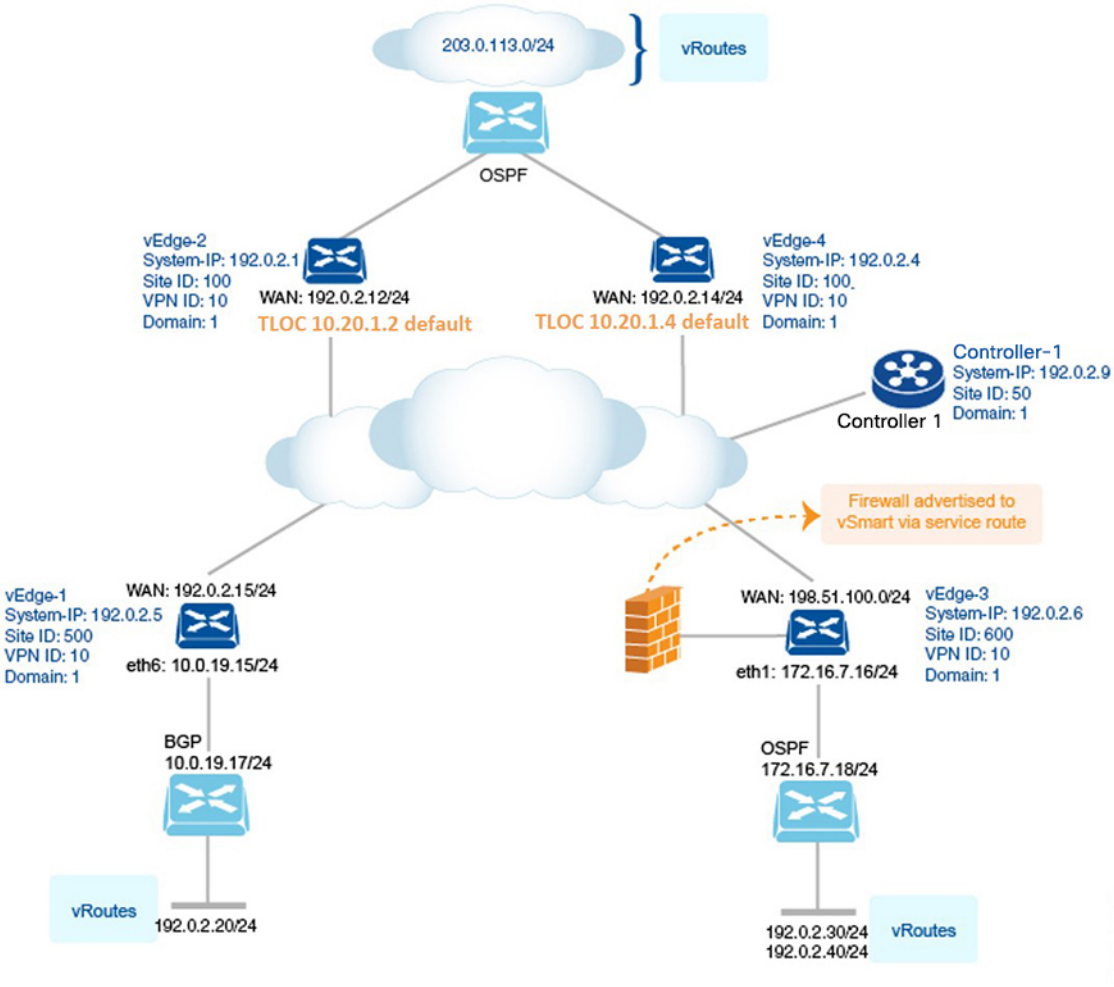
OMP advertises the following types of routes:

- OMP routes (also called vRoutes)—Prefixes that establish reachability between end points that use the OMP-orchestrated transport network. OMP routes can represent services in a central data center, services at a branch office, or collections of hosts and other end points in any location of the overlay network. OMP routes require and resolve into TLOCs for functional forwarding. In comparison with BGP, an OMP route is the equivalent of a prefix carried in any of the BGP AFI/SAFI NLRI fields (Address Family Indicator (AFI), Subsequent Address Family Identifiers (SAFI), Network Layer Reachability Information (NLRI)) fields).
- Transport locations (TLOCs)—Identifiers that tie an OMP route to a physical location. The TLOC is the only entity of the OMP routing domain that is visible to the underlying network, and it must be reachable via routing in the underlying network. A TLOC can be directly reachable via an entry in the routing table of the physical network, or it can be represented by a prefix residing on the outside of a NAT device and must be included in the routing table. In comparison with BGP, the TLOC acts as the next hop for OMP routes.
- Service routes—Identifiers that tie an OMP route to a service in the network, specifying the location of the service in the network. Services include firewalls, Intrusion Detection Systems (IDPs), and load balancers. Service route information is carried in both service and OMP routes.

(OMP also advertises policies configured on the Cisco Catalyst SD-WAN Controllers that are executed on Cisco vEdge devices including application-routing policy, cflowd flow templates, and data policy. For more information, see *Policy Overview*.)

The following figure illustrates the three types of OMP routes.

Figure 1: Different Types of OMP Routes



OMP Routes

Each device at a branch or local site advertises OMP routes to the Cisco Catalyst SD-WAN Controllers in its domain. These routes contain routing information that the device has learned from its site-local network.

A Cisco Catalyst SD-WAN device can advertise one of the following types of site-local routes:

- Connected (also known as direct)
- Static
- BGP
- OSPF (inter-area, intra-area, and external)
- IS-IS

OMP routes advertise the following attributes:

- TLOC—Transport location identifier of the next hop for the vRoute. It is similar to the BGP NEXT_HOP attribute. A TLOC consists of three components:
  - System IP address of the OMP speaker that originates the OMP route
  - Color to identify the link type
  - Encapsulation type on the transport tunnel
- Origin—Source of the route, such as BGP, OSPF, connected, and static, and the metric associated with the original route.
- Originator—OMP identifier of the originator of the route, which is the IP address from which the route was learned.
- Preference—Degree of preference for an OMP route. A higher preference value is more preferred.
- Service—Network service associated with the OMP route.
- Site ID—Identifier of a site within the Cisco Catalyst SD-WAN overlay network domain to which the OMP route belongs.
- Tag—Optional, transitive path attribute that an OMP speaker can use to control the routing information it accepts, prefers, or redistributes.
- VPN—VPN or network segment to which the OMP route belongs.

You configure some of the OMP route attribute values, including the system IP, color, encapsulation type, carrier, preference, service, site ID, and VPN. You can modify some of the OMP route attributes by provisioning control policy on the Cisco Catalyst SD-WAN Controller.

TLOC Routes

TLOC routes identify transport locations. These are locations in the overlay network that connect to physical transport, such as the point at which a WAN interface connects to a carrier. A TLOC is denoted by a 3-tuple that consists of the system IP address of the OMP speaker, a color, and an encapsulation type. OMP advertises each TLOC separately.

TLOC routes advertise the following attributes:

- TLOC private address—Private IP address of the interface associated with the TLOC.
- TLOC public address—NAT-translated address of the TLOC.
- Carrier—An identifier of the carrier type, which is generally used to indicate whether the transport is public or private.
- Color—Identifies the link type.
- Encapsulation type—Tunnel encapsulation type.
- Preference—Degree of preference that is used to differentiate between TLOCs that advertise the same OMP route.
- Site ID—Identifier of a site within the Cisco Catalyst SD-WAN overlay network domain to which the TLOC belongs.

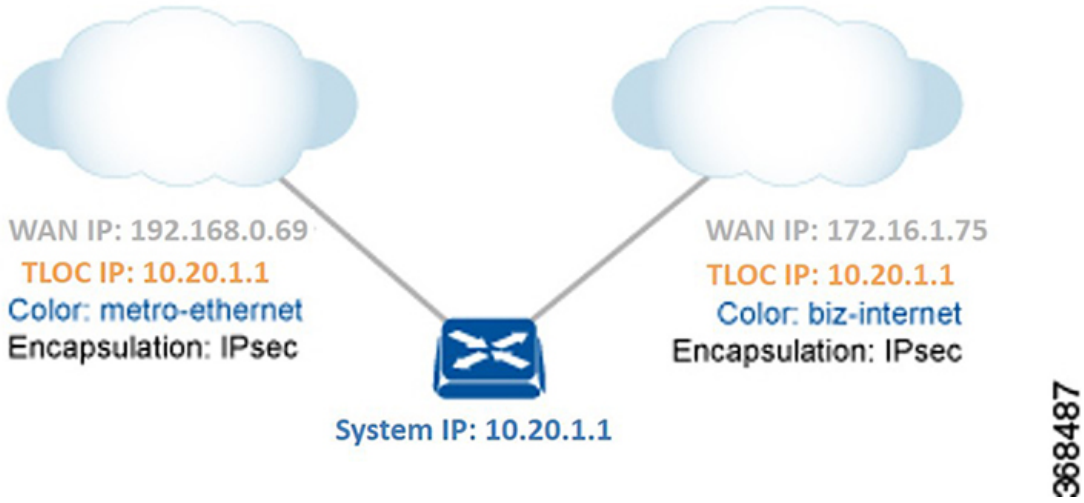
- Tag—Optional, transitive path attribute that an OMP speaker can use to control the flow of routing information toward a TLOC. When an OMP route is advertised along with its TLOC, both or either can be distributed with a community TAG, to be used to decide how to send traffic to or receive traffic from a group of TLOCs.
- Weight—Value that is used to discriminate among multiple entry points if an OMP route is reachable through two or more TLOCs.

The IP address used in the TLOC is the fixed system address of the device itself. The reason for not using an IP address or an interface IP address to denote a TLOC is that IP addresses can move or change; for example, they can be assigned by DHCP, or interface cards can be swapped. Using the system IP address to identify a TLOC ensures that a transport end point can always be identified regardless of IP addressing.

The link color represents the type of WAN interfaces on a device. The Cisco Catalyst SD-WAN solution offers predefined colors, which are assigned in the configuration of the devices. The color can be one of default, 3g, biz-internet, blue, bronze, custom1, custom2, custom3, gold, green, lte, metro-ethernet, mpls, private1, private2, public-internet, red, or silver.

The encapsulation is that used on the tunnel interface. It can be either IPsec or GRE.

Figure 2: Router Attributes



The diagram to the right shows a device that has two WAN connections and hence two TLOCs. The system IP address of the router is 10.20.1.1. The TLOC on the left is uniquely identified by the system IP address 10.20.1.1, the color metro-ethernet, and the encapsulation IPsec, and it maps to the physical WAN interface with the IP address 192.168.0.69. The TLOC on the right is uniquely identified by the system IP address 10.20.1.1, the color biz-internet, and the encapsulation IPsec, and it maps to the WAN IP address 172.16.1.75.

You configure some of the TLOC attributes, including the system IP address, color, and encapsulation, and you can modify some of them by provisioning control policy on the Cisco Catalyst SD-WAN Controller. See *Centralized Control Policy*.

Service Routes

Service routes represent services that are connected to a Cisco vEdge device or to the local-site network in which the Cisco vEdge device resides. The Cisco vEdge device advertises these routes to Cisco Catalyst SD-WAN Controllers using service address family NLRI. See *Service Chaining*.

OMP Route Advertisements for Cisco Catalyst SD-WAN Controllers

Table 1: Feature History

Feature Name	Release Information	Description
Increased OMP Path Limit for Cisco Catalyst SD-WAN Controllers	Cisco SD-WAN Release 20.5.1	This feature extends the limit on the number of OMP routes that can be exchanged between Cisco Catalyst SD-WAN Controllers to 128. Prior to this release, the limit was 16.

Overview

The transport location (TLOC) information is advertised to the OMP peers including Cisco Catalyst SD-WAN Controllers and its local-site branches. Starting from Cisco SD-WAN Release 20.5.1, the limit on the number of OMP paths that can be exchanged between Cisco Catalyst SD-WAN Controllers per VPN per prefix is extended to a maximum of 128.

Limitations

- Multitenant Cisco Catalyst SD-WAN Controllers only support global OMP configuration.
- The number of paths that are shared is dependent upon factors such as memory and the organization of internal data structure.

Configure Path Limit

The following example shows how to configure the number of paths that a Cisco Catalyst SD-WAN Controller can send to another Cisco Catalyst SD-WAN Controller:

```
Device(config)# omp
Device(config-omp)# controller-send-path-limit 100
```

Use the **controller-send-path-limit** command to configure maximum 128 send path limit to be exchanged between Cisco Catalyst SD-WAN Controllers. Use the **no** form of this command to set the send path limit to default. The default configuration enables the controllers to send the information of all the paths available up to maximum of 128.



Note We recommend using the default configuration, which sends information about all available paths, subject to a limit of 128 paths. This ensures that you have network visibility across controllers.

We recommend not to change the path limit frequently. For any changes on the peers, Cisco Catalyst SD-WAN Controller performs a full route database update. This leads to complete network updates.

For more information about configuring path limits, see controller-send-path-limit command page.

OMP Route Redistribution

OMP automatically redistributes the following types of routes that it learns either locally or from its routing peers:

- Connected
- Static
- OSPF intra-area routes
- OSPF inter-area routes
- OSPFv3 intra-area routes (Address-Family IPv6)
- OSPFv3 inter-area routes (Address-Family IPv6)

To avoid routing loops and less than optimal routing, redistribution of following types of routes requires explicit configuration:

- BGP
- EIGRP
- LISP
- IS-IS
- OSPF external routes
- OSPFv3 external route (Address-Family IPv6)
- OSPFv3 all routes (Address-Family IPv4)

The **advertise network**<*ipv4-prefix*> command can be used to advertise a specific prefix when a non-OMP route corresponding to the prefix is present in the VRF IPv4 routing table. Note that this command is only supported for **address-family ipv4**.

The following is an example for advertise network configuration:

```
omp
  no shutdown
  graceful-restart
  address-family ipv4 vrf 1
   advertise connected
   advertise static
   advertise network X.X.X.X/X
  !
```

To avoid propagating excessive routing information from the edge to the access portion of the network, the routes that devices receive via OMP are not automatically redistributed into the other routing protocols running on the routers. If you want to redistribute the routes received via OMP, you must enable this redistribution locally on each device.

OMP sets the origin and sub-origin type in each OMP route to indicate the route's origin (see the table below). When selecting routes, the Cisco Catalyst SD-WAN Controller and the router take the origin type and subtype into consideration.

To configure redistribution of OSPF routes into OMP for VRF1, you need to configure **advertise ospf route-map <route-map-name> external**. The OSPF internal routes are redistributed into OMP by default without any explicit configuration.

The following example shows the redistribution of OSPF external routes on all VRFs:

```
omp
  no shutdown
```

```
 ecmp-limit       6
 graceful-restart
 no as-dot-notation
 timers
  holdtime               300
  graceful-restart-timer 120
 exit
 address-family ipv4
  advertise ospf external <-- This configuration implies OSPF Inter-Area/Intra-Area routes
& External routes are redistributed into OMP
  advertise connected
  advertise static
 !
```

The following example shows the redistribution of OSPF external routes for a specific VRF:

```
omp
 no shutdown
 ecmp-limit       6
 graceful-restart
 no as-dot-notation
 timers
  holdtime               300
  graceful-restart-timer 120
 exit
 address-family ipv4 vrf 1
  advertise ospf external
  advertise ospf route-map RLB
 !
```

With the **external** keyword, the configuration applies the supplied route-map to both external and internal OSPF routes (Intra-Area/Inter-Area).

The following example shows the redistribution of OSPFv3 external routes:

```
omp
 no shutdown
 ecmp-limit       6
 graceful-restart
 no as-dot-notation
 timers
  holdtime               300
  graceful-restart-timer 120
 exit
  address-family ipv6
  advertise ospfv3
  advertise ospf external

 !
```



Note Starting from Cisco IOS XE Catalyst SD-WAN Release 17.7.2, the real-time display of omp routes received and advertised in Cisco SD-WAN Manager are limited to only 4001 routes to avoid excessive CPU usage.

Table 2:

OMP Route Origin Type	OMP Route Origin Subtype
BGP	External Internal

OMP Route Origin Type	OMP Route Origin Subtype
Connected	—
OSPF	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
Static	—
IS-IS	Level 1 and level 2

OMP also carries the metric of the original route. A metric of 0 indicates a connected route.

Administrative Distance

Administrative distance is the metric used to select the best path when there are two or more different routes to the same destination from multiple routing protocols. When the Cisco Catalyst SD-WAN Controller or the router is selecting the OMP route to a destination, it prefers the one with the lowest administrative distance value.

The following table lists the default administrative distances used by the Cisco Catalyst SD-WAN devices:

Table 3:

Protocol	Administrative Distance
Connected	0
Static	1
NAT (NAT and static routes cannot coexist in the same VPN; NAT overwrites static routes)	1
Learned from DHCP	1
EIGRP Summary	5
EBGP	20
EIGRP	Internal: 90, External: 170
OSPF	110
OSPFv3	110
IS-IS	115
IBGP	200
OMP	250

OMP Best-Path Algorithm

Cisco Catalyst SD-WAN devices advertise their local paths to the Cisco Catalyst SD-WAN Controller using OMP. Depending on the network topology, some paths might be advertised from multiple devices. Cisco Catalyst SD-WAN devices use the following algorithm to choose the best path:

Table 4: Best Path Algorithm

Step	Applies to	Description
1	Edge devices Cisco Catalyst SD-WAN Controller	**Path validity** Check whether the OMP path is valid. If not, ignore it.
2	Edge devices Cisco Catalyst SD-WAN Controller	**Active vs. stale paths** Prefer an active path over a stale path. An active path is a one from a peer with which an OMP session is up. A stale path is one from a peer with which an OMP session is in Graceful Restart mode. **Note** A stale path will only be advertised if the stale version is similar to the Route Information Base (RIB) version. Otherwise, the stale path is dropped.
3	Edge devices	**Administrative distance** Select the OMP path with the lower administrative distance. Example: A path that the device learns locally via BGP would be preferred over a path that it learns from a Cisco SD-WAN Controller via OMP. For information about administrative distance, see Administrative Distance, on page 10.
4	Edge devices Cisco Catalyst SD-WAN Controller	**OMP path preference** Select the OMP path with the higher OMP path preference value.
5	Cisco Catalyst SD-WAN Controller	**Access region** Cisco SD-WAN Controller drops advertisement from border router (BR) to BR in the same region.
6	Edge devices	**Core region** Cisco SD-WAN Controller allows advertisement between BRs in the same access region, but receiving BR drops advertisement.
7	Multi-Region Fabric scenario only Edge devices	**Region path length** Compare region-path-length. Prefer lower. If **region-path-length-ignore** is configured, then skip this step. (This addresses secondary regions in Multi-Region Fabric.)

Step	Applies to	Description
8	Multi-Region Fabric scenario only Border routers	**Access region vs. core region** Prefer access region paths over core region paths.
9	Edge devices	**Direct vs. transport gateway path** Prefer a direct path over a transport gateway path. This step can be modified by the transport gateway path preference options, which can (a) cause the transport gateway path to be preferred, or (b) result in the paths to be considered equal. See Configure the Transport Gateway Path Preference in the *Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide.*
10	Multi-Region Fabric scenario only Edge devices	**Multi-Region Fabric subregion comparison** • Prefer paths from the router's own subregion. • When comparing two paths that are not from the router's subregion, prefer a path that is not part of any subregion.
11	Multi-Region Fabric scenario only Edge devices	**Border router preference** Prefer a path with a higher border router preference value.
12	Edge devices	**Derived affinity** Prefer a path with a lower derived affinity value.
13	Edge devices with an affinity preference configured	**Affinity preference** Depending on the affinity preference configured on the device, prefer a path whose affinity is earlier in the preference list (higher priority). If the device uses **affinity-preference-auto**, then prefer a path with a numerically lower affinity group. **Note** When comparing two paths with similar reorigination types, one with an affinity value and one without, prefer the path with an affinity value.
14	Edge devices	**TLOC preference** Select an OMP path with a higher TLOC preference value. **Note** With TLOC preference and AAR policy configured, outbound and inbound traffic may follow different paths when the preferred TLOC goes out of SLA. For outbound traffic, tunnels in SLA will be preferred regardless of TLOC preference; however, TLOC preference still dictates inbound path selection.

Step	Applies to	Description
15	Edge devices Cisco Catalyst SD-WAN Controller	**Origin type and subtype** Compare the origin type and subtype, and select the first match in the following list: • Connected • Static • EIGRP Summary • BGP External • EIGRP Internal • OSPF/OSPFv3 Intra-area • OSPF/OSPFv3 Inter-area • IS-IS Level 1 • EIGRP External • OSPF/OSPFv3 External (External OSPF Type1 is preferred over External OSPF Type2) • IS-IS Level 2 • BGP Internal • Unknown
16	Edge devices Cisco Catalyst SD-WAN Controller	**Origin metric** Select an OMP path that has a lower origin metric.
17	Cisco Catalyst SD-WAN Controller	**Path source** Prefer a path sourced from an edge router over the same path coming from a Cisco Catalyst SD-WAN Controller.
18	Edge devices Cisco Catalyst SD-WAN Controller	**Private IP address** If the router IDs are equal, a Cisco IOS XE Catalyst SD-WAN device selects the OMP path with the lower private IP address. If a Cisco Catalyst SD-WAN Controller receives the same prefix from two different sites and if all attributes are equal, it chooses both of them.



Note From all equal cost multi-paths for a given prefix selected as a best-paths and accepted by policy, advertise not more than number of paths specified in send-path-limit.

Here are some examples of choosing the best path:

- In a network with two WAN Edge devices and four Cisco SD-WAN Controllers, WAN Edge 1 by default forms control connections with two Cisco SD-WAN Controllers—Cisco SD-WAN Controller 1 and Cisco SD-WAN Controller 2, for example. If either Cisco SD-WAN Controller 1 or Cisco SD-WAN Controller 2 fails, WAN Edge 1 automatically establishes a new connection with one of the backup controllers, such as Cisco SD-WAN Controller 3 or Cisco SD-WAN Controller 4.

  For instance, if WAN Edge 1 loses its connection to Cisco SD-WAN Controller 1, it connects to Cisco SD-WAN Controller 3 while maintaining its session with Cisco SD-WAN Controller 2. After this failover, Cisco SD-WAN Controller 2 may stop advertising certain routes to WAN Edge 1. Specifically, if WAN Edge 2 originates prefix A and connects to Cisco SD-WAN Controller 2, Cisco SD-WAN Controller 2 may not advertise prefix A to WAN Edge 1, because it now learns the prefix from Cisco SD-WAN Controller 3 instead of directly from WAN Edge 2.

  As a result, WAN Edge 1 will only receive prefix A from Cisco SD-WAN Controller 3.

  This is the expected day 1 behavior: a Cisco SD-WAN Controller does not re-advertise a route learned from another Cisco SD-WAN Controller to a WAN Edge device if that device already receives the same route directly from a Cisco SD-WAN Controller.

  However, this behavior is not always deterministic. Sometimes, the route may still be advertised, especially after you run the **clear sdwan omp all command** or disable graceful restart (GR).

- A Cisco Catalyst SD-WAN Controller receives an OMP path to 10.10.10.0/24 via OMP from a Cisco vEdge device with an origin code of OSPF, and it also receives the same path from another Cisco Catalyst SD-WAN Controller, also with an origin code of OSPF. If all other things are equal, the best-path algorithm chooses the path that came from the Cisco vEdge device.

- A Cisco Catalyst SD-WAN Controller learns the same OMP path, 10.10.10.0/24, from two Cisco vEdge devices in the same site. If all other parameters are the same, both paths are chosen and advertised to other OMP peers. By default, up to four equal-cost paths are selected and advertised.

A Cisco vEdge device installs an OMP path in its forwarding table (FIB) only if the TLOC to which it points is active. For a TLOC to be active, an active BFD session must be associated with that TLOC. BFD sessions are established by each device which creates a separate BFD session with each of the remote TLOCs. If a BFD session becomes inactive, the Cisco Catalyst SD-WAN Controller removes from the forwarding table all the OMP paths that point to that TLOC.

OMP Graceful Restart

Graceful restart for OMP allows the data plane in the Cisco Catalyst SD-WAN overlay network to continue functioning if the control plane stops functioning or becomes unavailable. With graceful restart, if the Cisco SD-WAN Controller in the network goes down, or if multiple Cisco SD-WAN Controllers go down simultaneously, Cisco vEdge device can continue forwarding data traffic. They do this using the last known good information that they received from the Cisco SD-WAN Controller. When a Cisco SD-WAN Controller is again available, its DTLS connection to the device is re-established, and the device then receives updated, current network information from the Cisco SD-WAN Controller.

When OMP graceful restart is enabled, Cisco vEdge devices and a Cisco SD-WAN Controller (that is, two OMP peers) cache the OMP information that they learn from their peers. This information includes OMP routes, TLOC routes, service routes, IPsec SA parameters, and centralized data policies. When one of the OMP peers is no longer available, the other peer uses the cached information to continue operating in the network. So, for example, when a device no longer detects the presence of the OMP connection to a Cisco SD-WAN Controller, the device continues forwarding data traffic using the cached OMP information. The device also periodically checks whether the Cisco SD-WAN Controller has again become available. When it does come back up and the device re-establishes a connection to it, the device flushes its local cache and

considers only the new OMP information from the Cisco SD-WAN Controller to be valid and reliable. This same scenario occurs when a Cisco SD-WAN Controller no longer detects the presence of Cisco vEdge devices.



Note When a change to an OMP graceful restart configuration is made, the OMP session between the Cisco SD-WAN Controllers and the device is flapped. This causes all OMP routes belonging to different address families, such as TLOC, IPv4 or IPv6 unicast, IPv4 multicast, and other families to be withdrawn locally and relearned a few seconds later when the OMP session with the Cisco SD-WAN Controllers comes back up. As the TLOC routes are temporarily removed and added back, Bidirectional Forwarding Detection (BFD) sessions also flap momentarily. This is the expected behavior.

Each OMP peer independently sets up its graceful restart timer on Cisco IOS XE Catalyst SD-WAN Devices and Cisco SD-WAN Controllers. For example, consider a Cisco SD-WAN Controller with a graceful restart time of 300 seconds (5 minutes) and a Cisco IOS XE Catalyst SD-WAN Device configured with a timer of 600 seconds (10 minutes). In this scenario, the Cisco SD-WAN Controller retains the OMP routes learned from the Cisco IOS XE Catalyst SD-WAN Device for 10 minutes, reflecting the its configured timer sent during the OMP session setup. Conversely, the device retains the routes from the Cisco SD-WAN Controller for 5 minutes, which is the default timer value set on the controller and communicated during the session setup.

BGP and OSPF Routing Protocols

The Cisco Catalyst SD-WAN overlay network supports BGP and OSPF unicast routing protocols. These protocols can be configured on Cisco vEdge device in any VPN except for VPN 0 and VPN 512 to provide reachability to networks at their local sites. Cisco vEdge devices can redistribute route information learned from BGP and OSPF into OMP so that OMP can better choose paths within the overlay network.

When the local site connects to a Layer 3 VPN MPLS WAN cloud, the devices act as an MPLS CE devices and establish a BGP peering session to connect to the PE router in the L3VPN MPLS cloud.

When the devices at a local site do not connect directly to the WAN cloud but are one or more hops from the WAN and connect indirectly through a non-Cisco SD-WAN device, standard routing must be enabled on the devices' DTLS connections so that they can reach the WAN cloud. Either OSPF or BGP can be the routing protocol.

In both these topologies, the BGP or OSPF sessions run over a DTLS connection created on the loopback interface in VPN 0, which is the transport VPN that is responsible for carrying control traffic in the overlay network. The Cisco Catalyst SD-WAN Validator learns about this DTLS connection via the loop-back interface and conveys this information to the Cisco Catalyst SD-WAN Controller so that it can track the TLOC-related information. In VPN 0, you also configure the physical interface that connects the Cisco vEdge device to its neighbor—either the PE router in the MPLS case or the hub or next-hop router in the local site—but you do not establish a DTLS tunnel connection on that physical interface.

Configure Unicast Overlay Routing

This topic describes how to provision unicast overlay routing.

Service-Side Routing

Provisioning BGP and OSPF enables routing on the service side of the network.

To set up routing on a Cisco vEdge 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.

Because Cisco Catalyst SD-WAN Controllers never participate in a local site network, you never configure BGP or OSPF on these devices.

Transport-Side Routing

To enable communication between Cisco SD-WAN devices, you configure OSPF or BGP on a loopback interface in VPN 0. The loopback interface is a virtual transport interface that is the terminus of the DTLS and IPsec tunnel connections required for Cisco vEdge devices to participate in the overlay network.

To configure service-side and transport-side BGP using Cisco SD-WAN Manager, see *Configure BGP*. To configure service-side and transport-side BGP using the CLI, see the *Configure BGP Using CLI* topic.

Configure BGP

The Border Gateway Protocol (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 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:

1. Create a BGP feature template to configure BGP parameters.
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

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



Note In Cisco vManage Release 20.7.x and earlier releases, **Device Template** is titled **Device**.

3. Click **Create Template**
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you are creating the template.

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.
7. Create a template for VPNs **1** through **511**, and **513** through **65525**:
   a. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
   b. Click the **Service VPN** drop-down list.
   c. Under **Additional VPN Templates**, click **BGP**.
   d. 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.
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.
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

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
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

Parameter Name	Description
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

For service-side BGP, you might want to configure Overlay Management Protocol (OMP) to advertise to the Cisco Catalyst SD-WAN Controller any BGP routes that the device learns. By default, Cisco 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 or Cisco SD-WAN software.

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.

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

Configure Unicast Address Family

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: • 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 *prefix/length* 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 *prefix/length*.
	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.		

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

To change the AS number, perform the following steps:

1. Remove the BGP configuration. Wait for few seconds.
2. Configure the BGP again with changed global-as and the local-as configuration.

Configure BGP Neighbors

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

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

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.*Range:* 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.		

Configure Advanced Neighbor Parameter

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 **ge** *port*/*slot*.

Parameter Name	Description
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

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

Change the Scope of a Parameter Value

When you first open a feature template, for each parameter that has a default value, the scope is set to Default (a ), and the default setting or value is shown). 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	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. 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. To change the default key, type a new string and move the cursor out of the **Enter Key** box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
 Global	Enter a value for the parameter, and apply that value to all devices. Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.

Configure Advanced BGP Parameters

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.

Parameter Name	Description
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.

To save the feature, click **Save**.

Configure BGP Routing in a Service 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 a BGP Routing feature in a Service 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

Field	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: 1 through 255 Default: 20

b. Configure Unicast Address Family fields.

Table 6: 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
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	

Field	Description
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
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	

Field	Description
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 Neighor fields.

Table 7: 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.

Field	Description
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*	Choose one of the following options: • **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: • **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 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.

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*	Choose one of the following options: • **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: • **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.

What to do next

Also see Deploy a configuration group.

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 8: 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 9: 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
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	

Field	Description
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.
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.

Field	Description
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.
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 MPLS Interface.

Table 10: MPLS Interface

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

d. Configure Neighbor.

Table 11: 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.
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.
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.

Field	Description
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*	Choose one of the following options: • **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: • **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 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.

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*	Choose one of the following options: • **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: • **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 12: Advanced

Field	Description
Keepalive (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. This keepalive time is 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. This hold time is the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Compare MED	Enable this option to compare the router IDs among BGP paths to determine the active path.

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.

What to do next

Also see Deploy a configuration group.

Configure BGP Using CLI

The following section describes how to configure BGP for service-side and transport-side for unicast overlay routing:

Configure Service-Side Routing

To set up routing on the Cisco vEdge 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.

1. Configure a VPN.

   ```
   Device(config)# vpn vpn-id
   ```

   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.

2. Configure BGP to run in the VPN:

   a. Configure the local AS number:

      ```
      Device(config-vpn)# router bgp 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).

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

      ```
      Device(config-bgp)# neighbor address remote-as remote-as-number
      Device(config-bgp)# no shutdown
      ```

3. Configure a system IP address for the Cisco vEdge device:

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

Example of BGP Configuration on a vEdge Router

```
Device# show running-config system
system
  system-ip 10.1.2.3
!
Device# show running-config vpn 1
vpn 1
  router
    bgp 1
      neighbor 11.1.2.3
        no shutdown
        remote-as 2
       !
    !
  !
  ip route 0.0.0.0/0 10.0.16.13
!
```

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 Catalyst SD-WAN devices, then advertises the OMP routes to all the BGP routers in the service-side of the network.

```
Device(config)# vpn vpn-id router bgp
vEdge(config-bgp)# address-family ipv4-unicast redistribute omp [route-policy policy-name]
```

You can also redistribute routes learned from other protocols into BGP:

```
Device(config-bgp)# address-family ipv4-unicast redistribute (connected | nat |
natpool-outside | ospf | static) [route-policy policy-name]
```

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

```
vEdge(config-bgp)# neighbor ip-address
vEdge(config-neighbor)# address-family ipv4-unicast redistribute (connected | nat |
natpool-outside | omp | ospf | static)
vEdge(config-neighbor)# route-policy policy-name (in | out)
```

In the BGP route redistribution commands, the optional route policy is applied to the routes that are redistributed into BGP or routes that are redistributed out from BGP.

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

```
vEdge(config)# omp advertise bgp

vEdge(config)# vpn vpn-id omp advertise bgp
```

BGP Route Advertisements

By default, when BGP advertises routes into OMP, BGP advertises each prefix's metric. BGP can also advertise the prefix's AS path:

```
Device(config)# vpn vpn-id router bgp
vEdge(config-bgp)# propagate-aspath
```

When you configure BGP to propagate AS path information, the router sends AS path information to routers that are behind the vEdge router (in the service-side network) that are running BGP, and it receives AS path information from these routers. If you are redistributing BGP routes into OMP or into another protocol, or if you are advertising BGP routes to OMP, the AS path information is included in the advertised BGP routes. If you configure BGP AS path propagation on some but not all vEdge routers in the overlay network, the routers on which it is not configured receive the AS path information but they do not forward it to the BGP routers in their local service-side network. Propagating AS path information can help to avoid BGP routing loops.

In networks that have both overlay and underlay connectivity—for example, when vEdge routers are interconnected by both a Cisco SD-WAN overlay network and an MPLS underlay network—you can assign an AS number to OMP itself. For vEdge routers running BGP, this overlay AS number is included in the AS path of BGP route updates. To configure the overlay AS:

```
Device(config)# omp
Device(omp)# overlay-as 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). As a best practice, it is recommended that the overlay AS number be a unique AS number within both the overlay and the underlay networks. That use, select an AS number that is not used elsewhere in the network.

If you configure the same overlay AS number on multiple vEdge routers in the overlay network, all these routers are considered to be part of the same AS, and as a result, they do not forward any routes that contain the overlay AS number. This mechanism is an additional technique for preventing BGP routing loops in the network.

Configure Transport-Side Routing

To configure transport-side routing, you configure a loopback interface, the physical interface, and the routing protocol in VPN 0.

1. Configure a physical interface in VPN 0:

   ```
   Device(config)# vpn 0 interface geslot/port ip address address
   vedge(config-interface)# no shutdown
   ```

2. Configure a loopback interface in VPN 0:

   ```
   Device(config)# vpn 0 interface loopbacknumber ip address address
   Device(config-interface)# no shutdown
   Device(config-interface)# tunnel-interface color color
   ```

3. Configure a BGP instance in VPN 0:

   ```
   Device(config)# vpn 0 router bgp local-as-number
   ```

4. Create a policy for BGP to advertise the loopback interface address to its neighbors:

   ```
   vEdge(config)# policy lists prefix-list prefix-list-name ip-prefix prefix
   prefix is the IP address of the loopback interface.
   ```

   prefix is the IP address of the loopback interface.

5. Configure a route policy that affects the loopback interface's prefix:

   ```
   Device(config)# policy route-policy policy-name sequence number match address
   prefix-list-name
   Device(config)# policy route-policy policy-name sequence number action accept
   Device(config)# policy route-policy policy-name default-action reject
   ```

6. Reference the policy in the BGP instance. To apply the policy such that the loopback address is advertised to all BGP neighbors:

   ```
   Device(config)# vpn 0 router bgp local-as-number address-family ipv4-unicast redistribute
    connected route-policy policy-name
   ```

   To apply the policy only to a specific neighbor:

   ```
   Device(config)# vpn 0 router bgp local-as-number neighbor neighbor-address address-family
    ipv4-unicast redistribute connected route-policy policy-name out
   ```

   Specify **out** in the second command so that BGP advertises the loopback prefix out to the neighbor.

Example of BGP Transport-Side Configuration

Here is an example of a minimal BGP transport-side routing configuration in which the loopback address is advertised to all the vEdge router's BGP neighbors. Note that even though we did not configure any services on the tunnel interface, these services are associated with the tunnel by default and are included in the configuration. Because services affect only physical interfaces, you can ignore them on loopback interfaces.

```
vEdge# show running-config vpn 0
vpn 0
 router
  bgp 2
   router-id 172.16.255.18
   timers
    keepalive 1
    holdtime  3
   !
   address-family ipv4-unicast
    redistribute connected route-policy export_loopback
   !
   neighbor 10.20.25.16
    no shutdown
    remote-as 1
    timers
     connect-retry          2
     advertisement-interval 1
    !
   !
  !
 !
 interface ge0/1
  ip address 10.20.25.18/24
  no shutdown
 !
 interface loopback
  ip address 172.16.255.118/32
  tunnel-interface
   color lte
   allow-service dhcp
   allow-service dns
   allow-service icmp
   no allow-service sshd
   no allow-service ntp
   no allow-service stun
  !
  no shutdown
 !
!
policy
 lists
  prefix-list loopback_prefix
```

```
  ip-prefix 172.16.255.118/32
 !
!
route-policy export_loopback
 sequence 10
  match
   address loopback_prefix
  !
  action accept
  !
 !
 default-action reject
!
!
```

Configure OSPFv3 IPv4 Routing 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 a Transport VPN feature in a Transport and Management profile for VPN 0 or VPN 512. Create and configure a Service VPN feature in a Service profile for VPNs 1 through 511, and 513 through 65530.

Step 3 Create and configure an OSPFv3 IPV4 routing feature.

a. Configure Basic Configuration.

Table 13: Basic Configuration

Field	Description
Router ID	Enter the OSPFv3 IPv4 router ID, in decimal four-part dotted notation. This is the IP address associated with the router for OSPF adjacencies.
Reference Bandwidth (Mbps)	Specify the reference bandwidth for the OSPF auto-cost calculation for the interface. Range: 1 through 4294967 Mbps Default: 100 Mbps
RFC 1583 Compatible	By default, the OSPF calculation is done per RFC 1583. Disable this option to calculate the cost of summary routes based on RFC 2328.

Field	Description
Originate	Enable this option to generate a default external route into an OSPF routing domain. When you enable this option, the following fields appear: • **Always**: Enable this option to always advertise the default route in an OSPF routing domain. • **Default Metric**: Set the metric used to generate the default route. Range: 0 through 16777214 Default: 10 • **Metric Type**: Choose to advertise the default route as an OSPF Type 1 external route or an OSPF Type 2 external route.
SPF Calculation Delay (milliseconds)	Specify the amount of time between when the first change to a topology is received until performing the SPF calculation. Range: 1 through 600000 milliseconds (60 seconds) Default: 200 milliseconds
Initial Hold Time (milliseconds)	Specify the amount of time between consecutive SPF calculations. Range: 1 through 600000 milliseconds (60 seconds) Default: 1000 milliseconds
Maximum Hold Time (milliseconds)	Specify the longest time between consecutive SPF calculations. Range: 1 through 600000 Default: 10000 milliseconds (60 seconds)
Distance for External Routes	Specify the OSPF route administration distance for routes learned from other domains. Range: 1 through 255 Default: 110
Distance for Inter-Area Routes	Specify the OSPF route administration distance for routes coming from one area into another. Range: 1 through 255 Default: 110
Distance for Intra-Area Routes	Specify the OSPF route administration distance for routes within an area. Range: 0 through 255 Default: 110

b. Configure Interface.

Table 14: Interface Settings

Field	Description
Add Interface	Configure the properties of an interface. Configure the area range of an interface in an OSPFv3 area.
Name	Enter the name of the interface, in the format **`geslot/port`** or **`loopback number`**.
Cost (optional)	Specify a number for the Type 3 summary LSA. OSPF uses this metric during its SPF calculation to determine the shortest path to a destination. Range: 0 through 16777214
Authentication Type (optional)	Specify the SPI and authentication key if you use IPSec SHA1 authentication type. • **no-auth**: Select no authentication. • **ipsec-sha1**: Enter the value for the IPSEC Secure Hash Algorithm 1 (SHA-1) authentication.
SPI (optional)	Specifies the Security Policy Index (SPI) value. Range: 256 through 4294967295
Authentication Key (optional)	Provide a value for the authentication key. When IPSEC SHA-1 authentication is used, the key must be 40 hex digits long.
Passive Interface (optional)	Specify whether to set the OSPFv3 interface to be passive. A passive interface advertises its address, but does not actively run the OSPFv3 protocol. Default: Disabled
Add Range	Configure the area range of an interface in an OSPF area.
IP Address*	Enter the IP address.
Subnet Mask*	Enter the subnet mask.
Cost	Specify a number for the Type 3 summary LSA. OSPF uses this metric during its SPF calculation to determine the shortest path to a destination. Range: 0 through 16777214
No-advertise*	Enable this option to not advertise the Type 3 summary LSAs.

c. Configure Redistribute.

d. Configure Area Parameters.

Table 15: Area

Field	Description
Add Area	

Field	Description
Area Number*	Enter the number of the OSPF area. Range: 32-bit number
Set the area type	Choose the type of OSPF area: • Stub • NSSA

What to do next

Also see Deploy a configuration group.

Configure OSPF

Use the OSPF template for all Cisco Catalyst SD-WAN devices.

To configure OSPF on a device using Cisco SD-WAN Manager templates:

1. Create an OSPF feature template to configure OSPF parameters. OSPF 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 the Cisco Catalyst SD-WAN devices when the router is not directly connected to the WAN cloud. Create separate OSPF templates for the two OSPF routing types.
2. Create a VPN feature template to configure VPN parameters for either service-side OSPF routing (in any VPN other than VPN 0 or VPN 512) or transport-side OSPF routing (in VPN 0). See the VPN help topic for more information.

Create an OSPF Template

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

Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, select the type of device for which you are creating the template. 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 **OSPF**.

c. From the **OSPF** drop-down list, click **Create Template**. The OSPF template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OSPF parameters.

6. Create a template for VPNs in this Range: 1 to 65525, excluding 512. For details see the VRF range behavior change described here.
   a. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
   b. Click the **Service VPN** drop-down list.
   c. Under **Additional VPN Templates**, click **OSPF**.
   d. From the **OSPF** drop-down list, click **Create Template**. The OSPF template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OSPF parameters.
7. 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.
8. 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.

When you first open a feature template, for each parameter that has a default value, the scope is set to **Default** (indicated by a check mark), and the default setting or value is shown. To change the default or to enter a value, click the scope drop-down list to the left of the parameter field and choose one of the following:

Table 16:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	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 Cisco SD-WAN device to a device template . 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 Cisco SD-WAN device to a device template. For more information, see *Create a Template Variables Spreadsheet*. To change the default key, type a new string and move the cursor out of the **Enter Key** box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
Global (indicated by a globe icon)	Enter a value for the parameter, and apply that value to all devices. Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.

Configure Basic OSPF

To configure basic OSPF, select **Basic Configuration** and then configure the following parameters. All these parameters are optional. For OSPF to function, you must configure area 0, as described below.

Table 17:

Parameter Name	Description
Router ID	Enter the OSPF router ID in decimal four-part dotted notation. This is the unique 32-bit identifier associated with the OSPF router for Link-State Advertisements (LSAs) and adjacencies.
Distance for External Routes	Specify the OSPF route administration distance for routes learned from other domains. *Range:* 0 through 255 *Default:* 110
Distance for Inter-Area Routes	Specify the OSPF route administration distance for routes coming from one area into another. *Range:* 0 through 255 *Default:* 110
Distance for Intra-Area Routes	Specify the OSPF route administration distance for routes within an area. *Range:* 0 through 255 *Default:* 110

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

Redistribute Routes into OSPF

To redistribute routes learned from other protocols into OSPF on Cisco SD-WAN devices, choose **Redistribute** > **Add New Redistribute** and configure the following parameters:

Table 18:

Parameter Name	Description
Protocol	Choose the protocol from which to redistribute routes into OSPF. Choose from BGP, Connected, NAT, OMP, EIGRP and Static.
Route Policy	Enter the name of a localized control policy to apply to routes before they are redistributed into OSPF.

To add another OSPF route redistribution policy, click the plus sign (+).

To remove an OSPF route redistribution policy from the template configuration, click **the trash icon** to the right of the entry.

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

Configure OSPF To Advertise a Maximum Metric

To configure OSPF to advertise a maximum metric so that other devices do not prefer the Cisco vEdge device as an intermediate hop in their Shortest Path First (SPF) calculation, choose **Maximum Metric (Router LSA)** > **Add New Router LSA** and configure the following parameters:

Table 19:

Parameter Name	Description
Type	Choose a type: • **Administrative**—Force the maximum metric to take effect immediately through operator intervention. • **On-Startup**—Advertise the maximum metric for the specified time.
Advertisement Time	If you selected **On-Startup**, specify the number of seconds to advertise the maximum metric after the router starts up. *Range:* 0, 5 through 86400 seconds *Default:* 0 seconds (the maximum metric is advertised immediately when the router starts up)

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

Configure OSPF Areas

To configure an OSPF area within a VPN on a Cisco SD-WAN device, choose **Area** > **Add New Area**. For OSPF to function, you must configure area 0.

Table 20:

Parameter Name	Description
Area Number	Enter the number of the OSPF area. *Range:* 32-bit number
Set the Area Type	Choose the type of OSPF area, Stub or NSSA.
No Summary	Click **On** to not inject OSPF summary routes into the area.
Translate	If you configured the area type as NSSA, choose when to allow Cisco Catalyst SD-WAN devices that are ABRs (area border routers) to translate Type 7 LSAs to Type 5 LSAs: • **Always**—Router always acts as the translator for Type 7 LSAs. That is no other router, even if it is an ABR, can be the translator. If two ABRs are configured to always be the translator, only one of them actually ends up doing the translation. • **Candidate**—Router offers translation services, but does not insist on being the translator. • **Never**—Translate no Type 7 LSAs.

To save the new area, click **Add**.

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

Configure Interfaces in an OSPF Area

To configure the properties of an interface in an OSPF area, choose **Area** > **Add New Area** > **Add Interface**. In the **Add Interface** popup, configure the following parameters:

Table 21:

Parameter Name	Description
Interface Name	Enter the name of the interface, in the format **ge** *slot*/*port* or **loopback** *number*.
Hello Interval	Specify how often the router sends OSPF hello packets. *Range:* 1 through 65535 seconds *Default:* 10 seconds
Dead Interval	Specify how often the Cisco vEdge device must receive an OSPF hello packet from its neighbor. If no packet is received, the Cisco vEdge deviceassumes that the neighbor is down. *Range:* 1 through 65535 seconds *Default:* 40 seconds (4 times the default hello interval)
LSA Retransmission Interval	Specify how often the OSPF protocol retransmits LSAs to its neighbors. *Range:* 1 through 65535 seconds *Default:* 5 seconds
Interface Cost	Specify the cost of the OSPF interface. *Range:* 1 through 65535

To configure advanced options for an interface in an OSPF area, in the **Add Interface** popup, click **Advanced Options** and configure the following parameters:

Table 22:

Parameter Name	Description
Designated Router Priority	Set the priority of the router to be elected as the designated router (DR). The router with the highest priority becomes the DR. If the priorities are equal, the node with the highest router ID becomes the DR or the backup DR. *Range:* 0 through 255 *Default:* 1
OSPF Network Type	Choose the OSPF network type to which the interface is to connect: • Broadcast network—WAN or similar network. • Point-to-point network—Interface connects to a single remote OSPF router. • Non-broadcast—Point-to-multipoint. *Default:* Broadcast
Passive Interface	Click **On** or **Off** to specify whether to set the OSPF interface to be passive. A passive interface advertises its address, but does not actively run the OSPF protocol. *Default:* Off

Parameter Name	Description
Authentication	Specify the authentication and authentication key on the interface to allow OSPF to exchange routing update information securely.
• **Authentication Type**	Choose the authentication type: • Simple authentication—Password is sent in clear text. • Message-digest authentication—MD5 algorithm generates the password.
• **Authentication Key**	Enter the authentication key. Plain text authentication is used when devices within an area cannot support the more secure MD5 authentication. The key can be 1 to 32 characters.
• **Message Digest**	Specify the key ID and authentication key if you are using message digest (MD5).
• **Message Digest Key ID**	Enter the key ID for message digest (MD5 authentication). It can be 1 to 32 characters.
• **Message Digest Key**	Enter the MD5 authentication key in clear text or as an AES-encrypted key. It can be from 1 to 255 characters.

To save the interface configuration, click **Save**.

To save the new area, click **Add**.

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

Configure an Interface Range for Summary LSAs

To configure the properties of an interface in an OSPF area, choose **Area** > **Add New Area** > **Add Range**. In the **Area Range** popup, click **Add Area Range**, and configure the following parameters:

Table 23:

Parameter Name	Description
Address	Enter the IP address and subnet mask, in the format *prefix*/*length* for the IP addresses to be consolidated and advertised.
Cost	Specify a number for the Type 3 summary LSA. OSPF uses this metric during its SPF calculation to determine the shortest path to a destination. *Range:* 0 through 16777215
No Advertise	Click **On** to not advertise the Type 3 summary LSAs or **Off** to advertise them.

To save the area range, click **Save**.

To save the new area, click **Add**.

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

Configure Other OSPF Properties

To configure other OSPF properties, click **Advanced** and configure the following properties:

Table 24:

Parameter Name	Description
Reference Bandwidth	Specify the reference bandwidth for the OSPF auto-cost calculation for the interface. *Range:* 1 through 4294967 Mbps *Default:* 100 Mbps
RFC 1538 Compatible	By default, the OSPF calculation is done per RFC 1583. Click **Off** to calculate the cost of summary routes based on RFC 2328.
Originate	Click **On** to generate a default external route into an OSPF routing domain: • Always—Click **On** to always advertise the default route in an OSPF routing domain. • Default metric—Set the metric used to generate the default route. *Range:* 0 through 16777214 *Default:* 10 • Metric type—Select to advertise the default route as an OSPF Type 1 external route or an OSPF Type 2 external route.
SPF Calculation Delay	Specify the amount of time between when the first change to a topology is received until performing the SPF calculation. *Range*: 0 through 600000 milliseconds (60 seconds) *Default*: 200 milliseconds
Initial Hold Time	Specify the amount of time between consecutive SPF calculations. *Range*: 0 through 600000 milliseconds (60 seconds) *Default*: 1000 milliseconds
Maximum Hold Time	Specify the longest time between consecutive SPF calculations. *Range*: 0 through 600000 *Default*: 10000 milliseconds (60 seconds)
Policy Name	Enter the name of a localized control policy to apply to routes coming from OSPF neighbors.

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

Configure OSPF Using CLI

This topic describes how to configure basic service-side and transport-side OSPF for Unicast overlay routing.

Configure Basic Service-Side OSPF

To set up routing on the Cisco vEdge 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.

To configure basic service-side OSPF functionality:

1. Configure a VPN for the OSPF network:

   ```
   vEdge(config)# vpn vpn-id
   ```

   vpn-id can be any VPN number except VPN 0 and VPN512. VPN 0 is the transport VPN and carries only control traffic, and VPN 512 is the management interface.

2. Configure OSPF area 0 and the interfaces that participate in that area:

   ```
   vEdge(config-vpn)# router ospf
   vEdge(config-ospf)# area 0
   vEdge(config-area-0)# interface  interface-name
   vEdge(config-interface)# ip-address  address
   vEdge(config-interface)# no shutdown
   vEdge (ospf-if)#  exit
   ```

3. Redistribute OMP routes into OSPF:

   ```
   vEdge(config-ospf)# redistribute omp
   ```

   By default, OMP routes are not redistributed into OSPF.

4. Repeat Steps 1 through 3 for any additional VPNs.

5. If desired, configure OMP to advertise to the Cisco Catalyst SD-WAN Controller any BGP and OSPF external routes that the Cisco vEdge device has learned:

   ```
   vEdge(config)# omp
   vEdge(config-omp)# advertise bgp
   vEdge(config-omp)# advertise ospf external
   ```

Example of Basic Service-Side OSPF Configuration

This configuration sets up VPN 10 with two interfaces, **ge2/0** and **ge3/0**. It enables OSPF routing on those interfaces in area 0, and it redistributes the OMP routes from the Cisco Catalyst SD-WAN Controller into OSPF.

```
vpn 10
  router
    ospf
      redistribute omp
      area 0
        interface ge2/0
        exit
      interface ge3/0
       exit
     exit
    !
  !
  interface ge2/0
    ip address 10.0.5.12/24
    no shutdown
  !
  interface ge3/0
    ip address 10.0.2.12/24
    no shutdown
  !
```

Configure OSPF Transport-Side Routing

To configure transport-side routing, you configure a loopback interface, the physical interface, and the routing protocol in VPN 0.

To configure OSPF transport-side routing:

1. Configure a physical interface in VPN 0:

```
vEdge(config)# vpn 0 interface geslot/port ip address address
vEdge(config-interface)# no shutdown
```

2. Configure a loopback interface in VPN 0 as a tunnel interface:

```
vEdge(config)# vpn 0 interface loopbacknumber ip address address
vEdge(config-interface)# no shutdown
vEdge(config-interface)# tunnel-interface color color
```

3. Configure an OSPF instance in VPN 0:

```
vEdge(config)# vpn 0 router ospf
```

4. Add the physical and loopback interfaces to the OSPF area:

```
vEdge(config-ospf)# area number interface geslot/port
vEdge(config-area)# interface loopbacknumber
```

Example of Transport-Side OSPF Configuration

Here is any example of a minimal OSPF transport-side routing configuration. Note that even though we did not configure any services on the tunnel interface, these services are associated with the tunnel by default and are included in the configuration. Because services affect only physical interfaces, you can ignore them on loopback interfaces.

```
vEdge# show running-config vpn 0
vpn 0
 router
  ospf
   router-id 172.16.255.11
   timers spf 200 1000 10000
   area 0
    interface ge0/1
    exit
    interface loopback1
    exit
   exit
  !
 !
 interface ge0/1
  ip address 10.0.26.11/24
  no shutdown
 !
 interface loopback1
  ip address 10.0.101.1/32
  tunnel-interface
   color lte
   allow-service dhcp
   allow-service dns
   allow-service icmp
   no allow-service sshd
   no allow-service ntp
   no allow-service stun
  !
  no shutdown
 !
!
```

Configure OMP 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 a OMP feature in a System profile.

a. Configure Basic Configuration.

Table 25: Basic Configuration

Field	Description
Graceful Restart Enable	Enable graceful restart. By default, the graceful restart for OMP is enabled.
Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. A advertises routes to Cisco Catalyst SD-WAN Controllers, and the controllers redistribute the learned routes, advertising each route-TLOC tuple. A can have up to four TLOCs, and by default advertises each route-TLOC tuple to the Cisco Catalyst SD-WAN Controller. If a local site has two s, a Cisco Catalyst SD-WAN Controller could potentially learn eight route-TLOC tuples for the same route. If the configured limit is lower than the number of route-TLOC tuples, the best route or routes are advertised. Range: 1 through 16 Default: 4
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco Catalyst SD-WAN Controller that can be installed in the local route table of the Cisco IOS XE Catalyst SD-WAN device. By default, a installs a maximum of four unique OMP paths into its route table. Range: 1 through 16 Default: 4
Advertisement Interval (In Second)	Specify the time between OMP update packets. Range: 0 through 65535 seconds Default: 1 second We recommend you to configure 5 seconds on edge devices and 20 seconds on vSmart.

Field	Description
Hold Time(In Second)	Specify how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed. Range: 0 through 65535 seconds Defaults, by Cisco Catalyst SD-WAN Control Components release: • 20.18.x and later: 300 seconds • 20.16.x: 5400 seconds • 20.12.1 to 20.15.x: 300 seconds • Before 20.12.1: 60 seconds Defaults, by Cisco IOS XE Catalyst SD-WAN release: • 17.18.1 and later: 300 seconds • 17.16.x: 5400 seconds
EOR Timer(In Second)	Specify how long to wait after an OMP session has gone down and then come back up to send an end-of-RIB (EOR) marker. After this marker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. Range: 1 through 3600 seconds (1 hour) Default: 300 seconds (5 minutes)
Overlay AS	Specify a BGP AS number that OMP advertises to the BGP neighbors of the router.
Shutdown	Enable this option to disable OMP and disable the Cisco Catalyst SD-WAN overlay network. OMP is enabled by default.
OMP Admin Distance Ipv4	To advertise a route over OMP, configure the OMP administrative distance for the IPv4 address lower than the leaked route administrative distance.
OMP Admin Distance Ipv6	To advertise a route over OMP, configure the OMP administrative distance for the IPv6 address lower than the leaked route administrative distance.

b. Configure Timers.

Table 26: Timers

Field	Description
Graceful Restart(In Second)	Specify how often the OMP information cache is flushed and refreshed. A timer value of 0 disables OMP graceful restart. Range: 0 through 604800 seconds (168 hours, or 7 days) Default: 43200 seconds (12 hours)

c. Configure Advertise.

Table 27: Advertise

Field	Description
Advertise Ipv4 BGP	Enable this option to advertise BGP routes to OMP. By default, BGP routes are not advertised to OMP.
Advertise Ipv4 OSPF	Enable this option to advertise external OSPF routes to OMP. By default, external OSPF routes are not advertised to OMP.
Advertise Ipv4 OSPF v3	Enable this option to advertise external OSPFv3 routes to OMP. By default, external OSPFv3 routes are not advertised to OMP.
Advertise Ipv4 Connected	Enable this option to advertise connected routes to OMP. By default, connected routes are not advertised to OMP.
Advertise Ipv4 Static	Enable this option to advertise static routes to OMP. By default static routes are not advertised to OMP.
Advertise Ipv4 LISP	Enable this option to advertise LISP routes to OMP. By default, LISP routes are not advertised to OMP.
Advertise Ipv4 ISIS	Enable this option to advertise IS-IS routes to OMP. By default, IS-IS routes are not advertised to OMP.
Advertise Ipv4 EIGRP	Enable this option to advertise EIGRP routes to OMP. By default, EIGRP routes are not advertised to OMP.
Advertise Ipv6 BGP	Enable this option to advertise BGP routes to OMP. By default, BGP routes are not advertised to OMP.
Advertise Ipv6 OSPF	Enable this option to advertise external OSPF routes to OMP. By default, external OSPF routes are not advertised to OMP.
Advertise Ipv6 Connected	Enable this option to advertise connected routes to OMP. By default, connected routes are not advertised to OMP.
Advertise Ipv6 Static	Enable this option to advertise static routes to OMP. By default static routes are not advertised to OMP.
Advertise Ipv6 LISP	Enable this option to advertise LISP routes to OMP. By default, LISP routes are not advertised to OMP.
Advertise Ipv6 ISIS	Enable this option to advertise IS-IS routes to OMP. By default, IS-IS routes are not advertised to OMP.
Advertise Ipv6 EIGRP	Enable this option to advertise EIGRP routes to OMP. By default, EIGRP routes are not advertised to OMP.

d. Configure Best Path.

Table 28: Best Path

Field	Description
Treat Hierarchical and Direct Paths Equally	(Minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.13.1) In a Multi-Region Fabric scenario, if using secondary regions, enable this option to enable packets to use all available paths rather than only direct paths. By default, when a direct path is available to reach a destination, the overlay management protocol (OMP) enables only the direct path to the routing forwarding layer because the direct path uses fewer hops. This logic is part of route optimization. The result is that the forwarding layer, which includes application-aware routing policy, can only use the direct path. **Treat Hierarchical and Direct Paths Equally** disables this comparison of the number of hops so that traffic can use either the direct secondary-region path (fewer hops) or the primary-region path (more hops). When you disable the comparison of the number of hops, OMP applies equal-cost multi-path routing (ECMP) to all routes, and packets can use all available paths.
Transport Gateway Path Behavior	(Minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.13.1) Choose one of the following: • **Prefer Transport Gateway Path**: For devices that can connect through a transport gateway, use only the transport gateway paths, even if other paths are available. • **Do ECMP Between Direct and Transport Gateway Paths**: For devices that can connect through a transport gateway and through direct paths, apply ECMP to all available paths.
Site Type	(Minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.13.1) If you configure a value for **Transport Gateway Path Behavior**, this field appears. Optionally, choose one or more site types to apply the transport gateway path behavior only to those site types.

What to do next

Also see Deploy a configuration group.

Configure OMP using templates

Use the OMP template to configure OMP parameters for all Cisco vEdge devices, and for Cisco Catalyst SD-WAN Controllers.

OMP is enabled by default on all Cisco vEdge devices, Cisco SD-WAN Manager NMSs, and Cisco Catalyst SD-WAN Controllers, so there is no need to explicitly enable OMP. OMP must be operational for the Cisco SD-WAN overlay network to function. If you disable it, you disable the overlay network.

Note

- Route advertisements in OMP are done either by applying the configuration at the global level or at the specific VPN level. For more information about route advertisements in OMP, see the *Configure OMP Advertisements* section in this topic.

Create OMP Template

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

Note In Cisco vManage Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you're creating the template.
6. To create a custom template for OMP, choose the **Factory_Default_OMP_Template** and click **Create Template**. The OMP template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OMP parameters. You may need to click an operation or the plus sign (+) to display more fields.
7. 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.
8. 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.

When you first open a feature template, for each parameter that has a default value, the scope is set to **Default** (indicated by a check mark), and the default setting or value is shown. 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:

Table 29:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	Use a device-specific value for the parameter. For device-specific parameters, you can't enter a value in the feature template. You enter the value when you attach a Cisco SD-WAN device to a device template. 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 Cisco SD-WAN device to a device template. To change the default key, type a new string and move the cursor out of the **Enter Key** box. Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.
Global (indicated by a globe icon)	Enter a value for the parameter, and apply that value to all devices. Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.

Configure Basic OMP Options

To configure basic OMP options, click **Basic Configuration** and configure the following parameters. All parameters are optional.

Table 30:

Parameter Name	Description
Graceful Restart for OMP	Ensure that Yes is selected to enable graceful restart. By default, graceful restart for OMP is enabled.
Overlay AS Number	Specify a BGP AS number that OMP advertises to the router's BGP neighbors.
Graceful Restart Timer	Specify how often the OMP information cache is flushed and refreshed. A timer value of 0 disables OMP graceful restart. Range: 0 to 604800 seconds (168 hours, or 7 days) Default: 43200 seconds (12 hours)

Parameter Name	Description
Number of Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. Cisco vEdge devices advertise routes to Cisco Catalyst SD-WAN Controllers, and the controllers redistributes the learned routes, advertising each route-TLOC tuple. A Cisco vEdge device can have up to eight TLOCs, and by default advertises each route-TLOC tuple to the Cisco Catalyst SD-WAN Controller. If a local site has two Cisco vEdge devices, a Cisco Catalyst SD-WAN Controller could potentially learn eight route-TLOC tuples for the same route. If the configured limit is lower than the number of route-TLOC tuples, the best route or routes are advertised. Range: 1 to 16 Default: 4
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco Catalyst SD-WAN Controller that can be installed in the Cisco vEdge device's local route table. By default, a Cisco vEdge device installs a maximum of four unique OMP paths into its route table. Range: 1 to 16 Default: 4
Send Backup Paths (on Cisco Catalyst SD-WAN Controllers only)	Click **On** to have OMP advertise backup routes to Cisco vEdge devices. By default, OMP advertises only the best route or routes. If you configure to send backup paths, OMP also advertises the first non-best route in addition to the best route or routes.
Shutdown	Ensure that **No** is chosen to enable to the Cisco SD-WAN overlay network. Click **Yes** to disable OMP and disable the Cisco SD-WAN overlay network. OMP is enabled by default.
Discard Rejected (on Cisco Catalyst SD-WAN Controllers only)	Click **Yes** to have OMP discard routes that have been rejected on the basis of policy. By default, rejected routes aren't discarded.

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

Configure OMP Timers

To configure OMP timers, click **Timers** and configure the following parameters:

Table 31:

Parameter Name	Description
Advertisement Interval	Specify the time between OMP Update packets. Range: 0 to 65535 seconds Default: 1 second We recommend you to configure 5 seconds on edge devices and 20 seconds on vSmart.

Parameter Name	Description
Hold Time	Specify how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed. Range: 0 to 65535 seconds Defaults, by Cisco Catalyst SD-WAN Control Components release: • 20.18.x and later: 300 seconds • 20.16.x: 5400 seconds • 20.12.1 to 20.15.x: 300 seconds • Before 20.12.1: 60 seconds Defaults, by Cisco IOS XE Catalyst SD-WAN release: • 17.18.1 and later: 300 seconds • 17.16.x: 5400 seconds
EOR Timer	Specify how long to wait after an OMP session has gone down and then come back up to send an end-of-RIB (EOR) marker. After this marker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. Range: 1 to 3600 seconds (1 hour) Default: 300 seconds (5 minutes)

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

Configure OMP Advertisements



Note Route advertisements in OMP are done either by applying the configuration at the global level or at the specific VPN level.

To advertise routes learned locally by the Cisco vEdge device to OMP, click **Advertise** and configure the following parameters:

Table 32:

Parameter Name	Description
Advertise	Click **On** or **Off** to enable or disable the Cisco vEdge device advertising to OMP the routes that it learns locally: • BGP—Click **On** to advertise BGP routes to OMP. By default, BGP routes are not advertised to OMP. • Connected—Click **Off** to disable advertising connected routes to OMP. By default, connected routes are advertised to OMP. • OSPF—Click **On** and click **On** again in the External field that appears to advertise external OSPF routes to OMP. OSPF inter-area and intra-area routes are always advertised to OMP. By default, external OSPF routes aren't advertised to OMP. • Static—Click **Off** to disable advertising static routes to OMP. By default static routes are advertised to OMP. To configure per-VPN route advertisements to OMP, use the VPN feature template.

Click **Save**.

Configure OMP Using CLI

By default, OMP is enabled on all Cisco vEdge devices and Cisco Catalyst SD-WAN Controllers. OMP must be operational for Cisco SD-WAN overlay network to function. If you disable it, you disable the overlay network.

OMP support in Cisco SD-WAN includes the following:

- IPv4 and IPv6 protocols, which are both turned on by default for VPN 0
- OMP route advertisements to BGP, EIGRP, OSPF, connected routes, static routes, and so on

Configure OMP Graceful Restart

OMP graceful restart is enabled by default on Cisco Catalyst SD-WAN Controllers and Cisco SD-WAN devices. OMP graceful restart has a timer that tells the OMP peer how long to retain the cached advertised routes. When this timer expires, the cached routes are considered to be no longer valid, and the OMP peer flushes them from its route table.

The default timer is 43,200 seconds (12 hours), and the timer range is 1 through 604,800 seconds (7 days). To modify the default timer value:

```
Device(config-omp)# timers graceful-restart-timer seconds
```

To disable OMP graceful restart:

```
Device(config-omp)# no omp graceful-restart
```

The graceful restart timer is set up independently on each OMP peer; that is, it's set up separately on each Cisco vEdge Device and Cisco Catalyst SD-WAN Controller. To illustrate what this means, let's consider a Cisco SD-WAN Controller that uses a graceful restart time of 300 seconds, or 5 minutes, and a Cisco vEdge

Device that is configured with a timer of 600 seconds (10 minutes). Here, Cisco Catalyst SD-WAN Controller retains the OMP routes learned from that device for 10 minutes—the graceful restart timer value that is configured on the device and that the device has sent to Cisco Catalyst SD-WAN Controller during the setup of the OMP session. The Cisco vEdge Device retains the routes it learns from the Cisco SD-WAN Controller for 5 minutes, which is the default graceful restart time value that is used on the Cisco Catalyst SD-WAN Controller and that the controller sent to the device, also during the setup of the OMP session.

While a Cisco Catalyst SD-WAN Controller is down and a Cisco vEdge Device is using cached OMP information, if you reboot the device, it loses its cached information and hence will not be able to forward data traffic until it is able to establish a control plane connection to Cisco Catalyst SD-WAN Controller.

Advertise Routes to OMP

By default, a Cisco vEdge Device advertises connected, static routes, and OSPF inter-area and intra-area routes to OMP, and hence to Cisco Catalyst SD-WAN Controller responsible for the device's domain. The device doesn't advertise BGP or OSPF external routes to OMP.

To have the device advertise these routes to OMP, and hence to Cisco Catalyst SD-WAN Controller responsible for the device's domain, use the advertise command:

Route advertisements in OMP are done either by applying the configuration at the global level or at the specific VPN level. To enable certain protocol route advertisements in all VPNs, you must add the configuration at the global level as shown in the example below.

```
Device# config
Device(config)# omp
Device(config-omp)# advertise bgp
Device(config-omp)# commit
```

To enable route advertisements for a certain protocol in only a few VPNs, you must remove any global-level configuration and add a per-VPN-level configuration as shown below:

```
Device# config
Device(config)# omp
Device(config-omp)# no advertise bgp
Device(config)# vpn 2
Device(config-vpn-2)# omp advertise bgp
Device(config-omp)# vpn 4
Device(config-vpn-4)# omp advertise bgp
Device(config-omp)# commit
```

To disable certain protocol route advertisements in all or a few VPNs, you should make sure that the configuration is present at neither the global level nor the VPN level.

For OSPF, the route type can be **external**.

The **bgp**, **connected**, **ospf**, and **static** options advertise all learned or configured routes of that type to OMP. To advertise a specific route instead of advertising all routes for a protocol, use the **network** option, specific the prefix of the route to advertise.

For individual VPNs, you can aggregate routes from the specified prefix before advertising them into OMP. By default, the aggregated prefixes and all individual prefixes are advertised. To advertise only the aggregated prefix, include the **aggregate-only** option.

Route advertisements that you set with the **omp advertise** command apply to all VPNs configured on the device. Route advertisements that you set with the **vpn omp advertise** command apply only to the specific VPN. If you configure route advertisements with both commands, they are both applied.

By default, when BGP advertises routes into OMP, BGP advertises each prefix's metric. BGP can also advertise the prefix's AS path:

```
Device(config)# vpn vpn-id router bgp
Device(config-bgp)# propagate-aspath
```

When you configure BGP to propagate AS path information, the device sends AS path information to devices that are behind the Cisco vEdge Devices (in the service-side network) that are running BGP, and it receives AS path information from these routers. If you are redistributing BGP routes into OMP, the AS path information is included in the advertised BGP routes. If you configure BGP AS path propagation on some but not all devices in the overlay network, the devices on which it isn't configured receive the AS path information but they don't forward it to the BGP routers in their local service-side network. Propagating AS path information can help to avoid BGP routing loops.

In networks that have both overlay and underlay connectivity—for example, when devices are interconnected by both a Cisco SD-WAN overlay network and an MPLS underlay network—you can assign as AS number to OMP itself. For devices running BGP, this overlay AS number is included in the AS path of BGP route updates. To configure the overlay AS:

```
Device(config)# omp
Device(omp)# overlay-as as-number
```

You can specify the AS number in 2-byte ASDOT notation (1–65535) or in 4-byte ASDOT notation (1.0 through 65535.65535). As a best practice, it's recommended that the overlay AS number be a unique AS number within both the overlay and the underlay networks. That use, select an AS number that isn't used elsewhere in the network.

If you configure the same overlay AS number on multiple devices in the overlay network, all these devices are considered to be part of the same AS, and as a result, they do not forward any routes that contain the overlay AS number. This mechanism is an additional technique for preventing BGP routing loops in the network.

Configure the Number of Advertised Routes

A Cisco vEdge Device can have up to eight WAN interfaces, and each WAN interface has a different TLOC. (A WAN interface is any interface in VPN 0 (or transport VRF) that is configured as a tunnel interface. Both physical and loopback interfaces can be configured to be tunnel interfaces.) This means that each router can have up to eight TLOCs. The device advertises each route–TLOC tuple to the Cisco Catalyst SD-WAN Controller.

The Cisco Catalyst SD-WAN Controller redistributes the routes it learns from Cisco vEdge Devices, advertising each route–TLOC tuple. If, for example, a local site has two devices, a Cisco Catalyst SD-WAN Controller could potentially learn eight route–TLOC tuples for the same route.

By default, Cisco vEdge Devices and Cisco Catalyst SD-WAN Controllers advertises up to four equal-cost route–TLOC tuples for the same route. You can configure devices to advertise from 1 to 16 route–TLOC tuples for the same route:

```
Device(config-omp)# send-path-limit 14
```

Beginning with Cisco Catalyst SD-WAN Control Components Release 20.8.x, you can configure a Cisco SD-WAN Controller operating in a Hierarchical SD-WAN environment to advertise from 1 to 32 route-TLOC tuples to edge devices for the same route.

Beginning with Cisco SD-WAN Controllers Release 20.9.x, you can configure a Cisco SD-WAN Controller in any Cisco SD-WAN environment to advertise from 1 to 32 route-TLOC tuples to edge devices for the same route.

If the limit is lower than the number of route–TLOC tuples, the Cisco vEdge Device or Cisco Catalyst SD-WAN Controller advertises the best routes.

Configure the Number of Installed OMP Paths

Cisco vEdge Devices install OMP paths that they received from the Cisco Catalyst SD-WAN Controller into their local route table. By default, a Cisco vEdge Devices installs a maximum of four unique OMP paths into its route table. You can modify this number:

```
Device(config-omp)# ecmp-limit 2
```

The maximum number of OMP paths installed can range from 1 through 16.

Configure the OMP Hold Time

The OMP hold time determines how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed.

We recommend that you configure OMP hold time to 300 seconds. The range is 0 to 65,535 seconds.

To modify the OMP hold time interval:

```
Device(config-omp)# timers holdtime 300
```

Defaults, by Cisco Catalyst SD-WAN Control Components release:

- 20.18.x and later: 300 seconds
- 20.16.x: 5400 seconds
- 20.12.1 to 20.15.x: 300 seconds
- Before 20.12.1: 60 seconds

Defaults, by Cisco IOS XE Catalyst SD-WAN release:

- 17.18.1 and later: 300 seconds
- 17.16.x: 5400 seconds

The keepalive timer is one-third the hold time and isn't configurable.

If the local device and the peer have different hold time intervals, the higher value is used.

If you set the hold time to 0, the keepalive and hold timers on the local device and the peer are set to 0.

The hold time must be at least two times the hello tolerance interval set on the WAN tunnel interface in VPN 0. To configure the hello tolerance interface, use the hello-tolerance command.

Configure the OMP Update Advertisement Interval

By default, OMP sends Update packets once per second. To modify this interval:

```
Device(config-omp)# timers advertisement-interval 5000
```

The interval can be in the range 0 through 65535 seconds.

Configure the End-of-RIB Timer

After an OMP session goes down and then comes back up, an end-of-RIB (EOR) marker is sent after 300 seconds (5 minutes). After this maker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. To modify the EOR timer:

```
Device(config-omp)# timers eor-timer 300
```

The time can be in the range 1 through 3600 seconds (1 hour).

Verify OMP Configuration Using the CLI

Verify OMP Routes

Table 33: Feature History

Feature Name	Release Information	Description
Verify OMP routes prefix	Cisco SD-WAN Release 20.8.1	The verify keyword is added to "show omp route <prefix>" CLI to validate the availability of route on Cisco vEdge devices.

Use the **show omp verify-routes** command to verify if a route prefix is available. This command helps to reduce the number of steps needed for troubleshooting an OMP prefix by verifying the received and installed RIB and FIB entries corresponding TLOCs and BFD sessions. For complete details, see show omp verify-routes command.

The following is a sample output from the **show omp verify-routes** command that displays the verification information for route prefixes:

```
Device# show omp verify-routes vpn 1 10.2.2.0/24
Codes Route/TLOC Status:
C    -> chosen
I    -> installed
Red  -> redistributed
Rej  -> rejected
L    -> looped
R    -> resolved
S    -> stale
Ext  -> extranet
Inv  -> invalid
Stg  -> staged
O    -> On-demand inactive
U    -> TLOC unresolved
Codes Rib Status:
    F -> fib, S -> selected, I -> inactive,
    B -> blackhole, R -> recursive, L -> import

               PATH                    ATTRIBUTE
STATUS      BFD      RIB
FROM PEER       ID     LABEL   STATUS      TYPE       TLOC IP         COLOR   ENCAP   TLOC
PREFERENCE  STATUS   STATUS
-------------------------------------------------------------------------------------------------
172.16.255.19    8      1005   C,I,R    installed    172.16.255.11   lte    ipsec   C,I,R
    -        up       F,S
172.16.255.19    9      1005   C,R      installed    172.16.255.11   3g     ipsec   C,R
    -        up        -
```

Verify OMP Peer Sessions

Note Starting from Cisco SD-WAN Release 20.8.1 **show support omp peer** command is added for Cisco vEdge devices.

The following is a sample output from the **show support omp peer** command displaying the active OMP peer sessions information on Cisco SD-WAN Controllers or Cisco vEdge devices:

```
Device# show support omp peer peer-ip 172.16.255.41
========================================
            PEERS for CONTEXT 172.16.255.41
========================================
 Local address: 172.16.255.41
 Looking up Peer: 172.16.255.5
 Peer: 172.16.255.5 (0x7fd197ee1800), Type: vSmart, Site: 200, Region-id-set: None, Domain:
 1, Overlay: 1, Legit: yes
        State: Up, version: 1, Control-Up: yes, Staging: no, flags: 0x21
        CAP: BR: no, TGW: no
       Multithreading- down: no, move-marker: no, update-gen: no, work-queue: no, needs_upd:
 0x0
        buffer ev: 0x0x7fd197aca580
        fd: 21
        Hello timer: Enabled (e: 2, c: 20, md: 20 lmd: 0)  Hold timer: Enabled (e: 43 v:
60 c: 60)
        Connect retry: Disabled (e: -1 v: 2 c: 2)  Adv. timer: Enabled (e: 1 v: 1 c: 1)
        Down-pending: Disabled (e: -1 v: 1 c: 1)
    EOR interval: 300 EOR timer: Disabled (e: -1 v: 300)

For complete sample output, see show omp peer sessions.
```

Configure OSPFv3 IPv4 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 a Transport VPN in a Transport and Management profile or a Service VPN feature in a Service profile.

Step 3 Create and configure a OSPFv3 IPv4.

a. Configure Basic Configuration.

Table 34: Basic Settings

Field	Description
Router ID	Enter the OSPF router ID, in decimal four-part dotted notation. This value is the IP address that is associated with the router for OSPF adjacencies. Default: No Router ID is configured.
Add Redistribute	

Field	Description
Protocol	Choose the protocol from which to redistribute routes into OSPFv3, for all OSPFv3 sessions. • **Connected** • **Static** • **Nat-route** • **BGP**
Select Route Policy	Enter the name of a localized control policy to apply to routes before they are redistributed into OSPF.

b. Configure Area Parameters.

Table 35: Area

Field	Description
Area Number*	Enter the number of the OSPFv3 area. Allowed value: Any 32-bit integer
Area Type	Choose the type of OSPFv3 area: • **Stub -** no external routes • **NSSA**: not-so-stubby area, allows external routes • **Normal** **Note** You can't enter a value for **Area type** if you have entered 0 as a value for **Area Number**.
Interface	
Add Interface	Configure the properties of an interface in an OSPFv3 area.
Name*	Enter the name of the interface. Examples of interface names: GigabitEthernet0/0/1, GigabitEthernet0/1/2.1, GigabitEthernet0, or Loopback1.
Cost	Specify a number for the Type 3 summary link-state advertisement (LSA). OSPFv3 uses this metric during its SPF calculation to determine the shortest path to a destination. Range: 0 through 16777215
Authentication Type	Specify the SPI and authentication key if you use IPSec SHA1. • **no-auth**: Select no authentication. • **ipsec-sha1**: Enter the value for the IPSEC Secure Hash Algorithm 1 (SHA-1) authentication.

Field	Description
SPI	Specifies the Security Policy Index (SPI) value. Range: 256 through 4294967295
Authentication Key	Provide a value for the authentication key. When IPSEC SHA-1 authentication is used, the key must be 40 hex digits long.
Passive Interface	Specify whether to set the OSPFv3 interface to be passive. A passive interface advertises its address, but does not actively run the OSPFv3 protocol. Default: Disabled
IPv4 Range	
Add IPv4 Range	Configure the area range of an interface in an OSPFv3 area.
Network Address*	Enter the IPv4 address.
Subnet Mask*	Enter the subnet mask.
No Advertise*	Enable this option to not advertise the Type 3 summary LSAs.
Cost	Specify the cost of the OSPFv3 interface. Range: 1 through 65535

c. Configure Advanced Parameters.

Table 36: Advanced

Field	Description
Route Policy	Enter the name of a localized control policy to apply to routes coming from OSPFv3 neighbors.
Reference Bandwidth (Mbps)	Specify the reference bandwidth for the OSPFv3 autocost calculation for the interface. Range: 1 through 4294967 Mbps Default: 100 Mbps
RFC 1583 Compatible	By default, the OSPFv3 calculation is done per RFC 1583. Disable this option to calculate the cost of summary routes based on RFC 2328.

Field	Description
Originate	Enable this option to generate a default external route into an OSPF routing domain. When you enable this option, the following fields appear: • **Always**: Enable this option to always advertise the default route in an OSPF routing domain. • **Default Metric**: Set the metric used to generate the default route. Range: 0 through 16777214 Default: 10 • **Metric Type**: Choose to advertise the default route as an OSPF Type 1 external route or an OSPF Type 2 external route.
Distance	Define the OSPFv3 route administration distance based on route type. Default: 100
Distance for External Routes	Set the OSPFv3 distance for routes learned from other domains. Range: 0 through 255 Default: 110
Distance for Inter-Area Routes	Set the distance for routes coming from one area into another. Range: 0 through 255 Default: 110
Distance for Intra-Area Routes	Set the distance for routes within an area. Range: 0 through 255 Default: 110
SPF Calculation Timers	Configure the amount of time between when OSPFv3 detects a topology and when it runs its SPF algorithm.
SPF Calculation Delay (milliseconds)	Specify the amount of time between when the first change to a topology is received until performing the SPF calculation. Range: 1 through 600000 ms (600 seconds) Default: 200 ms
Initial Hold Time (milliseconds)	Specify the amount of time between consecutive SPF calculations. Range: 1 through 600000 ms (600 seconds) Default: 1000 ms
Maximum Hold Time (milliseconds)	Specify the longest time between consecutive SPF calculations. Range: 1 through 600000 ms (600 seconds) Default: 10000 ms (10 seconds)

Field	Description
Maximum Metric (Router LSA)	Configure OSPFv3 to advertise a maximum metric so that other routers do not prefer this Cisco vEdge Device as an intermediate hop in their Shortest Path First (SPF) calculation. • **Immediately**: Force the maximum metric to take effect immediately, through operator intervention. • **On-startup**: Advertise the maximum metric for the specified number of seconds after the router starts up. Range: 5 through 86400 seconds Maximum metric is disabled by default.

What to do next

Also see Deploy a configuration group.

Configure OSPFv3 IPv6 Routing 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 a Transport VPN feature in a Transport and Management profile or a Service VPN feature in a Service profile.

Step 3 Create and configure a OSPFv3 IPv6 feature.

a. Configure Basic Configuration.

Table 37: Basic Settings

Field	Description
Router ID	Enter the OSPF router ID, in decimal four-part dotted notation. This value is the IP address that is associated with the router for OSPF adjacencies. Default: No Router ID is configured.
Add Redistribute	

Field	Description
Protocol	Choose the protocol from which to redistribute routes into OSPFv3, for all OSPFv3 sessions. • **Connected** • **Static** • **BGP**
Select Route Policy	Enter the name of a localized control policy to apply to routes before they are redistributed into OSPF.

b. Configure Area Parameters.

Table 38: Area

Field	Description
Area Number*	Enter the number of the OSPFv3 area. Allowed value: Any 32-bit integer
Area Type	Choose the type of OSPFv3 area: • **Stub**: No external routes • **NSSA**: Not-so-stubby area, allows external routes • **Normal** **Note** You can't enter a value for **Area type** if you have entered 0 as a value for **Area Number**.
Interface	
Add Interface	Configure the properties of an interface in an OSPFv3 area.
Name*	Enter the name of the interface. Examples of interface names: GigabitEthernet0/0/1, GigabitEthernet0/1/2.1, GigabitEthernet0, or Loopback1.
Cost	Specify a number for the Type 3 summary link-state advertisement (LSA). OSPFv3 uses this metric during its SPF calculation to determine the shortest path to a destination. Range: 0 through 16777215
Authentication Type	Specify the SPI and authentication key if you use IPSec SHA1. • **no-auth**: Select no authentication. • **ipsec-sha1**: Enter the value for the IPSEC Secure Hash Algorithm 1 (SHA-1) authentication.

Field	Description
SPI	Specifies the Security Policy Index (SPI) value. Range: 256 through 4294967295
Authentication Key	Provide a value for the authentication key. When IPSEC SHA-1 authentication is used, the key must be 40 hex digits long.
Passive Interface	Specify whether to set the OSPFv3 interface to be passive. A passive interface advertises its address, but does not actively run the OSPFv3 protocol. Default: Disabled
IPv6 Range	
Add IPv6 Range	Configure the area range of an interface in an OSPFv3 area.
Network Address*	Enter the IPv6 address.
Subnet Mask*	Enter the subnet mask.
No Advertise*	Enable this option to not advertise the Type 3 summary LSAs.
Cost	Specify the cost of the OSPFv3 interface. Range: 1 through 65535

c. Configure Advanced Parameters.

Table 39: Advanced

Field	Description
Route Policy	Enter the name of a localized control policy to apply to routes coming from OSPFv3 neighbors.
Reference Bandwidth (Mbps)	Specify the reference bandwidth for the OSPFv3 autocost calculation for the interface. Range: 1 through 4294967 Mbps Default: 100 Mbps
RFC 1583 Compatible	By default, the OSPFv3 calculation is done per RFC 1583. Disable this option to calculate the cost of summary routes based on RFC 2328.

Field	Description
Originate	Enable this option to generate a default external route into an OSPF routing domain. When you enable this option, the following fields appear: • **Always**: Enable this option to always advertise the default route in an OSPF routing domain. • **Default Metric**: Set the metric used to generate the default route. Range: 0 through 16777214 Default: 10 • **Metric Type**: Choose to advertise the default route as an OSPF Type 1 external route or an OSPF Type 2 external route.
Distance	Define the OSPFv3 route administration distance based on route type. Default: 100
Distance for External Routes	Set the OSPFv3 distance for routes learned from other domains. Range: 0 through 255 Default: 110
Distance for Inter-Area Routes	Set the distance for routes coming from one area into another. Range: 0 through 255 Default: 110
Distance for Intra-Area Routes	Set the distance for routes within an area. Range: 0 through 255 Default: 110
SPF Calculation Timers	Configure the amount of time between when OSPFv3 detects a topology and when it runs its SPF algorithm.
SPF Calculation Delay (milliseconds)	Specify the amount of time between when the first change to a topology is received until performing the SPF calculation. Range: 1 through 600000 ms (600 seconds) Default: 200 ms
Initial Hold Time (milliseconds)	Specify the amount of time between consecutive SPF calculations. Range: 1 through 600000 ms (600 seconds) Default: 1000 ms
Maximum Hold Time (milliseconds)	Specify the longest time between consecutive SPF calculations. Range: 1 through 600000 ms (600 seconds) Default: 10000 ms (10 seconds)

Field	Description
Maximum Metric (Router LSA)	Configure OSPFv3 to advertise a maximum metric so that other routers do not prefer this vEdge router as an intermediate hop in their Shortest Path First (SPF) calculation. • **Immediately**: Force the maximum metric to take effect immediately, through operator intervention. • **On-startup**: Advertise the maximum metric for the specified number of seconds after the router starts up. Range: 5 through 86400 seconds Maximum metric is disabled by default.

What to do next

Also see Deploy a configuration group.

Configure EIGRP Routing 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 Configure a EIGRP Routing feature in Service Profile.

a. Configure basic settings.

Table 40: Basic Configuration

Parameter Name	Description
Autonomous System ID *	Enter the local autonomous system (AS) number. Range: 1 through 65535 Default: None
Network	
IP Address*	Enter the IPv4 address.
Mask*	Enter the subnet mask.
Interface	

Parameter Name	Description
Add Interface	Provide values for the following fields: • **AF Interface**: Enter a value for the Address Family (AF) interface. • **Shutdown**: Enables the interface to run EIGRP by default. Toggle ON to disable the interface. • **Add Summary Address**: Enter an IPv4 address and choose a subnet mask.

b. Configure IPv4 unicast address family.

Table 41: IPv4 Unicast Address Family

Parameter Name	Description
Protocol *	Select one of the protocols from which to redistribute routes into EIGRP, for all EIGRP sessions: • **bgp**: Redistribute Border Gateway Protocol (BGP) routes into EIGRP. • **connected**: Redistribute connected routes into EIGRP. • **nat-route**: Redistribute network address translation (NAT) routes into EIGRP. • **omp**: Redistribute Overlay Management Protocol (OMP) routes into EIGRP. • **ospf**: Redistribute Open Shortest Path First (OSPF) routes into EIGRP. **Note** From Cisco IOS XE Catalyst SD-WAN Release 16.12.1b and later, you can set metric values for redistribution by using the CLI add-on feature template. Use the following command: `redistribute ospf 1 metric 1000000 1 1 1 1500` For more information, see CLI Add-on Feature Templates. • **ospfv3**: OSPFv3 routes into EIGRP. • **static**: Redistribute static routes into EIGRP.
Route Policy *	Enter the name of the route policy to apply to redistributed routes.

c. Configure authentication.

Table 42: Authentication

Parameter	Description
MD5*	**MD5 Key ID**: Enter an MD5 key ID to compute an MD5 hash over the contents of the EIGRP packet using that value.
	MD5 Authentication Key: Enter an MD5 authentication key to use an encoded MD5 checksum in the transmitted packet.
	Authentication Key: A 256-byte unique key that is used to compute the Hashed Message Authentication Code (HMAC) and is known both by the sender and the receiver of the message.
HMAC-SHA-256	**Authentication Key**: A 256-byte unique key that is used to compute the HMAC and is known both by the sender and the receiver of the message.

d. Configure advanced settings.

Table 43: Advanced Configuration

Parameter Name	Description
Hold Time (seconds)	Set the interval after which EIGRP considers a neighbor to be down. The local router then terminates the EIGRP session to that peer. This acts as the global hold time. Range: 0 through 65535 Default: 15 seconds
Hello Interval (seconds)	Set the interval at which the router sends EIGRP hello packets. Range: 0 through 65535 Default: 5 seconds
Route Policy	Enter the name of an EIGRP route policy.
Filter	Toggle **ON** to filter routes that do not match the policy.

What to do next

Also see Deploy a configuration group.