



Unicast Overlay Routing

The overlay network is controlled by the Cisco SD-WAN Overlay Management Protocol (OMP), which is at the heart of Cisco SD-WAN overlay routing. This solution allows the building of scalable, dynamic, on-demand, and secure VPNs. The Cisco 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 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 11](#)

Supported Protocols

This section explains the protocols supported for unicast routing.

OMP Routing Protocol

The Cisco SD-WAN Overlay Management Protocol (OMP) is the protocol responsible for establishing and maintaining the Cisco 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 vSmart 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 vSmart 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 vSmart 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 vSmart 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, on page 32](#).

OMP advertises the following types of routes:

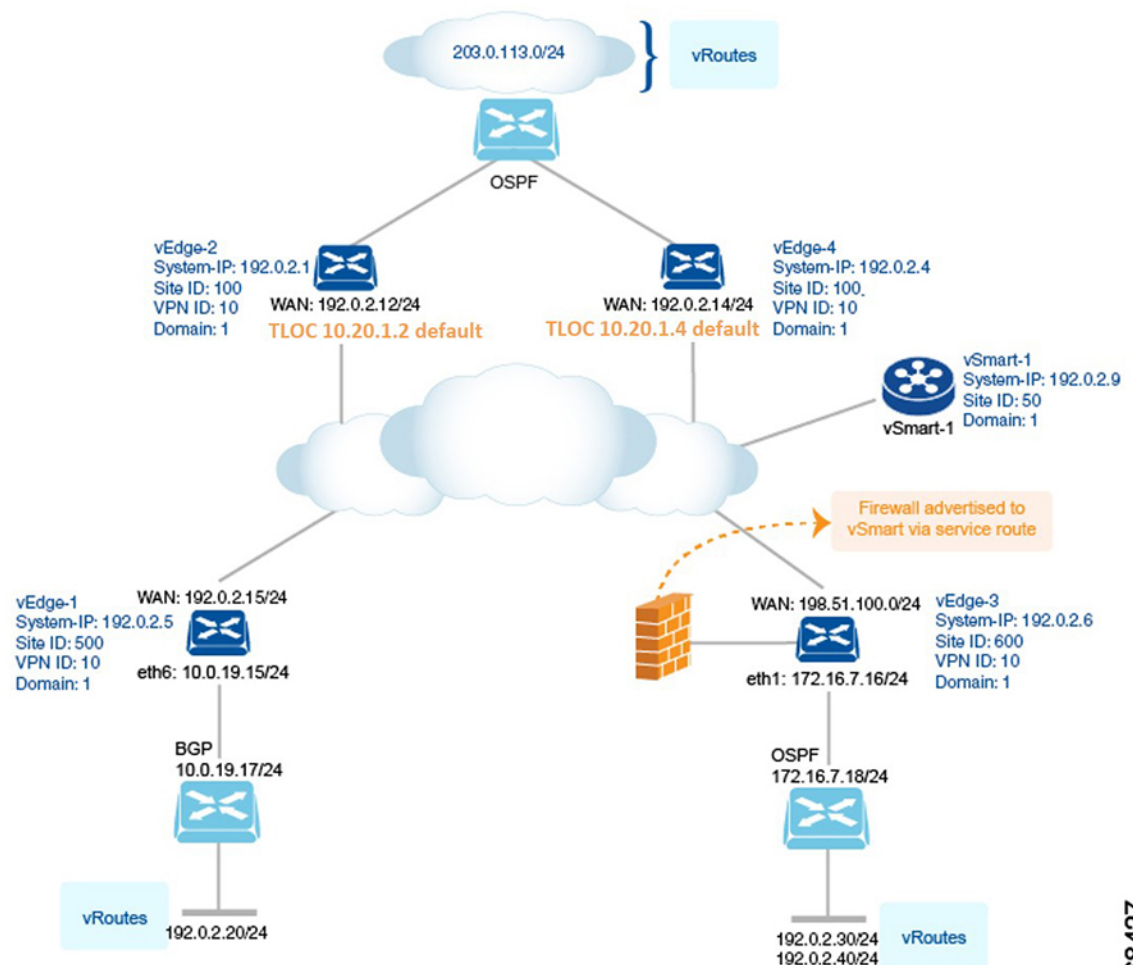
- OMP routes (also called vRoutes)—Prefixes that establish reachability between end points that use the OMP-orchestrated transport network. OMP routes can represent services in a central data center, services at a branch office, or collections of hosts and other end points in any location of the overlay 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 vSmart 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.



368427

OMP Routes

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

A Cisco 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 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 vSmart 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.

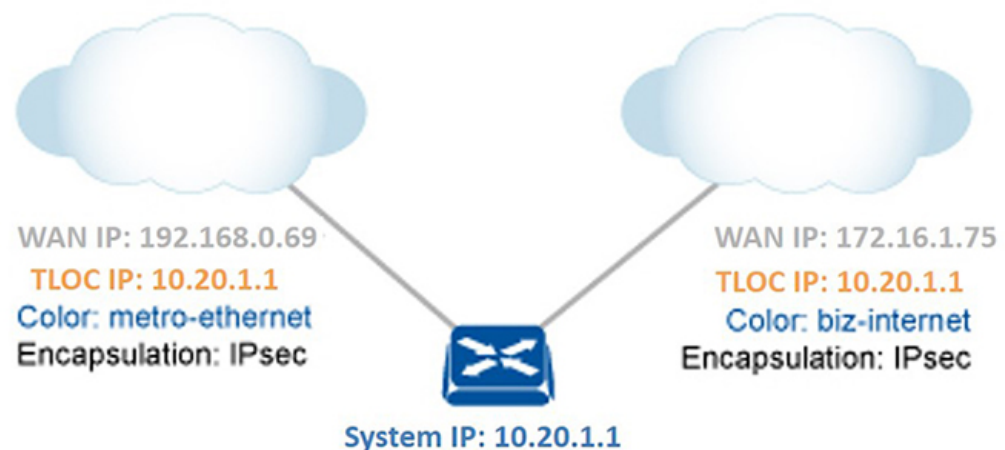
TLOC routes advertise the following attributes:

- 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 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 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.



368487

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 vSmart 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 vSmart Controllers using service address family NLRI. See *Service Chaining*.

OMP Route Advertisements for Cisco vSmart Controllers

Table 1: Feature History

Feature Name	Release Information	Description
Increased OMP Path Limit for Cisco vSmart 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 vSmart 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 vSmart 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 vSmart Controllers per VPN per prefix is extended to a maximum of 128.

Limitations

- Multitenant Cisco vSmart 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 vSmart Controller can send to another Cisco vSmart 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 vSmart 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 vSmart Controller performs a full route database update. This leads to complete network updates.

For more information, 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
!
```

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 vSmart Controller and the router take the origin type and subtype into consideration.



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

Table 2:

OMP Route Origin Type	OMP Route Origin Subtype
BGP	External Internal
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 vSmart 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 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

Protocol	Administrative Distance
EIGRP	Internal: 90, External: 170
OSPF	110
OSPFv3	110
IS-IS	115
IBGP	200
OMP	250

OMP Best-Path Algorithm and Loop Avoidance

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

1. Select an ACTIVE route. An ACTIVE route is preferred over a STALE route. An active route is a route from a peer with which an OMP session is UP. A stale route is a route from a peer with which an OMP session is in Graceful Restart mode.



Note A stale route will only be advertised if the stale version is similar to the Route Information Base (RIB) version. Otherwise, the stale route is dropped.

2. Check whether the OMP route is valid. If not, ignore it.
3. If the OMP route is valid and if it has been learned from the same Cisco SD-WAN device, select the OMP route with the lower administrative distance.
4. If the administrative distances are equal, select the OMP route with the higher OMP route preference value.
5. On Edge devices only, if the OMP route preference values are equal, select the OMP route with the higher TLOC preference value.
6. If the TLOC preference values are equal, compare the origin type and subtype, and select one in the following order (select the first match).
 - 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

7. If the origin type is the same, select the OMP route that has the lower origin metric.



Note All routes in the following steps are considered equal, but are sorted as per below criteria:

8. Edge router sourced route is preferred over the same route coming from Cisco vSmart Controller.
9. If the origin types are the same, select the OMP route with the lower router ID.
10. If the router IDs are equal, a Cisco vEdge device selects the OMP route with the lower private IP address. If a Cisco vSmart 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 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 route:

- A Cisco vSmart Controller receives an OMP route 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 route from another Cisco vSmart Controller, also with an origin code of OSPF. If all other things are equal, the best-path algorithm chooses the route that came from the Cisco vEdge device.
- A Cisco vSmart Controller learns the same OMP route, 10.10.10.0/24, from two Cisco vEdge devices in the same site. If all other parameters are the same, both routes are chosen and advertised to other OMP peers. By default, up to four equal-cost routes are selected and advertised.

A Cisco vEdge device installs an OMP route 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 vSmart Controller removes from the forwarding table all the OMP routes that point to that TLOC.

OMP Graceful Restart

Graceful restart for OMP allows the data plane in the Cisco SD-WAN overlay network to continue functioning if the control plane stops functioning or becomes unavailable. With graceful restart, if the Cisco vSmart controller in the network goes down, or if multiple Cisco vSmart 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 vSmart controller. When a Cisco vSmart 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 vSmart controller.

When OMP graceful restart is enabled, Cisco vEdge devices and a Cisco vSmart 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 vSmart controller, the device continues forwarding data traffic using the cached OMP information. The device also periodically checks whether the Cisco vSmart 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 vSmart controller to be valid and reliable. This same scenario occurs when a Cisco vSmart 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 vSmart 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 vSmart 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.

BGP and OSPF Routing Protocols

The Cisco 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 vBond Orchestrator learns about this DTLS connection via the loop-back interface and conveys this information to the Cisco vSmart 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 vSmart 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 vManage, 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 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 SD-WAN devices use VRFs in place of VPNs. However, the following steps still apply to configure Cisco IOS XE SD-WAN devices through Cisco vManage. When you complete the configuration, the system automatically maps the VPN configurations to VRF configurations.

To configure the BGP routing protocol using Cisco vManage 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 vManage 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. To create a template for VPNs **1** through **511**, and **513** through **65530**:
 - 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 vSmart 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:			
<ul style="list-style-type: none"> • For service-side BGP routing, choose OMP. By default, OMP routes are not redistributed into BGP. • For transport-side BGP routing, choose Connected, and then under Route Policy, specify a route policy that has BGP advertise the loopback interface address to its neighbors. 				
Route Policy	Enter the name of the route policy to apply to redistributed routes.			
Click Add to save the redistribution information.				
Network	Click Network > New Network .			
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.		
	Network Prefix	Enter a network prefix, in the format <i>prefix/length</i> to be advertised by BGP.		
	Click Add to save the network prefix.			

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

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

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. <i>Range:</i> 1 through 65535 minutes	
	Warning Only	Click On to display a warning message without restarting the BGP connection.	
	Route Policy In	Click On and specify the name of a route policy that will have the prefixes from the neighbour.	
Route Policy Out	Click On and specify the name of a route policy that will have the prefixes sent to the neighbour.		
Shutdown	Click On to enable the connection to the BGP neighbor.		

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

Configure Advanced BGP Parameters

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

Parameter Name	Description
Hold Time	<p>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.</p> <p>Range: 0 through 65535 seconds</p> <p>Default: 180 seconds (three times the keepalive timer)</p>
Keepalive	<p>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.</p> <p>Range: 0 through 65535 seconds</p> <p>Default: 60 seconds (one-third the hold-time value)</p>
Compare MED	<p>Click On to always compare MEDs regardless of whether the peer ASs of the compared routes are the same.</p>
Deterministic MED	<p>Click On to compare multiple exit discriminators (MEDs) from all routes received from the same AS, regardless of when the route was received.</p>
Missing MED as Worst	<p>Click On to consider a path as the worst path if the path is missing a MED attribute.</p>
Compare Router ID	<p>Click On to compare the device IDs among BGP paths to determine the active path.</p>

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 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 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 vSmart Controller. Doing so allows the Cisco vSmart 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 OSPF

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

To configure OSPF on a device using Cisco vManage 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 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 vManage 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. To create a template for VPNs 1 through 511, and 513 through 65530:
 - 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 4:

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

Configure 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 5:

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. <i>Range: 0 through 255 Default: 110</i>
Distance for Inter-Area Routes	Specify the OSPF route administration distance for routes coming from one area into another. <i>Range: 0 through 255 Default: 110</i>

Parameter Name	Description
Distance for Intra-Area Routes	Specify the OSPF route administration distance for routes within an area. <i>Range: 0 through 255 Default: 110</i>

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 6:

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 7:

Parameter Name	Description
Type	Choose a type: <ul style="list-style-type: none"> • 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. <i>Range: 0, 5 through 86400 seconds Default: 0 seconds (the maximum metric is advertised immediately when the router starts up)</i>

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 8:

Parameter Name	Description
Area Number	Enter the number of the OSPF area. <i>Range:</i> 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 SD-WAN devices that are ABRs (area border routers) to translate Type 7 LSAs to Type 5 LSAs: <ul style="list-style-type: none"> • 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 9:

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. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 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 device assumes that the neighbor is down. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 40 seconds (4 times the default hello interval)

Parameter Name	Description
LSA Retransmission Interval	Specify how often the OSPF protocol retransmits LSAs to its neighbors. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 5 seconds
Interface Cost	Specify the cost of the OSPF interface. <i>Range:</i> 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 10:

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. <i>Range:</i> 0 through 255 <i>Default:</i> 1
OSPF Network Type	Choose the OSPF network type to which the interface is to connect: <ul style="list-style-type: none"> • Broadcast network—WAN or similar network. • Point-to-point network—Interface connects to a single remote OSPF router. • Non-broadcast—Point-to-multipoint. <i>Default:</i> 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. <i>Default:</i> Off
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: <ul style="list-style-type: none"> • 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.

Parameter Name	Description
• 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 11:

Parameter Name	Description
Address	Enter the IP address and subnet mask, in the format <i>prefix/length</i> 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. <i>Range: 0 through 16777215</i>
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 12:

Parameter Name	Description
Reference Bandwidth	Specify the reference bandwidth for the OSPF auto-cost calculation for the interface. <i>Range: 1 through 4294967 Mbps Default: 100 Mbps</i>
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.

Parameter Name	Description
Originate	<p>Click On to generate a default external route into an OSPF routing domain:</p> <ul style="list-style-type: none"> • 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. <i>Range: 0 through 16777214 Default: 10</i> • 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	<p>Specify the amount of time between when the first change to a topology is received until performing the SPF calculation.</p> <p><i>Range: 0 through 600000 milliseconds (60 seconds) Default: 200 milliseconds</i></p>
Initial Hold Time	<p>Specify the amount of time between consecutive SPF calculations.</p> <p><i>Range: 0 through 600000 milliseconds (60 seconds) Default: 1000 milliseconds</i></p>
Maximum Hold Time	<p>Specify the longest time between consecutive SPF calculations.</p> <p><i>Range: 0 through 600000 Default: 10000 milliseconds (60 seconds)</i></p>
Policy Name	<p>Enter the name of a localized control policy to apply to routes coming from OSPF neighbors.</p>

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 vSmart 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 vSmart 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 an 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

Use the OMP template to configure OMP parameters for all Cisco vEdge devices, and for Cisco vSmart Controllers.

OMP is enabled by default on all Cisco vEdge devices, Cisco vManage NMSs, and Cisco vSmart 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, see the *Configure OMP Advertisements* section in this topic.

Create OMP Template

1. From the Cisco vManage 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 13:

Parameter Scope	Scope Description
Device Specific (indicated by a host icon)	<p>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.</p> <p>When you click Device Specific, the Enter Key box opens. This box displays a key, which is a unique string that identifies the parameter in a CSV file that you create. This file is an Excel spreadsheet that contains one column for each key. The header row contains the key names (one key per column), and each row after that corresponds to a device and defines the values of the keys for that device. You upload the CSV file when you attach a Cisco SD-WAN device to a device template.</p> <p>To change the default key, type a new string and move the cursor out of the Enter Key box.</p> <p>Examples of device-specific parameters are system IP address, hostname, GPS location, and site ID.</p>
Global (indicated by a globe icon)	<p>Enter a value for the parameter, and apply that value to all devices.</p> <p>Examples of parameters that you might apply globally to a group of devices are DNS server, syslog server, and interface MTUs.</p>

Configure Basic OMP Options

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

Table 14:

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. <i>Range:</i> 0 through 604800 seconds (168 hours, or 7 days) <i>Default:</i> 43200 seconds (12 hours)
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 vSmart Controllers, and the controllers redistributes the learned routes, advertising each route-TLOC tuple. A Cisco vEdge device can have up to four TLOCs, and by default advertises each route-TLOC tuple to the Cisco vSmart Controller. If a local site has two Cisco vEdge devices, a Cisco vSmart 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. <i>Range:</i> 1 through 16 <i>Default:</i> 4
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco vSmart 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. <i>Range:</i> 1 through 32 <i>Default:</i> 4
Send Backup Paths (on Cisco vSmart 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 vSmart 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 15:

Parameter Name	Description
Advertisement Interval	Specify the time between OMP Update packets. <i>Range: 0 through 65535 seconds Default: 1 second</i>
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. <i>Range: 0 through 65535 seconds Default: 60 seconds</i>
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. <i>Range: 1 through 3600 seconds (1 hour) Default: 300 seconds (5 minutes)</i>

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 16:

Parameter Name	Description
Advertise	<p>Click On or Off to enable or disable the Cisco vEdge device advertising to OMP the routes that it learns locally:</p> <ul style="list-style-type: none"> • 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. <p>To configure per-VPN route advertisements to OMP, use the VPN feature template.</p>

Click **Save**.

Configure OMP Using CLI

By default, OMP is enabled on all Cisco vEdge devices and Cisco vSmart 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 vSmart 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 vSmart Controller. To illustrate what this means, let's consider a vSmart 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 vSmart 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 vSmart Controller during the setup of the OMP session. The Cisco vEdge Device retains the routes it learns from the vSmart controller for 5 minutes, which is the default graceful restart time value that is used on the Cisco vSmart Controller and that the controller sent to the device, also during the setup of the OMP session.

While a Cisco vSmart 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 vSmart 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 vSmart 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 vSmart 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 an 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 vSmart Controller.

The Cisco vSmart 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 vSmart Controller could potentially learn eight route–TLOC tuples for the same route.

By default, Cisco vEdge Devices and Cisco vSmart 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 SD-WAN Controllers Release 20.8.x, you can configure a Cisco vSmart controller operating in a Hierarchical 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 vSmart Controller advertises the best routes.

Configure the Number of Installed OMP Paths

Cisco vEdge Devices install OMP paths that they received from the Cisco vSmart 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. The default OMP hold time is 60 seconds but it can be configured to up to 65,535 seconds. To modify the OMP hold time interval:

```
Device(config-omp)# timers holdtime 75
```

The hold time can be in the range 0 through 65535 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 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. 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 17: 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
STATUS   BFD      RIB      ATTRIBUTE
```

FROM PEER PREFERENCE	ID STATUS	LABEL STATUS	STATUS	TYPE	TLOC IP	COLOR	ENCAP	TLOC
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 vSmart 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](#).