Unicast Overlay Routing

Note To achieve simplification and consistency, the Cisco SD-WAN solution has been rebranded as Cisco Catalyst SD-WAN. In addition, from Cisco IOS XE SD-WAN Release 17.12.1a and Cisco Catalyst SD-WAN Release 20.12.1, the following component changes are applicable: **Cisco vManage** to **Cisco Catalyst SD-WAN Manager**, **Cisco vAnalytics** to **Cisco Catalyst SD-WAN Analytics**, **Cisco vBond** to **Cisco Catalyst SD-WAN Validator**, **Cisco vSmart** to **Cisco Catalyst SD-WAN Controller**, and **Cisco Controllers** to **Cisco Catalyst SD-WAN Control Components**. See the latest Release Notes for a comprehensive list of all the component brand name changes. While we transition to the new names, some inconsistencies might be present in the documentation set because of a phased approach to the user interface updates of the software product.

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 21

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 IOS XE Catalyst SD-WAN 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 VRF setting are usually located within a VRF domain, and they are protected so that they are not visible outside the VRF. In such a traditional architecture, it is a challenge to extend VRF 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 IOS XE Catalyst SD-WAN 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, on page 42.

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.

The following figure illustrates the two types of OMP routes.

Figure 1: Different Types of OMP Routes

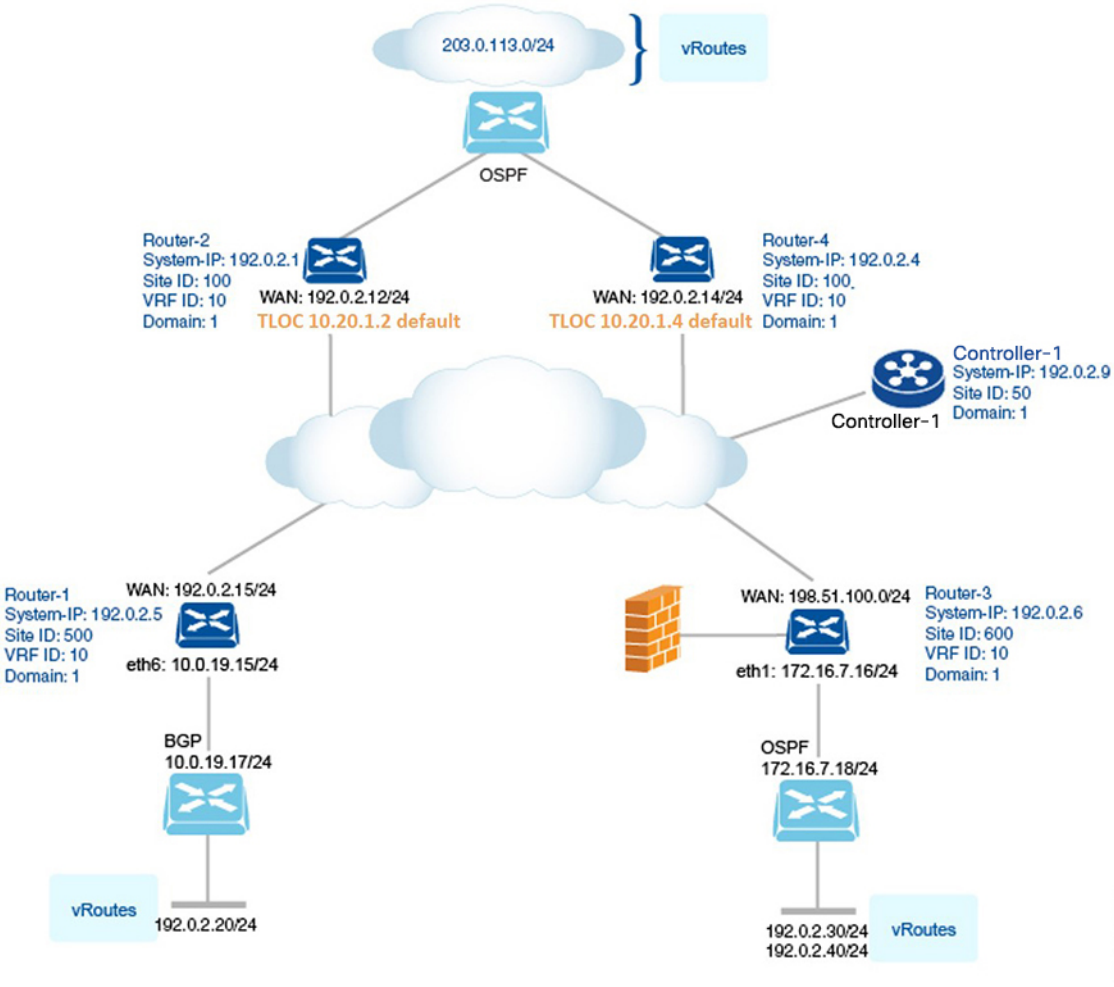

OMP Routes

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
- EIGRP
- LISP
- OSPF (inter-area, intra-area, and external)
- OSPFv3 (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.
- 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.
- VRF—VRF 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 VRF. 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

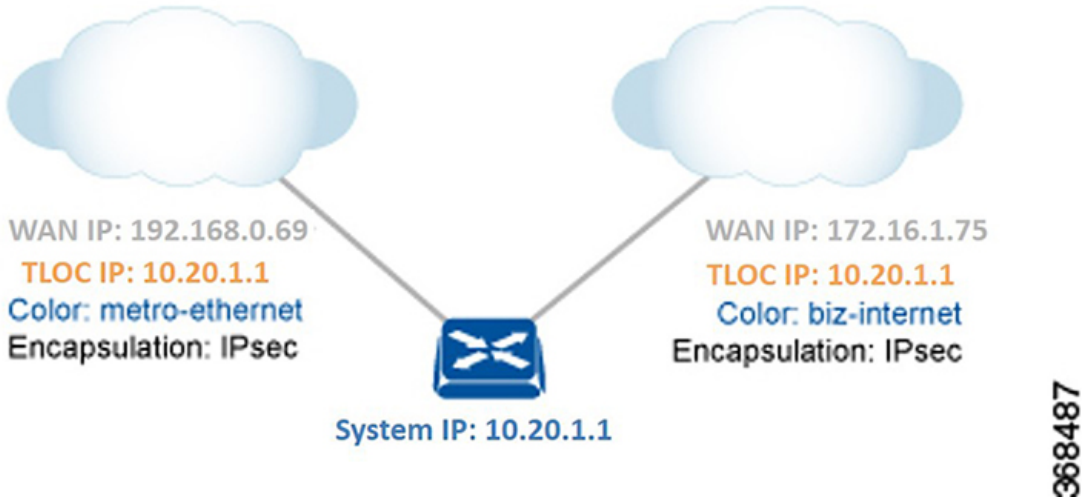

The diagram to the right shows a device that has two WAN connections and hence two TLOCs. The system IP address of the router is 10.20.1.1. The TLOC on the left is uniquely identified by the system IP address 10.20.1.1, the color metro-ethernet, and the encapsulation IPsec, and it maps to the physical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

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 IOS XE Catalyst SD-WAN Release 17.5.1a	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 IOS XE Catalyst SD-WAN Release 17.5.1a , 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               15
  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               15
  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               15
  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
OSPFv3	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
Static	—
EIGRP	• EIGRP Summary • EIGRP Internal • EIGRP External
LISP	—
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

Protocol	Administrative Distance
OSPFv3	110
IS-IS	115
IBGP	200
OMP	251

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.

Step	Applies to	Description
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.)
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.

Step	Applies to	Description
14	Edge devices	**TLOC preference** Select an OMP path with a higher TLOC preference value.
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:

- A Cisco Catalyst SD-WAN Controller receives an OMP path to 10.10.10.0/24 via OMP from a Cisco IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN device.
- A Cisco Catalyst SD-WAN Controller learns the same OMP path, 10.10.10.0/24, from two Cisco IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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.

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 IOS XE Catalyst SD-WAN devices in any VRF except for transport and management VRFs to provide reachability to networks at their local sites. Cisco IOS XE Catalyst SD-WAN 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 types of topologies, the BGP or OSPF sessions run over a DTLS connection created on the loopback interface in VRF 0, which is the transport VRF that is responsible for carrying control traffic in the overlay network. The Cisco Catalyst SD-WAN Validator learns about this DTLS connection via the loopback interface and conveys this information to the Cisco Catalyst SD-WAN Controller so that it can track the TLOC-related information. In VRF 0, you also configure the physical interface that connects the Cisco IOS XE Catalyst SD-WAN 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.

BGP Community Propagation

Table 5: Feature History

Feature Name	Release Information	Description

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

Feature Name	Release Information	Description
BGP Community Propagation	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature enables propagation of BGP communities between routing protocols during route redistribution. On one node, the OMP redistributes routes from BGP and on the other node, the BGP redistributes routes from OMP. In addition to configurable AS path attribute propagation, there is an option to propagate BGP communities. The BGP community propagation helps in propagating BGP communities between Cisco Catalyst SD-WAN sites, across VPNs using OMP redistribution. To propagate the BGP communities during route redestribution from OMP, use the **propagate-community** command.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.3.1a, the community propagation feature is supported. Without this option, no BGP communities are sent to the BGP neighbor, even if they are attached. With this feature, the Cisco IOS XE Catalyst SD-WAN device can start propagating the communities attached to the BGP entries to the neighbor. The BGP overlay is migrated to a Cisco Catalyst SD-WAN overlay where BGP route attributes are propagated between Cisco Catalyst SD-WAN sites across VPNs. For more information on **propagate-community** command, refer propagate-community.

Starting from Cisco IOS XE Catalyst SD-WAN Release 17.4.1a, you can manipulate communities when propagating communities from BGP to OMP and back from OMP to BGP using the `route-map` command. It defines the conditions for redistributing routes from one routing protocol into another routing protocol. Each **`route-map`** command has a list of `match` and `set` commands associated with it. The `match` commands specify the `match communities`, the conditions under which redistribution is allowed. The `set` commands specify the `set communities`, the particular redistribution actions to perform if the criteria enforced by the `match` commands are met. For more information on the commands, refer Command Reference Guide.

OSPFv3

Table 6: Feature History

Feature Name	Release Information	Description
OSPFv3 Support on Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 17.3.2 Cisco vManage Release 20.3.1	Open Shortest Path First version 3 (OSPFv3) is an IPv4 and IPv6 link-state routing protocol that supports IPv6 and IPv4 unicast address families.

OSPFv3 is a routing protocol for IPv4 and IPv6 address families. It is a link-state protocol that makes its routing decisions based on the states of the links that connect source and destination machines. The state of a link is a description of that interface and its relationship to its neighboring networking devices. The interface information includes the IPv6 prefix of the interface, the network mask, the type of network it is connected to, the devices connected to that network, and so on. This information is propagated in various type of link-state advertisements (LSAs).

Much of OSPFv3 is the same as in OSPF version 2. OSPFv3, which is described in RFC 5340, expands on OSPF version 2 to provide support for IPv6 routing prefixes and the larger size of IPv6 addresses.

For address family IPv6, OSPFv3 routes are referred to OSPF routes and OSPFv3 internal routes (intra-area and inter-area) are implicitly advertised to OMP. OSPFv3 external routes (both AS-External and NSSA) can be explicitly advertised in OMP using the advertise OSPF external configuration. This is consistent with OSPF routes in address family IPv4 where OSPF internal routes are implicitly advertised in OMP. Similarly, OSPF external routes can be explicitly advertised to OMP using the advertise OSPF external configuration.

For address family IPv4, OSPFv3 routes are referred to as OSPFv3 routes and OSPFv3 internal routes are not implicitly advertised in OMP. All OSPFv3 IPv4 routes can be advertised in OMP using the advertise OSPFv3 configuration. OSPFv3 integration in controller mode is not supported.

EIGRP

Cisco EIGRP (Enhanced Interior Gateway Routing Protocol) is a Cisco proprietary routing protocol. It is an open-standard Interior Gateway Protocol (IGP). EIGRP is an enhancement to the original Interior Gateway Routing Protocol (IGRP developed) by Cisco. EIGRP does not fully update if there are no changes in the network. This reduces the flooding activities in other IGPs. It also can use both equal cost and unequal cost paths, which is unique among IGPs.

EIGRP is supported only on Cisco IOS XE Catalyst SD-WAN devices.

See Introduction to EIGRP for more information in EIGRP.

Benefits of EIGRP

- Increased network width from 15 to 100 hops
- Fast convergence
- Incremental updates, minimizing bandwidth
- Protocol-independent neighbor discovery
- Easy scaling

Limitations and Restrictions

- EIGRP is not supported on the transport side network on Cisco IOS XE Catalyst SD-WAN devices.
- EIGRP route match is not supported in Cisco SD-WAN Controller centralized control policy.

Routing Information Protocol (RIP)

Table 7: Feature History

Feature Name	Release Information	Description
RIPv2 Support on Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 17.7.1a Cisco vManage Release 20.7.1 Cisco SD-WAN Release 20.7.1	This feature enables you to configure RIPv2 on Cisco IOS XE Catalyst SD-WAN devices. Routers redistribute RIPv2 routes to Overlay Management Protocol (OMP) for advertisement in the Cisco Catalyst SD-WAN overlay, and to Open Shortest Path First version 3 (OSPFv3) for service-side routing.
RIPng (IPv6) Support on Cisco IOS XE Catalyst SD-WAN Devices	Cisco IOS XE Catalyst SD-WAN Release 17.8.1a Cisco vManage Release 20.8.1 Cisco SD-WAN Release 20.8.1	This feature adds support for IPv6 addresses and prefixes on Cisco IOS XE Catalyst SD-WAN devices. It also supports redistribution of connect, static, Overlay Management Protocol (OMP), and Open Shortest Path First (OSPF) routes into Routing Information Protocol next generation (RIPng).

Information About Routing Information Protocol Support

The Routing Information Protocol (RIP) uses broadcast or multicast User Datagram Protocol (UDP) data packets to exchange routing information. RIP is a commonly used routing protocol in small to medium TCP/IP networks. RIP uses a distance-vector algorithm to calculate routes. Cisco IOS software sends routing information updates every 30 seconds, which is termed as advertising. RIP sends routing-update messages at regular intervals, and when the network topology changes.

RIPv2 (RIP for IPv4)

In the Cisco IOS software implementation of RIP Version 2 (RIPv2), each RIP process maintains a local database. The RIP local database contains a set of best-cost RIP routes that are learned from all the networking devices neighboring to RIP-enabled routers. Route redistribution allows routes to be specified by a prefix, using a route map and prefix list.

The Cisco implementation of RIPv2 supports plain text and message digest algorithm 5 (MD5) authentication, route summarization, classless interdomain routing (CIDR), and variable-length subnet masks (VLSMs). If you are sending and receiving RIPv2 packets, we recommend that you enable RIP authentication on an interface because RIPv1 does not support authentication. Plain text authentication is the default authentication in every RIPv2 packet.

By default, the software receives RIP Version 1 (RIPv1) and RIPv2 packets, but sends only RIPv1 packets. You can configure the software to receive and send only RIPv1 packets. Alternatively, you can configure the software to receive and send only RIPv2 packets. To override the default behavior, you can configure the RIP

version that an interface sends. Similarly, you can also control how packets that are received from an interface are processed. RIP v2 is supported on both service side and transport side.

Note For network configuration, we recommend that you use Classful IP Network ID Addressing.

See Configure Routing Information Protocol Using the CLI for more details on configurations using the CLI.

RIPng (RIP for IPv6)

Routing Information Protocol next generation (RIPng) is a UDP-based protocol for communicating routing information that is used to compute routes through IPv6 networks. RIP enhancements for IPv6, which are detailed in RFC 2080, include support for IPv6 addresses and prefixes.

RIPng as an Interior Gateway Protocol (IGP) supports redistribution of the following:

- OMP routes into RIP
- RIP routes into OMP
- RIP routes into OSPFv3
- OSPFv3 routes into RIP
- Static routes into RIP
- RIP routes into static
- Connect routes into RIP
- RIP routes into connect

Each router that implements RIPng requires a routing table containing the following fields:

- The IPv6 prefix of the destination.
- Metric: Total cost of the metric advertised for the address.
- Route Tag: A route attribute that must be advertised and redistributed with the route.
- Next-hop IPv6 address of the destination.
- Various timers associated with the routes.

When not in Virtual Routing and Forwarding (VRF) mode, every IPv6 RIPng process and the configuration that is associated with it keeps all the routes in the same routing table. The IPv6 RIPng VRF-aware support enables isolation, modularity, and potential performance improvement by reducing the number of routes stored in a single routing table. It also allows a network administrator to create different RIP routing tables and share the same protocol configuration that is stored in a single RIP protocol configuration block.

RIPng in large networks is prone to routing loops, making the traffic take a longer path. To avoid route looping, RIP and RIPng routes are identified using the well-known OMP RIP tag.

The following figure illustrates the RIPv2 and RIPng OMP route tagging process:

Figure 3: RIPv2 and RIPng Topology

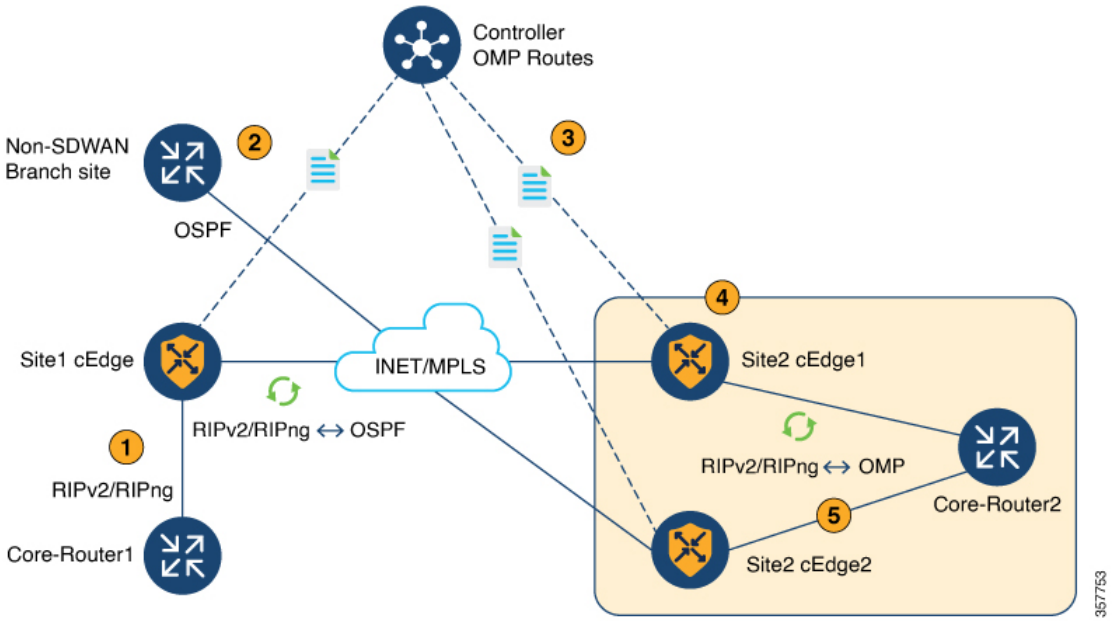

1. Core-Router1 advertises RIPv2 and RIPng routes to Site1.

   As a general rule, the RIPv2 and RIPng routes have a default administrative distance of 120. The default administrative distance for OMP routes is 251.

2. The RIPv2 and RIPng route is redistributed and advertised in OMP.
3. The Cisco Catalyst SD-WAN Controller advertises an OMP route to the other branch.
4. Site-2 Edge1 router adds an OMP route tag of a unique value of 44270, and redistributes the OMP-learned route into RIPv2 and RIPng.
5. When the Site-2 Edge2 router receives this route with the tag 44270, it will *not* install this route because it is already learning a route through OMP, which has AD 251.

   If the OMP route is withdrawn, the Site-2 Edge2 router installs the route, which is learned through the RIPv2 and RIPng protocol through service-side VPN with the tag 44270, into the routing table with an administrative distance of 252 (one value higher than that of OMP).

   Additionally, a Cisco Catalyst SD-WAN tagged route will not be readvertised in OMP when the RIPv2 and RIPng route is redistributed to OMP.

See Configure RIPng Using the CLI for more details on RIPng configurations using the CLI.

Prerequisites for Using Routing Information Protocol

- Version 2 must be configured to send and receive only RIPv2 packets. By default, RIP Version 1 (RIPv1) and RIP Version 2 (RIPv2) packets are received, but only RIPv1 packets are sent.

Restrictions for Using Routing Information Protocol

RIPv2 (IPv4)

RIP uses hop count as the metric to rate the value of different routes. Hop count is the number of devices that can be traversed in a route. A directly connected network has a metric of zero; an unreachable network has a metric of 16. This limited metric range makes RIP unsuitable for large networks.

RIPng (IPv6)

- Only the **sdwan** keyword can be used to configure the IPv6 RIP routing process name (*ripng-instance*) in the configuration commands.
- VRF-aware support in IPv6 RIP allows only one RIP instance at a given time. More than one RIP instance is not allowed.
- You can configure RIPng on only GigabitEthernet, TenGigabitEthernet, and VLAN interfaces.

Configure Unicast Overlay Routing

This topic describes how to provision unicast overlay routing.

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 IOS XE Catalyst SD-WAN devices to participate in the overlay network.

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

Configure BGP

The Border Gateway Protocl (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. 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.

Parameter Name	Description
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
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 MPLS Interface

Table 8: Feature History

Feature Name	Release Information	Description
MPLS-BGP Support on the Service Side	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This features allows you to enable support on Multiprotocol Label Switching (MPLS). Multiple Service VPNs use inter autonomous system (AS) BGP labelled path to forward the traffic, which in turn helps scaling the service side VPNs with less control plane signaling. Label distribution for a given VPN routing and forwarding (VRF) instance on a given device can be handled by Border Gateway Protocol (BGP).

The Cisco IOS XE Catalyst SD-WAN devices support Multiprotocol Label Switching (MPLS) to enable multiple protocol environment. MPLS offers an extremely scalable, protocol agnostic, data-carrying mechanism that transfers data packets with assigned labels across the network through virtual links. Extensions of the

BGP protocol can be used to manage an MPLS path. The Cisco IOS XE Catalyst SD-WAN devices also have the capability of BGP MPLS VPN Option B.

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

To configure an MPLS interface, do the following:

- Click **MPLS Interface**.
- Enter the interface name in the **Interface Name** field.
- You can click on + to add more interfaces and save the configuration.

Configure Label Range

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

Sample configuration:

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

Configure Route Targets

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

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

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

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

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

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*.
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

Parameter Name	Description
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)

Parameter Name	Description
Keepalive	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local device is still active and should be considered available. This keepalive time is the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Compare MED	Click **On** to always compare MEDs regardless of whether the peer ASs of the compared routes are the same.
Deterministic MED	Click **On** to compare multiple exit discriminators (MEDs) from all routes received from the same AS, regardless of when the route was received.
Missing MED as Worst	Click **On** to consider a path as the worst path if the path is missing a MED attribute.
Compare Router ID	Click **On** to compare the device IDs among BGP paths to determine the active path.
Multipath Relax	Click **On** to have the BGP best-path process select from routes in different in ASs. By default, when you are using BGP multipath, the BGP best path process selects from routes in the same AS to load-balance across multiple paths.

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

Configure BGP Using CLI

This is an example of a BGP configuration on a Cisco IOS XE Catalyst SD-WAN device for releases before Cisco IOS XE Catalyst SD-WAN Release 17.4.1a.

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


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


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

Note Starting from Cisco IOS XE Catalyst SD-WAN Release 17.4.1a, the following changes apply to BGP configuration under non-VRF address-family:

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

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

Following is the sample BGP configuration for Cisco IOS XE Catalyst SD-WAN Release 17.4.1a and later:

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

Verify BGP Redistribute Route in OMP

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

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

```
vm5# show sdwan omp routes 192.168.0.0/16 detail
------------------------------------------------
```

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

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

1. Configure a VPN.

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

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

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

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

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

Example of BGP Configuration on a SD-WAN IOS XE Router

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

!
address-family ipv4 unicast
neighbor 10.20.25.16 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 10.20.25.16 activate
neighbor 10.20.25.16 send-community extended
exit-address-family
!
address-family vpnv6 unicast
neighbor 10.20.25.16 activate
neighbor 10.20.25.16 send-community extended
exit-address-family
!
address-family ipv4 unicast vrf 1
redistribute connected
redistribute static
exit-address-family
!
address-family ipv6 unicast vrf 1
redistribute connected
redistribute omp

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

exit-address-family
```

Example of configuring route targets:

```
vrf config

vrf definition 1
rd 1:1

!
address-family ipv4

route-target export 200:1

route-target import 100:1
```

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

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.

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

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

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

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

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

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

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

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

You can configure the Cisco IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN devices in the overlay network. You can advertise BGP routes either globally or for a specific VRF:

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

Configure OSPF

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

Note Cisco IOS XE 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 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 transport-side routing to enable communication between the Cisco Catalyst SD-WAN devices when the router is not directly connected to the WAN cloud.
2. Create a VPN feature template to configure VPN parameters for 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. 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 9:

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 10:

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.

Parameter Name	Description
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 11:

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 IOS XE Catalyst SD-WAN 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 12:

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.

Parameter Name	Description
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 13:

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 14:

Parameter Name	Description
Interface Name	Enter the name of the interface, in the format **ge** *slot*/*port* or **loopback** *number*.

Parameter Name	Description
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 IOS XE Catalyst SD-WAN device must receive an OSPF hello packet from its neighbor. If no packet is received, the Cisco IOS XE Catalyst SD-WAN 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 15:

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
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.

Parameter Name	Description
• **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 16:

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 17:

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 OSPF for Unicast overlay routing.

Configure Basic Service-Side OSPF

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

Note When configuring OSPF from the CLI, ensure that the OSPF process id (PID) and the VRF ID match for OMP redistribution of OSPF to work for the specified VRF. The process ID is the ID of the OSPF process to which the interface belongs. The process ID is local to the router and is used as an identifier of the local OSPF process.

Here is an example of configuring service-side OSPF on a Cisco IOS XE Catalyst SD-WAN device.

```
config-transaction
 router ospf 1 vrf1
  auto-cost reference-bandwidth 100
  max-metric router-lsa
  timers throttle spf 200 1000 10000
  router-id 172.16.255.15
  default-information originate
  distance ospf external 110
  distance ospf inter-area110
  distance ospf intra-area110
  distredistribute connected subnets route-map route_map
  exit
 interface GigabitEthernet0/0/1
  no shutdown
  arp timeout 1200
  vrf forwarding 1
  ip address 10.1.100.14 255.255.255.0
  ip redirects
  ip mtu 1500
  ip ospf 1 area 0
  ip ospf network broadcast
  mtu 1500
  negotiation auto
  exit
```

Configure OMP

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

OMP is enabled by default on all Cisco IOS XE Catalyst SD-WAN 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 VRF level. For more information about route advertisements in OMP, see the *Configure OMP Advertisements* section in this topic.
- 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.

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 18:

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 19:

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 through 604800 seconds (168 hours, or 7 days) *Default:* 43200 seconds (12 hours)
Number of Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. s advertise routes to Cisco Catalyst SD-WAN Controllers, and the controllers redistributes the learned routes, advertising each route-TLOC tuple. A Cisco IOS XE Catalyst SD-WAN device 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 Cisco IOS XE Catalyst SD-WAN 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 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 Cisco IOS XE Catalyst SD-WAN device's local route table. By default, a Cisco IOS XE Catalyst SD-WAN device installs a maximum of four unique OMP paths into its route table. *Range:* 1 through 16 *Default:* 4
Send Backup Paths (on Cisco Catalyst SD-WAN Controllers only)	Click **On** to have OMP advertise backup routes to Cisco IOS XE Catalyst SD-WAN 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 20:

Parameter Name	Description
Advertisement Interval	Specify the time between OMP Update packets. *Range:* 0 through 65535 seconds *Default:* 1 second
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 through 65535 seconds *Default:* 60 seconds Starting from Cisco Catalyst SD-WAN Manager Release 20.12.1, the default hold time is 300 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 through 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 VRF level.

To advertise routes learned locally by the Cisco IOS XE Catalyst SD-WAN device to OMP, click **Advertise** and configure the following parameters:

Table 21:

Parameter Name	Description
Advertise	Click **On** or **Off** to enable or disable the Cisco IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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:

- IPv6 service routes
- IPv4 and IPv6 protocols, which are both turned on by default
- 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-transaction
Device(config)# sdwan
Device(config-omp)# timers graceful-restart-timer seconds
```

To disable OMP graceful restart:

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

The graceful restart timer is set up independently on each OMP peer; that is, it's set up separately on each Cisco IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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

Table 22: Feature History

Feature Name	Release Information	Description
OMP Route Aggregation	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature is an enhancement where OMP route aggregation is performed only for the routes that are configured for route redistribution to avoid black hole routing. This enhancement is applicable for OSPF, Connected, Static, BGP and other protocols only if the redestribution is requested.

By default, a Cisco IOS XE Catalyst SD-WAN device advertises connected routes, static routes, OSPF inter-area, OSPF intra-area routes, OSPFv3 IPv6 intra-area routes, and OSPF IPv6 inter-area routes are advertised to OMP for Cisco Catalyst SD-WAN Controller, that is responsible for the device's domain.

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.

Note Configuration of route advertisements in OMP can be done either by applying the configuration at the global level or at the specific VRF level.

The example below enables OMP advertisement of BGP routes for all VRFs. To enable protocol route advertisements for OMP protocol for all VRFs, add the configuration at the global level.

```
Device(config)# sdwan
Device(config-sdwan)# omp
Device(config-omp)# address-family ipv4
Device(config-ipv4)# advertise bgp
```

To enable protocol route advertisements for a few VRFs, remove the global-level configuration using **no advertise bgp** command and add a per-VRF-level configuration:

```
Device(config)# sdwan
Device(config-sdwan)# omp
Device(config-omp)# address-family ipv4
Device(config-ipv4)# no advertise bgp
```

```
Device(config-ipv4)# address-family ipv4 vrf 2
Device(config-vrf-2)# advertise bgp
Device(config-vrf-2)# address-family ipv4 vrf 4
Device(config-vrf-4)# advertise bgp
Device(config-vrf-4)# commit
```

Note To disable certain protocol route advertisements for all or for a few VRFs, ensure that the configuration is present at neither the global level nor the VRF level.

To configure the routes that the device advertises to OMP for all VRFs configured on the device:

```
config-transaction
 sdwan
  omp
   address-family ipv4
    advertise ospf external
    advertise bgp
    advertise eigrp
    advertise connected
    advertise static
    exit
  address-family ipv6
   advertise ospf external
   advertise bgp
   advertise eigrp
   advertise connected
   advertise static
   exit
```

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, and specify the prefix of the route to advertise.

To configure the routes that the device advertises to OMP for a specific VRF on the device:

```
config-transaction
 sdwan
  omp
   address-family ipv4 vrf 1
    advertise aggregate prefix 10.0.0.0/8
    advertise ospf external
    advertise bgp
    advertise eigrp
    advertise connected
    advertise static
    exit
  address-family ipv6 vrf 1
   advertise aggregate 2001:DB8::/32
   advertise ospf external
   advertise bgp
   advertise eigrp
   advertise connected
   advertise static
   exit
```

For individual VRFs, routes from the specified prefix can be aggregated after advertising them into OMP using **advertise** *protocol* config command. By default, the aggregated prefixes and all individual prefixes are advertised. To advertise only the aggregated prefix, include the **aggregate-only** option as shown below.

```
config-transaction
 sdwan
  omp
   address-family ipv4 vrf 1
    advertise aggregate 10.0.0.0/8 aggregate-only
    exit
```

Note Route advertisements in OMP are done either by applying configuration at the global level or to specific VRFs. The specific VRF configuration doesn't override global-VRF configuration in OMP.

When BGP advertises routes into OMP, it advertises each prefix's metric. BGP can also advertise the prefix's AS path.

```
config-transaction
 router bgp 200
 address-family ipv4 vrf 11
  neighbor 10.20.1.0 remote-as 200
  propagate-aspath
  exit
```

When you configure BGP to propagate AS path information, the device sends AS path information to devices that are behind the Cisco IOS XE Catalyst SD-WAN devices (in the service-side network) that are running BGP, and it receives AS path information from these routers. If you're 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's not 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:

```
config-transaction
 sdwan
  omp
   overlay-as 55
   exit
```

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 don't 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN 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 IOS XE Catalyst SD-WAN device or Cisco Catalyst SD-WAN Controller advertises the best routes.

Configure the Number of Installed OMP Paths

Cisco IOS XE Catalyst SD-WAN devices install OMP paths that they received from the Cisco Catalyst SD-WAN Controller into their local route table. By default, a Cisco IOS XE Catalyst SD-WAN 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.

Starting from Cisco Catalyst SD-WAN Manager Release 20.12.1, the default hold time is 300 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 transport VRF. 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).

Mapping Multiple BGP Communities to OMP Tags

Table 23: Feature History

Feature Name	Release Information	Description
Mapping Multiple BGP Communities to OMP Tags	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature allows you to display information about OMP routes on Cisco Catalyst SD-WAN Controller and Cisco IOS XE Catalyst SD-WAN devices. OMP routes carry information that the device learns from the routing protocols running on its local network, including routes learned from BGP and OSPF, as well as direct, connected, and static routes.

For more information on the **show sdwan omp routes** command, refer show sdwan omp routes.

Configure OSPFv3

To configure OSPFv3 routing protocol using Cisco SD-WAN Manager templates follow these steps:

1. Create an OSPFv3 feature template to configure OSPFv3 parameters.
2. Create a VPN feature template to configure VPN parameters for service-side routing (any VPN other than VPN 0 or VPN 512).
3. Create a device template and apply the templates to the correct devices.

Create an OSPFv3 Template

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

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

3. Click **Add Template** and choose a device from the list.
4. From the **Other Templates** section, choose **OSPFv3** and enter a name and a description for the template.
5. Choose **IPv4** or **IPv6**.

Basic Configuration

Click **Basic Configuration** to configure the basic details for the template.

Parameter Name	Description
Router ID	Enter the IP address of the router. For example: 10.20.1.1
Distance	Enter the administrative distance where you want the router to be installed.
External Routes	Specify the OSPFv3 route administrative distance for routes learned from other domains. Range: 0 through 255 Default: 110
Inter-Area Routes	Enter a value to apply as the OSPFv3 route administrative distance for routes coming from one area into another. Range: 0 through 255 Default: 110
Intra-Area Routes	Enter a value to apply as the OSPF route administrative distance for routes from a directly connected area. Range: 0 through 255 Default: 110
Timers Throttle SPF	Specify the shortest-path first (SPF) timer for throttling.
Table Map	Specify a route-map to modify route attributes or filter routes that OSPFv3 installs in the global or VRF routing table.
Filter	Click **On** to filter routes that are not accepted by the route-map specified for the table map.

Configure Redistribute Routes into OSPFv3

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

Table 24: Redistribution Parameters

Parameter Name	Value	Description
Mark as Optional Row	Click **Optional** 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 OSPFv3, for all OSPFv3 sessions.	
	bgp	Redistribute BGP routes into OSPFv3.
	connected	Redistribute connected routes into OSPFv3.
	nat-route	Redistribute NAT routes into OSPFv3.
	omp	Redistribute OMP routes into OSPFv3.
	eigrp	Redistribute EIGRP routes into OSPFv3.
	lisp	Redistribute LISP routes into OSPFv3.
	isis	Redistribute IS-IS routes into OSPFv3.
	ospf	Redistribute OSPF routes into OSPFv3. **Note** Redistribute of OSPF is supported only for IPv4 address family.
	static	Redistribute static routes into OSPFv3.
Route Policy *	Enter the name of the route policy to apply to redistributed routes.	

Click **Save**.

Configure OSPFv3 To Advertise a Maximum Metric

To configure OSPFv3 to advertise a maximum metric so that other devices do not prefer the Cisco IOS XE Catalyst SD-WAN 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 25:

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 after the startup.

Parameter Name	Description
Advertisement Time	If you chose **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)

Click **Save**.

Configure OSPFv3 Areas

To configure an OSPFv3 area within a VPN on a Cisco IOS XE Catalyst SD-WAN device, select **Area** > **Add New Area**. For OSPFv3 to function, you must configure area 0.

Table 26:

Parameter Name	Description
Area Number	Enter the number of the OSPFv3 area. *Range:* 32-bit number
Set the Area Type	Choose the type of OSPFv3 area. The options are: • normal • stub - no external routes • nssa - not-so-stubby area, allows external routes
No Summary	If you configured the area type as stub or NSSA, click **On** to prevent OSPFv3 summary routes from being injected into the area.
Translate	If you configured the area type as NSSA, choose when you allow devices configured as area border routers (ABR) to translate Type 7 LSAs to Type 5 LSAs: • Always—The 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—The router offers translation services, but does not insist on being the translator. • Never—The router never become the NSSA translator for Type 7 LSAs.

To configure the properties of an interface in an OSPFv3 area, choose **Area** > **Add New Area** > **Add Interface**. In the **Add Interface** pop-up, configure the following parameters:

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

Parameter Name	Description
Dead Interval	Specify the time interval within which the Cisco IOS XE Catalyst SD-WAN device must receive an OSPF hello packet from its neighbor. If no packet is received, the Cisco IOS XE Catalyst SD-WAN device assumes 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 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:

Parameter Name	Description
Address	Enter the IP address and prefix length, in the format *prefix*/*length* for the IP or IPv6 address twice to be consolidated and advertised. The address type is dependent on the address family.
Cost	Specify a number for the Type 3 summary LSA. OSPFv3 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 interface configuration, click **Save**.

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

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

Configure Advanced OSPFv3 Properties

To configure other OSPFv3 properties, click **Advanced**:

Table 27:

Parameter Name	Description
Reference Bandwidth (Mbps)	Specify the reference bandwidth for the OSPFv3 auto-cost calculation for the interface. *Range:* 1 through 4294967 Mbps *Default:* 100 Mbps

Parameter Name	Description
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—Choose the metric type, OSPF Type 1 external route or an OSPF Type 2 external route to advertise the default route.
SPF Calculation Delay (milliseconds)	Specify the 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 (milliseconds)	Specify the time between consecutive SPF calculations. *Range*: 0 through 600000 milliseconds (60 seconds) *Default*: 1000 milliseconds
Maximum Hold Time (milliseconds)	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 the routes installed by OSPFv3 into the global Route Information Base (RIB).
Filter	Filter inhibit OSPFv3 routes that do not match the policy from being installed in the global RIB.

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

Configure OSPFv3 Using CLI

To configure OSPFv3 on Cisco IOS XE SD-WAN devices on IPv4 and IPv6 address families:

```
config-transaction
router ospfv3 <vpn-id>
  !
  address-family ipv4 unicast vrf <vpn-id>
  router-id <ipv4-address-format>
  auto-cost reference-bandwidth <1-4294967>
  default-information originate [always] [route-map <route-map-name>] [metric <1-16777214>]

                                                  [metric-type {1|2}]
  distance <1-254>
  distance ospf {external <1-254> | intra-area <1-254> | inter-area <1-254>}
  timers throttle spf <1-600000> <1-600000> <1-600000>
  redistribute {bgp <1-4294967295>| connected | eigrp <vpn-id>| isis <vpn-id>| lisp |
nat-route | omp |
                        ospf <vpn-id> | static}
                     [route-map <route-map-name>]
  max-metric router-lsa [on-startup <5-86400>]
  table-map <route-map-name> [filter]
```

```
  area <1-4294967295> stub [no-summary]
  area <1-4294967295> nssa [no-summary] [translate type7 always]
  area <1-4294967295> range <ipv4-prefix-address> <ipv4-prefix-mask>  ! 192.168.0.1
255.255.255.0
                                                    [not-advertise | advertise] [cost
<1-16777214>]16777214
exit-address-family


address-family ipv6 unicast vrf <vpn-id>
  router-id <ipv4-address-format>
  auto-cost reference-bandwidth <1-4294967>
  default-information originate [always] [route-map <route-map-name>] [metric <1-16777214>]

                                                  [metric-type {1|2}]
  distance <1-254>
  distance ospf {external <1-254> | intra-area <1-254> | inter-area <1-254>}
  timers throttle spf <1-600000> <1-600000> <1-600000>
  redistribute {bgp <1-4294967295> | connected | eigrp <vpn-id>| isis <vpn-id>| lisp | omp
 |
                        static}
                       [route-map <route-map-name>]
  max-metric router-lsa [on-startup <5-86400>]
  table-map <route-map-name> [filter]
  area <1-4294967295> stub [no-summary]
  area <1-4294967295> nssa [no-summary] [translate type7 always]
  area <1-4294967295> range <ipv6-prefix>              ! 2001:DB8::/48
                                                    [not-advertise | advertise] [cost
<1-16777214>]
exit-address-family
```

OSPFv3 Table-Map Configuration

```
router ospfv3 1
 !
 address-family ipv4 unicast vrf 1
  redistribute omp route-map match-omp-tag
  table-map set-omp-tag
 exit-address-family
 !
 address-family ipv6 unicast vrf 1
  table-map set-omp-tag
  redistribute omp route-map match-omp-tag
 exit-address-family
!
route-map set-omp-tag permit 20
 set omp-tag 2000
route-map match-omp-tag permit 10
 match omp-tag 1000
 set metric 20
route-map match-omp-tag permit 20
 match omp-tag 2000
 set metric 30
route-map match-omp-tag deny 30
```

OSPFv3 IPsec Interface Authentication

The following example shows the OSPFv3 interface authentication configuration with an md5 key:

```
interface GigabitEthernet2
 vrf forwarding 1
```

```
 ip address 192.168.0.1 255.255.255.0
 negotiation auto
 ipv6 address 192.168.0.1/64
 ospfv3 authentication ipsec spi 256 md5 FEEDACEEDEADBEEFFEEDACEEDEADBEEF
 ospfv3 1 ipv6 area 0
 ospfv3 1 ipv4 area 0
 no mop enabled
 no mop sysid
!
```

The following example shows the OSPFv3 IPsec authentication configuration with a SHA1 key:

```
interface GigabitEthernet4
 vrf forwarding 1
 ip address 192.168.0.1 255.255.255.0
 negotiation auto
 ipv6 address 192.168.0.1/64
 ospfv3 authentication ipsec spi 300 sha1 FEEDACEEDEADBEEFFEEDACEEDEADBEEFFEEDACEE
 ospfv3 1 ipv4 area 0
 no mop enabled
 no mop sysid
```

Configure EIGRP

To configure the EIGRP routing protocol using Cisco SD-WAN Manager templates follow these steps:

1. Create an EIGRP feature template to configure EIGRP parameters.
2. Create a VPN feature template to configure VPN parameters for service-side routing (any VPN other than VPN 0 or VPN 512).
3. Create a device template and apply the templates to the correct devices.

Create an EIGRP Template

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

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

3. Click **Add Template** and select a device from the list.
4. From the **Other Templates** section, choose **EIGRP** and enter a name and a description for the template.

Basic Configuration

Click **Basic Configuration** to configure the local autonomous system (AS) number for the template.

Parameter Name	Description
Autonomous System ID *	Enter the local AS number. • **Range**: 1-65,535 • **Default**: None

Configure IP4 Unicast Address Family

To redistribute routes from one protocol (routing domain) into an EIGRP routing domain, click **New Redistribute** and enter the following parameter values:

Table 28: Redistribution Parameters

Parameter Name	Value	Description
Mark as Optional Row	Click **Optional** 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 *	Select 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** You can set metric values for redistribution using the CLI add-on feature template from Cisco IOS XE Catalyst SD-WAN Release 16.12.1b and later. Use the following command: `redistribute ospf 1 metric 1000000 1 1 1 1500` For more information, see CLI Add-on Feature Templates.
	static	Redistribute static routes into EIGRP.
Route Policy *	Enter the name of the route policy to apply to redistributed routes.	
Click **Add** to save the redistribution information.		

To advertise a prefix into the EIGRP routing domain, click **Network**, and then click **New Network** and enter the following parameter values:

Table 29: Configure Network

Parameter Name	Description
Mark as Optional Row	Click **Optional** 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. See Create a Template Variables Spreadsheet.
Network Prefix *	Enter the network prefix you want EIGRP to advertise in the format of *prefix/mask*.
Click **Add** to save the network prefix.	

Configure Advanced Parameters

To configure advanced parameters for EIGRP, click **Advanced** and configure the following parameter values:

Table 30: Advanced Parameters

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 65,535 • **Default**: 15 seconds
Hello Interval (seconds)	Set the interval at which the router sends EIGRP hello packets. • **Range**: 0 through 65,535 • **Default**: 5 seconds
Route Policy Name	Enter the name of an EIGRP route policy.

Configure Route Authentication Parameters

The IP Enhanced IGRP Route Authentication feature supports MD5 or HMAC-sha-256 authentication of routing updates from the EIGRP routing protocol. To configure authentication for EIGRP routes:

1. Click **Authentication**.
2. Click **Authentication** to open the **Authentication Type** field.
3. Choose **global** parameter scope.
4. From the drop-down list, choose **md5** or **hmac-sha-256**.

Parameter	Option	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 piece of information that is used to compute the HMAC and is known both by the sender and the receiver of the message.
Click **Add** to save the authentication parameters.		

Note To use a preferred route map, specify both an MD5 key (ID or auth key) and a route map.

Configure Interface Parameters

To configure interface parameters for EIGRP routes, click **Interface**, and enter the following parameter values:

Table 31: Interface Parameters

Parameter Name	Description
Mark as Optional Row	Click **Optional** 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.
Interface name	Enter the interface name(s) on which EIGRP should run.
Shutdown	**No** (the default) enables the interface to run EIGRP. **Yes** disables the interface.
Click **Add** to save the interfaces.	

Configure EIGRP Using CLI

Configure EIGRP on Cisco IOS XE Catalyst SD-WAN Devices

The following example shows the how to configure EIGRP on Cisco IOS XE Catalyst SD-WAN devices through CLI.

```
config-transaction
router eigrp vpn
 !
 address-family ipv4 unicast vrf 1 autonomous-system 100
  !
  topology base
   table-map foo filter
   redistribute omp
```

```
  exit-af-topology
  network 10.1.44.0 255.0.0.0
 exit-address-family
 !
 address-family ipv6 unicast vrf 1 autonomous-system 200
  !
  topology base
   table-map bar
   redistribute omp
  exit-af-topology
 exit-address-family
!
```

Example: Advertise EIGRP Routes to OMP

```
config-transaction
sdwan
 omp
  no shutdown
  graceful-restart
  address-family ipv4 vrf 1
   advertise eigrp
  !
  address-family ipv6 vrf 1
   advertise eigrp
  !
  address-family ipv4
   advertise connected
   advertise static
  !
 !
```

Verify EIGRP Configuration Using CLI

Configuration on Cisco IOS XE Catalyst SD-WAN Devices

The outputs of the following show commands show the EIGRP configuration on Cisco IOS XE Catalyst SD-WAN devices.

View IPv4 EIGRP routes on Cisco IOS XE Catalyst SD-WAN devices.

```
Device# show ip route vrf 1
m        192.168.22.22 [251/0] via 192.168.11.12, 00:28:00
      192.168.55.0/32 is subnetted, 1 subnets
D EX     192.168.55.55 [170/1] via 10.1.44.2, 00:33:58, GigabitEthernet3.2
      192.168.66.0/32 is subnetted, 1 subnets
B        192.168.66.66 [20/0] via 192.168.1.3, 00:33:57
      192.168.1.0/32 is subnetted, 3 subnets
D EX     192.168.1.3 [170/1] via 10.1.44.2, 00:33:58, GigabitEthernet3.2
m        192.168.1.33 [251/0] via 192.168.11.14 (3), 00:28:01
```

View IPv6 EIGRP routes on Cisco IOS XE Catalyst SD-WAN devices.

```
Device# show ipv6 route vrf 1
C   300:4::/64 [0/0]
     via GigabitEthernet3.2, directly connected
L   300:4::1/128 [0/0]
     via GigabitEthernet3.2, receive
D   2000:1:3::1/128 [90/1]
     via FE80::20C:29FF:FEF5:C767, GigabitEthernet3.2
L   FF00::/8 [0/0]
     via Null0, receive
```

```
cEdge4-Naiming#show ipv6 route vrf 1 2000:1:3::1/128
Routing entry for 2000:1:3::1/128
  Known via "eigrp 200", distance 90, metric 1
  OMP Tag 888, type internal
  Redistributing via omp
  Route count is 1/1, share count 0
  Routing paths:
    FE80::20C:29FF:FEF5:C767, GigabitEthernet3.2
      From FE80::20C:29FF:FEF5:C767
      Last updated 00:22:06 ago
```

View OMP routes in EIGRP on Cisco IOS XE Catalyst SD-WAN devices.

```
Device# show eigrp address-family ipv4 vrf 1 topology 192.168.44.4/24
EIGRP-IPv4 VR(vpn) Topology Entry for AS(100)/ID(192.168.1.44)
           Topology(base) TID(0) VRF(1)
EIGRP-IPv4(100): Topology base(0) entry for 192.168.44.4/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 1
  Descriptor Blocks:
  192.168.1.5, from Redistributed, Send flag is 0x0
      Composite metric is (1/0), route is External
      Vector metric:
        Minimum bandwidth is 0 Kbit
        Total delay is 0 picoseconds
        Reliability is 0/255
        Load is 0/255
        Minimum MTU is 0
        Hop count is 0
        Originating router is 192.168.1.44
      External data:
        AS number of route is 0
        External protocol is OMP-Agent, external metric is 4294967294
        Administrator tag is 0 (0x00000000)
```

Configure Routing Information Protocol (RIPv2) Using the CLI

You can configure RIPng using CLI device templates and CLI Add-on feature templates.

This section provides information about RIP configuration on Cisco IOS XE Catalyst SD-WAN devices.

Note

- Initial VRF routing table and address family submode configurations are required to verify RIP configurations using **show ip protocols** command.
- These commands can be run in any order.

- Configure the RIP routing process.

  Enable a RIP routing process and enter router configuration mode:

  ```
  Device# config-transaction
  Device(config)# router rip
  Device(config-router)#
  ```

- Configure the RIP VRF-aware support.

  Enter VRF address family configuration mode and enable IPv4 address prefixes:

```
Device(config)# router rip
Device(config-router)# address-family ipv4 vrf vrf-name
```

- Specify the RIP version.

  Specify RIP version as 2 to enable the device to send only RIP version 2 (RIPv2) packets:

  ```
  Device(config)# router rip
  Device(config-router)#  version {1|2}
  ```

- Configure RIP routes summarization

  Disable or restore the default behavior of automatic summarization of subnet routes into network-level routes used in router configuration mode:

  ```
  Device(config)# router rip
  Device(config-router)# auto-summary
  ```

- Validate the source IP address.

  Enable a router to perform validation checks on the source IP address of incoming RIP updates:

  ```
  Device(config)# router rip
  Device(config-router)# network ip-address
  Device(config-router)# validate-update-source
  ```

- Configure interpacket delay.

  Configure interpacket delay for outbound RIP updates, in milliseconds:

  ```
  Device(config)# router rip
  Device(config-router)# output-delay delay-value
  ```

- Redistribute the routes into the RIP routing process.

  Redistribute the specified routes into the IPv4 RIP routing process. We recommend the redistribution of protocols configuration only after configuring the source router protocols. The protocol argument can be one of these keywords—**bgp**, **connected**, **isis**, **eigrp**, **omp**, **ospf**, **ospfv3**, or **static**. In Cisco IOS XE Catalyst SD-WAN Release 17.7.1a, RIP Version 2 configurations in Cisco IOS XE Catalyst SD-WAN devices support OMP as a redistributed protocol.

  ```
  Device(config)# router rip
  Device(config-router)# redistribute protocol [metric metric-value] [route-map map-name]
  ```

- Filter the RIP-routing updates.

  Apply a prefix list to the RIP-routing updates that are received in or sent over an interface:

  ```
  Device(config)# router rip
  Device(config-router)# distribute-list prefix-list listname {in | out} [interface-type
   interface-number]
  ```

- Configure the RIP parameters.

  The network command is required to enable interfaces for RIP(v2), and to associate a network with a RIP routing process. There's no limit on the number of **network** commands that you can use on the router. For network configurations, we recommend that you use classful (Class A, Class B, Class C) IP network ID addressing.

  ```
  Device(config)# router rip
  Device(config-router)# network ip-address
  ```

  Define a neighboring device with which to exchange routing information:

```
Device(config)# router rip
Device(config-router)# network ip-address
Device(config-router)# neighbor ip-address bfd
```

Apply an offset list to routing metrics:

```
Device(config)# router rip
Device(config-router)# offset-list acl-number in offset[ interface-type |interface-name]
```

Adjust routing protocol timers:

```
Device(config)# router rip
Device(config-router)# timers basic update invalid holddown flush
```

- Customize a RIP.

  Define the maximum number of equal-cost routes that an IPv4 RIP can support:

  ```
  Device(config)# router rip
  Device(config-router)# maximum-paths number-paths
  ```

- Configure a route tag.

  By default, automatic RIPv2 route tag is enabled for redistributed OMP routes. When a router is installed by another Cisco IOS XE Catalyst SD-WAN device, the admin distance is set to 252 so that OMP routes are preferred over redistributed OMP routes:

  ```
  Device(config)# router rip
  Device(config-router)# omp-route-tag
  ```

- Configure the traffic.

  Configure traffic to use minimum-cost paths, and load splitting on multiinterfaces with equal-cost paths:

  ```
  Device(config)# router rip
  Device(config-router)# traffic-share min across-interfaces
  ```

Configuration Example

The following is a complete example of RIP configuration for Cisco IOS XE Catalyst SD-WAN devices using the CLI:

```
config-transaction
!
       vrf definition 172
       address-family ipv4
       exit-address-family
!
       router rip
       address-family ipv4 vrf 172
       distance 70
       omp-route-tag /* Default is enabled */
       default-information originate route-map RIP-MED
       version 2
       network 10.0.0.20    /* Only classful A, B, or C network. */
 distribute-list prefix v4KANYU-RIP in TenGigabitEthernet0/1/3.791
 redistribute rip v6kanyu metric 1 metric-type 1 route-map v6RED-RIP-OSPF1
 distribute-list prefix v4KANYU-RIP in TenGigabitEthernet0/1/3.792
 no auto-summary
!
```

Verify RIPv2 Configurations Using the CLI

You can verify RIP configurations using CLI or **IP Routes** window in Cisco SD-WAN Manager. The following is a sample output from the **show sdwan running | sec rip** command displaying the router RIP configurations:

```
Device# show sdwan running | sec rip
router rip
   version 2
   redistribute connected
   output-delay 20
   input-queue 20
!
 address-family ipv4 vrf 200
   redistribute connected
   redistribute omp metric 2
   network 56.0.0.0
   no auto-summary
   version 2
   exit-address-family
```

The following is a sample output from the **show ip route rip** command displaying RIP routes in the default routing table:

```
Device# show ip route rip
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is 10.0.5.13 to network 10.10.10.10

R     10.11.0.0/16 [120/1] via 172.16.1.2, 00:00:02, GigabitEthernet1
```

The following is a sample output from the **show ip route vrf** *vrf-id* **rip** command displaying RIP routes under the VRF table:

```
Device# show ip route vrf 1 rip
Routing Table: 1
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
H - NHRP, G - NHRP registered, g - NHRP registration summary
o - ODR, P - periodic downloaded static route, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
& - replicated local route overrides by connected

Gateway of last resort is not set
10.0.0.14/32 is subnetted, 1 subnets
R 10.14.14.14 [120/1] via 10.20.25.18, 00:00:18, GigabitEthernet5
```

The following is a sample output from the **show ip rip database** command displaying the contents of a RIP private database:

```
Device# show ip rip database
10.11.0.0/16    auto-summary
10.11.0.0/16
[1] via 172.16.1.2, 00:00:00, GigabitEthernet1
```

The following is a sample output from the **show ip rip neighbors** command displaying RIP Bidirectional Forwarding Detection (BFD) neighbors:

```
Device# show ip rip neighbors
BFD sessions created for the RIP neighbors
Neighbor      Interface         SessionHandle
10.10.10.2    GigabitEthernet1    1
```

The following is a sample output from the **show ip protocols** command using section RIP to display only is a RIP protocol configurations on the device:

```
Device# show ip protocols | sec rip
Routing Protocol is "rip"
  Outgoing update filter list for all interfaces is not set
  Incoming update filter list for all interfaces is not set
  Sending updates every 30 seconds, next due in 19 seconds
  Invalid after 180 seconds, hold down 180, flushed after 240
  Redistributing: rip
  Neighbor(s):
    10.1.1.2
  Default version control: send version 2, receive version 2
    Interface               Send  Recv  Triggered RIP  Key-chain
    GigabitEthernet1        2     2          No        none
    Loopback10              2     2          No        none
  Automatic network summarization is in effect
  Maximum path: 4
  Routing for Networks:
    10.11.0.1
  Routing Information Sources:
    Gateway         Distance      Last Update
    10.1.1.2             120      00:00:15
  Distance: (default is 120)
```

Configure RIPng Using the CLI

You can configure RIPng using CLI device templates and CLI Add-on feature templates.

This section provides information about RIPng configuration on Cisco IOS XE Catalyst SD-WAN devices.

Note Initial VRF routing table and address family submode configurations are required to verify RIP configurations using the **show ipv6 route vrf** command.

1. Configure IPv6 RIPng VRF-aware support.
   a. Enable VRF-aware support for IPv6 RIPng routing. It is mandatory for the RIPng to be configured within the service VPN.

      ```
      Device(config)# ipv6 rip vrf-mode enable
      ```

   b. Enable the forwarding of IPv6 unicast datagrams:

```
Device(config)# ipv6 unicast-routing
```

2. Configure the IPv6 RIPng routing process and enable router configuration mode for the IPv6 RIPng routing process:

Note For *ripng-instance*, use sdwan.

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)#
```

3. Enter VRF address family configuration mode and enable IPv6 address prefixes:

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)# address-family ipv6 vrf vrf-name
Device(config-ipv6-router-af)#
```

4. Define an administrative distance for routes that are inserted into a routing table:

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)# address-family ipv6 vrf vrf-name
Device(config-ipv6-router-af)# distance distance
```

5. Configure a route tag.

   By default, automatic RIPng route tag is enabled for redistributed OMP routes. When a Cisco IOS XE Catalyst SD-WAN device learns a RIPv2 and RIPng route with a unique SD-WAN tag (44270), the router installs the route with an administrative distance of 252, which is higher than the OMP distance (251), so that the OMP routes are preferred over redistributed OMP routes:

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)# omp-route-tag
```

6. Create an entry in the IPv6 prefix list:

```
Device(config)# ipv6 prefix-list list-name [seq seq-number] permit IPv6 prefix
(IP/length)
```

7. Apply a prefix list to IPv6 RIPng routing updates that are received or sent on an interface:

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)# distribute-list prefix-list prefix-list-name {in | out}
[interface-type | interface-number]
```

8. Redistribute the specified routes into the IPv6 RIPng routing process. The **rip** keyword and *ripng-instance* specify an IPv6 RIPng routing process.

```
Device(config)# ipv6 router rip ripng-instance
Device(config-rtr)# redistribute protocol [metric default-metric] [route-map
map-tag]
```

9. Configure the interface.

   a. Enable the specified IPv6 RIPng routing process on an interface:

Note For *ripng-instance*, use sdwan.

```
Device(config)# interface type number
Device(config-if)# ipv6 rip ripng-instance enable
```

b. (Optional) The IPv6 default route (::/0) distributes into the specified RIPng routing process updates sent out of the specified interface:

Note For *ripng-instance*, use sdwan.

```
Device(config)# interface type number
Device(config-if)# ipv6 rip ripng-instance default-information {only |
originate} [metric metric-value]
```

c. Set the IPv6 RIPng metric-offset for an interface.

Note For *ripng-instance*, use sdwan.

```
Device(config)# interface type number
Device(config-if)# ipv6 rip ripng-instance metric-offset metric-value
```

d. Configure the IPv6 RIPng to advertise summarized IPv6 addresses on an interface and to specify the IPv6 prefix that identifies the routes to be summarized.

Note For *ripng-instance*, use sdwan.

```
Device(config)# interface type number
Device(config-if)# ipv6 address {ipv6-prefix/prefix-length | prefix-name |
sub-bits/prefix-length}
Device(config-if)# ipv6 rip ripng-instance summary-address
{ipv6-prefix/prefix-length}
```

Configuration Example for RIPng

The following example shows a complete RIPng configuration for Cisco IOS XE Catalyst SD-WAN devices using the CLI:

```
config-transaction
!
   vrf definition 1
   address-family ipv6
   exit-address-family
!
  ipv6 rip vrf-mode enable
  ipv6 unicast-routing
!
  ipv6 prefix-list cisco seq 10 permit 2000:1::/64
!
```

```
   ipv6 router rip sdwan
      address-family ipv6 vrf 1
         distance 130
         omp-route-tag
         distribute-list prefix-list cisco in GigabitEthernet0/0/0
         redistribute omp metric 10
         exit-address-family
 !
    interface GigabitEthernet0/0/0
        ipv6 address 2001:DB8::/64
        ipv6 rip sdwan enable
        ipv6 rip sdwan default-information originate
        ipv6 rip sdwan metric-offset 10
        ipv6 rip sdwan summary-address 2001:90::1/32
!
```

Verify RIPng Configurations Using the CLI

The following is a sample output from the **show ipv6 route vrf** command displaying the router RIPng configurations:

```
Device# show ipv6 route vrf 1
IPv6 Routing Table - 1 - 11 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, ls - LISP destinations-summary, a - Application
       m - OMP
R   1100::/64 [120/2]
     via FE80::20C:29FF:FE2E:13FF, GigabitEthernet2
R   2000::/64 [120/2]
     via FE80::20C:29FF:FE51:762F, GigabitEthernet2
R   2001:10::/64 [120/2]
     via FE80::20C:29FF:FE82:D659, GigabitEthernet2
R   2500::/64 [252/11], tag 44270
     via FE80::20C:29FF:FEE1:5237, GigabitEthernet2
C   2750::/64 [0/0]
     via GigabitEthernet2, directly connected
L   2750::1/128 [0/0]
     via GigabitEthernet2, receive
R   2777::/64 [252/11], tag 44270
     via FE80::20C:29FF:FEE1:5237, GigabitEthernet2
m   2900::/64 [251/0]
     via 192.168.1.5%default
R   3000::/64 [120/2]
     via FE80::20C:29FF:FE2E:13FF, GigabitEthernet2
R   3400::/64 [252/11], tag 44270
     via FE80::20C:29FF:FE51:762F, GigabitEthernet2
L   FF00::/8 [0/0]
     via Null0, receive
```