



Overlay Management Protocol

- [Feature history for OMP, on page 1](#)
- [OMP routing mechanisms for Cisco SD-WAN overlay networks, on page 2](#)
- [OMP route advertisements, on page 3](#)
- [OMP paths, on page 8](#)
- [OMP route redistribution, on page 9](#)
- [OMP graceful restart, on page 16](#)
- [OMP vRoute advertisement optimization using system path, on page 18](#)
- [Configure OMP, on page 19](#)

Feature history for OMP

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

Table 1: Feature History

Feature Name	Release Information	Description
OMP Route Aggregation	Cisco IOS XE Catalyst SD-WAN Release 17.13.1a 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 redistribution is requested.

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 .
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.
OMP RIB-Out Optimization using TLOC Path	Cisco IOS XE Catalyst SD-WAN Release 17.18.2	This feature introduces an optimization to the OMP protocol by flattening vRoute RIB-Outs from a per-TLOC (Transport Locator) basis to a per-SysIP (System IP) basis. This significantly reduces memory consumption on Cisco Catalyst SD-WAN Controllers, improves overall scalability, and enhances convergence performance in large-scale SD-WAN deployments.

OMP routing mechanisms for Cisco SD-WAN overlay networks

The Cisco Catalyst SD-WAN Overlay Management Protocol (OMP) is a control protocol that

- establishes and maintains the Cisco Catalyst SD-WAN control plane
- exchanges routing, policy, and management information between Cisco Catalyst SD-WAN Controllers and Cisco IOS XE Catalyst SD-WAN devices in the overlay network
- orchestrates overlay network communication, including connectivity among network sites, service chaining, and VPN or VRF topologies, and
- distributes service-level routing information, related location mappings, data plane security parameters, and routing policy.

Operational details and architectural role of OMP

The Cisco IOS XE Catalyst SD-WAN devices automatically initiate OMP peering sessions between themselves. The two IP end points of the OMP session are the system IP addresses of the two devices.

OMP is a comprehensive 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. Edge devices then use OMP to receive routing, security, services, and policies for data plane connectivity and transport.

Limitations for OMP

OMP configuration

Multitenant Cisco Catalyst SD-WAN Controllers only support global OMP configuration.

Path limit

The number of paths that are shared is dependent upon factors such as memory and the organization of internal data structure.

OMP route advertisements

This topic describes OMP route advertisements and the types of routes that OMP advertises to.

On Cisco Catalyst SD-WAN Controllers and Cisco IOS XE Catalyst SD-WAN devices, OMP advertises to its peers and services the routes it has learned from its local site, along with their corresponding transport location mappings, which are called TLOCs. These routes, known as OMP routes or vRoutes, are tuples consisting of the route and its associated TLOC. The Cisco Catalyst SD-WAN Controllers learn the topology of the overlay network and the services available in the network through OMP routes.

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.

In an overlay networking environment, OMP uses a different notion of routing peers than in traditional networks. Logically, the overlay environment consists of a centralized controller and several 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 do not advertise routing information to each other using OMP or any other method. The OMP peering sessions between the centralized controller and the edge devices are exclusively for exchanging control plane traffic. They do not transmit data traffic.

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

When route maps are configured for protocol redistribution, AS path and community configuration—such as AS path prepend—are not supported. To configure and apply the AS path for redistributed OMP routes, use 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 applying the configuration at the global level or at the specific VPN level. To configure route advertisements to OMP at the global level, use the OMP feature template. On the other hand, to configure route advertisements to OMP at the specific VPN level, use the VPN feature template. For more information about configuring route advertisements to OMP, see [Configure OMP using templates](#).



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



Note In releases prior to Cisco IOS XE Catalyst SD-WAN Release 17.18.1a, controllers experience high CPU and memory usage in large-scale environments. It is recommended to avoid running commands that generate large outputs. Such commands can place substantial load on the system, potentially causing temporary delays or disruptions in command execution. High resource utilization in these situations may result in longer convergence times or, in some cases, process instability.

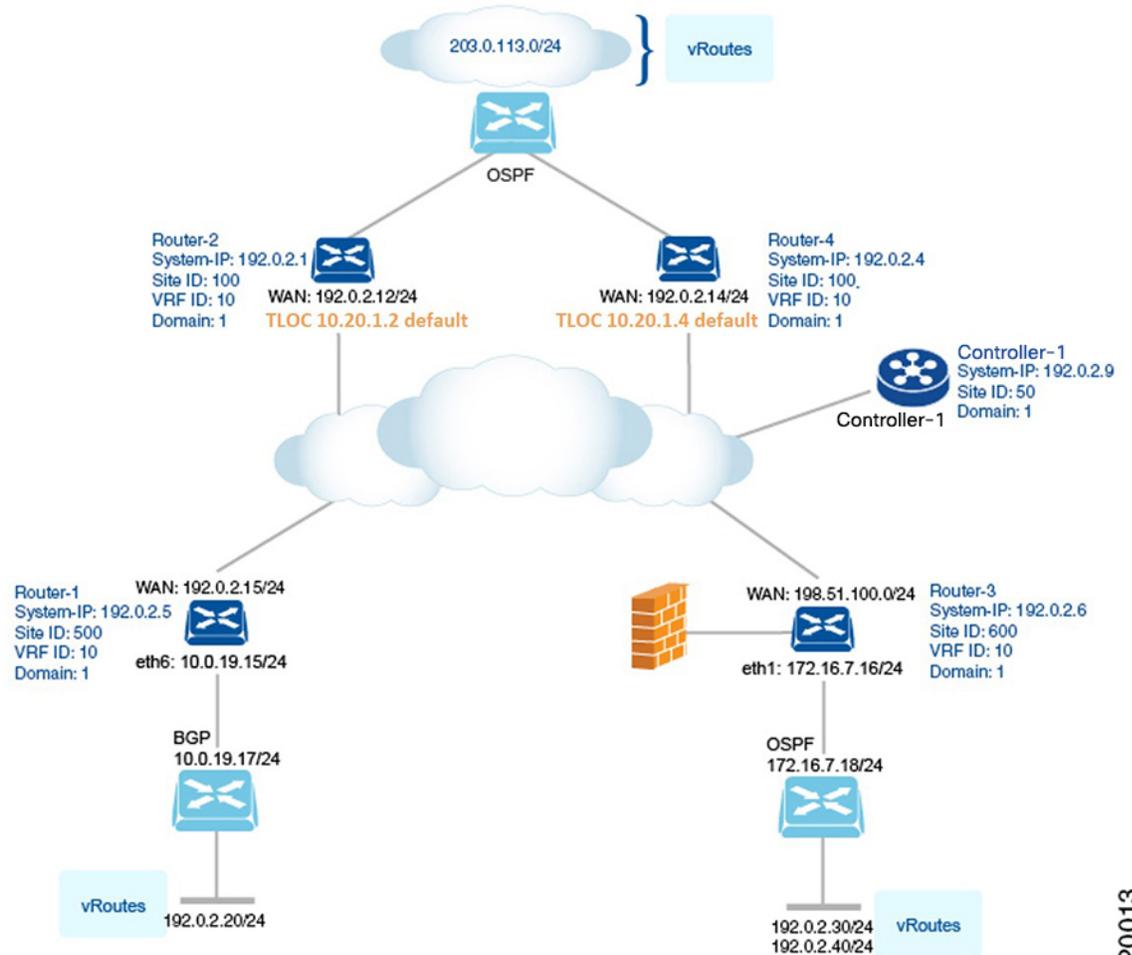
The system is designed to recover automatically once CPU and memory activity return to normal. To minimize operational impact, monitor system resources and schedule commands thoughtfully.

OMP advertises these 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.

This figure illustrates the two types of OMP routes.

Figure 1: Different types of OMP routes



520013

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 IOS XE 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 these 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.
- System IP: System IP address of the OMP speaker that originates the OMP route. Only applicable to optimized advertisement. For more information, see [OMP vRoute advertisement optimization using system path, on page 18](#).
- 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 these 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.

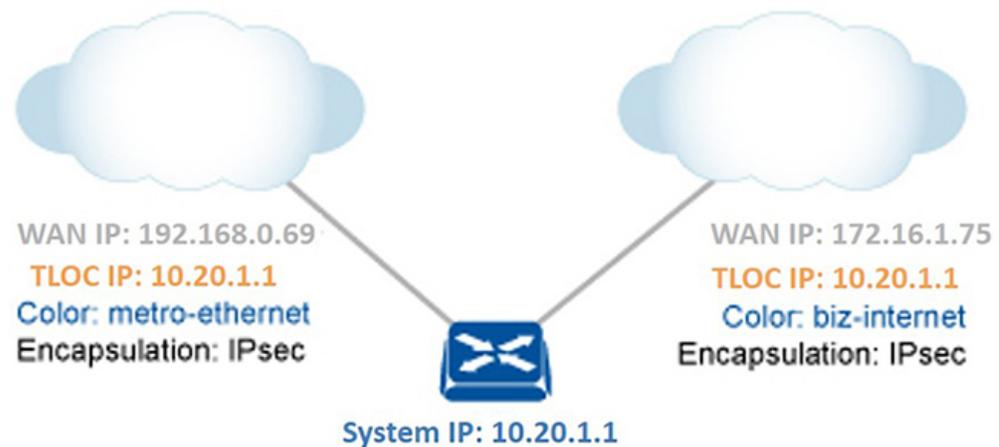
- Preference: Degree of preference that is used to differentiate between TLOCs that advertise the same OMP route.
- Site ID: Identifier of a site within the Cisco Catalyst SD-WAN overlay network domain to which the TLOC belongs.
- Tag: Optional, transitive path attribute that an OMP speaker can use to 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 the one 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. 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. 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](#).

368487

OMP paths

The transport location (TLOC) information is advertised to the OMP peers including Cisco Catalyst SD-WAN Controller and its local-site branches. 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.

Configure path limit using CLI commands

Configure the number of paths a Cisco Catalyst SD-WAN Controller can send to another Cisco Catalyst SD-WAN Controller.

Before you begin

Ensure **OMP send-path-limit** is set to a value equal to or greater than the number of TLOCs used for any advertised prefix.

Procedure

- Step 1** Create a CLI add-on profile or a CLI add-on template.
- Step 2** Enter the **omp** mode.
- Step 3** Use the **controller-send-path-limit** command to configure the send path limit to be exchanged between Cisco Catalyst SD-WAN Controllers.

You can configure a maximum of 128 send path limit.

```
# controller-send-path-limit <path-limit>
```

Example:

```
# controller-send-path-limit 100
```

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 route redistribution determines which routes OMP automatically learns and advertises, and which routes require explicit configuration for redistribution to prevent routing issues.

OMP automatically redistributes these routes after learning it 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 these routes require explicit configuration:

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

Advertise network

The **advertise network** *<ipv4-prefix>* command advertises a specific prefix when a non-OMP route corresponding to the prefix is present in the VRF IPv4 routing table.



Note This command is only supported for **address-family ipv4**.

This 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 route origin type and sub-type

OMP sets the origin and sub-origin type in each OMP route to indicate the route's origin. The Cisco Catalyst SD-WAN Controller and the router consider the origin type and subtype when selecting routes.

OMP route origin type	OMP route origin subtype
BGP	External Internal
Connected	—
OSPF	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
OSPFv3	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
Static	—
EIGRP	<ul style="list-style-type: none"> • 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.

Configure the redistribution of OSPF routes using CLI commands

Configure the redistribution of OSPF routes into OMP for VRF1.

Procedure

Step 1 Configure `advertise ospf route-map <route-map-name> external`.

The OSPF internal routes are redistributed into OMP by default without any explicit configuration.

Example:

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

```
omp
  no shutdown
  ecmp-limit      6
  graceful-restart
  no as-dot-notation
  timers
    holdtime      300
    graceful-restart-timer 120
  exit
```

```

address-family ipv4
  advertise ospf external <-- This configuration implies OSPF Inter-Area/Intra-Area routes & External
  routes are redistributed into OMP
  advertise connected
  advertise static
!
```

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

Example:

```

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

Step 2 Use the **external** keyword in the configuration to apply the supplied route-map to both external and internal OSPF routes (Intra-Area/Inter-Area).

Example:

This example shows the redistribution of OSPFv3 external routes:

```

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

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

Administrative distance

Administrative distance is a metric that

- is used to select the best path when there are two or more different routes to the same destination from multiple routing protocols, and
- is used by Cisco Catalyst SD-WAN Controller or the router to prefer the OMP route to a destination with the lowest administrative distance value.

This table lists the default administrative distances used by the Cisco IOS XE Catalyst SD-WAN devices:

Protocol	Administrative distance
Connected	0
Static	1
NAT (NAT and static routes cannot coexist in the same VPN; NAT overwrites static routes)	1
Learned from DHCP	1
EIGRP Summary	5
EBGP	20
EIGRP	Internal: 90, External: 170
OSPF	110
OSPFv3	110
IS-IS	115
IBGP	200
OMP	251

OMP best path algorithm

Cisco IOS XE 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 IOS XE Catalyst SD-WAN devices use this algorithm to choose the best path:

Table 2: Best path algorithm

Step	Applies to	Description
1	Edge devices Cisco Catalyst SD-WAN Controller	Path validity Checks whether the OMP path is valid. If not, ignores it.

Step	Applies to	Description
2	Edge devices Cisco Catalyst SD-WAN Controller	Active vs. stale paths Prefers an active path over a stale path. An active path is the 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 is only 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 Selects 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 Catalyst SD-WAN Controller via OMP. For information about administrative distance, see Administrative distance, on page 11 .
4	Edge devices Cisco Catalyst SD-WAN Controller	OMP path preference Selects the OMP path with the higher OMP path preference value.
5	Cisco Catalyst SD-WAN Controller	Access region Cisco Catalyst SD-WAN Controller drops advertisement from border router (BR) to BR in the same region.
6	Edge devices	Core region Cisco Catalyst SD-WAN Controller allows advertisement between BRs in the same access region, but receives BR drops advertisement.
7	Multi-Region Fabric scenario only Edge devices	Region path length Compares region-path-length and prefers lower. If region-path-length-ignore is configured, then skips this step. (This addresses secondary regions in Multi-Region Fabric.)
8	Multi-Region Fabric scenario only Border routers	Access region vs. core region Prefers access region paths over core region paths.

Step	Applies to	Description
9	Edge devices	<p>Direct vs. transport gateway path</p> <p>Prefers a direct path over a transport gateway path.</p> <p>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 <i>Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide</i>.</p>
10	Multi-Region Fabric scenario only Edge devices	<p>Multi-Region Fabric subregion comparison</p> <ul style="list-style-type: none"> • Prefers paths from the router's own subregion. • When comparing two paths that are not from the router's subregion, prefers a path that is not part of any subregion.
11	Multi-Region Fabric scenario only Edge devices	<p>Border router preference</p> <p>Prefers a path with a higher border router preference value.</p>
12	Edge devices	<p>Derived affinity</p> <p>Prefers a path with a lower derived affinity value.</p>
13	Edge devices with an affinity preference configured	<p>Affinity preference</p> <p>Depending on the affinity preference configured on the device, prefers a path whose affinity is earlier in the preference list (higher priority). If the device uses affinity-preference-auto, then it prefers a path with a numerically lower affinity group.</p> <p>Note When comparing two paths with similar reorigination types, one with an affinity value and one without, prefers the path with an affinity value.</p>
14	Edge devices	<p>TLOC preference</p> <p>Select an OMP path with a higher TLOC preference value.</p> <p>Note With TLOC preference and AAR policy configured, outbound and inbound traffic may follow different paths when the preferred TLOC goes out of SLA. For outbound traffic, tunnels in SLA will be preferred regardless of TLOC preference; however, TLOC preference still dictates inbound path selection.</p>

Step	Applies to	Description
15	Edge devices Cisco Catalyst SD-WAN Controller	Origin type and subtype Compares the origin type and subtype, and selects the first match from this list: <ul style="list-style-type: none"> • 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 Selects an OMP path that has a lower origin metric.
17	Cisco Catalyst SD-WAN Controller	Path source Prefers 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 that are selected as best-paths and accepted by policy, advertise no more than the number of paths specified in send-path-limit.

Choosing the best path

Examples for choosing the best path:

Control connection failover and route advertisement

In a setup with two WAN Edge devices and four Cisco SD-WAN Controllers:

WAN Edge 1 forms control connections with two controllers by default (e.g., Controller 1 and Controller 2). If one of these controllers fails (e.g., Controller 1), WAN Edge 1 automatically reconnects to a backup controller (e.g., Controller 3) while maintaining its session with Controller 2.

After failover, route advertisements may change. If WAN Edge 2 originates prefix A and is connected to Controller 2, Controller 2 may not advertise prefix A to WAN Edge 1 if WAN Edge 1 now learns this prefix through Controller 3.

Thus, WAN Edge 1 only receives prefix A from Controller 3.

Day 1 Expected Behavior: A Cisco SD-WAN Controller does not re-advertise a route learned from another Cisco SD-WAN Controller to a WAN Edge device if that device already receives the same route directly from a Cisco SD-WAN Controller.

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

Best path selection

When a Cisco SD-WAN Controller receives an OMP path (for example, 10.10.10.0/24) via OMP from a Cisco IOS XE Catalyst SD-WAN device (with an origin code of OSPF) and also from another Controller (with the same origin code), the best-path algorithm selects the path that the Cisco IOS XE Catalyst SD-WAN device sends directly, assuming all other factors are equal.

When a Controller learns the same OMP path (such as 10.10.10.0/24) from two Cisco IOS XE Catalyst SD-WAN devices at the same site, it chooses both paths and advertises them to other OMP peers, as long as all parameters are equal. By default, the Controller advertises up to four equal-cost paths.



Note 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 SD-WAN Controller removes from the forwarding table all the OMP paths that point to that TLOC.

OMP graceful restart

OMP graceful restart

OMP graceful restart is a control plane resiliency mechanism that

- allows Cisco IOS XE Catalyst SD-WAN devices to continue forwarding data traffic when the Cisco Catalyst SD-WAN Controller is unavailable

- uses cached OMP information (such as routes, TLOCs, service routes, and policies) to maintain data plane operations, and
- synchronizes updated network information when the controller connection is restored.

When OMP graceful restart is enabled, both Cisco IOS XE Catalyst SD-WAN devices and Cisco Catalyst SD-WAN Controllers cache OMP information received from their peers. This cache includes OMP routes, TLOC routes, service routes, IPsec SA parameters, and centralized data policies.

How graceful restart events work

Summary

OMP graceful restart enables devices and controllers to maintain forwarding operations and synchronize network state after connectivity is restored.

The key components involved in the process are:

- Cisco IOS XE Catalyst SD-WAN device: Maintains cached OMP information and forwards data traffic when the controller is unavailable.
- Cisco Catalyst SD-WAN Controller: Maintains cached OMP information and resumes synchronization when device connectivity is restored.
- OMP session: Facilitates communication and state synchronization between devices and controllers.

Workflow

These stages describe how OMP graceful restart enables devices and controllers to continue forwarding traffic using cached information during control plane outages and to resynchronize network state when connectivity is restored.

1. A device detects loss of OMP connection to a Cisco SD-WAN Controller and continues forwarding data traffic using cached OMP information.
2. The device periodically checks whether the controller is available.
3. When the controller becomes available again and the OMP session is re-established, the device flushes its local cache and accepts only the new OMP information from the controller.
4. The same process occurs in reverse if a Cisco Catalyst SD-WAN Controller loses connection to a device: the controller uses cached information until the device is available again, then updates its state upon reconnection.

Result

Data traffic continues to be forwarded using cached information during control plane outages, and devices/controllers automatically resynchronize when connectivity is restored.

OMP graceful restart timers

This section explains how OMP handles graceful restart timers configuration.

Each OMP peer independently configures its graceful restart timer on both Cisco IOS XE Catalyst SD-WAN devices and Cisco SD-WAN Controllers.

For example:

- If a controller is set to 300 seconds (5 minutes) and a device to 600 seconds (10 minutes), the controller retains OMP routes from the device for 10 minutes (per device's timer), and the device retains routes from the controller for 5 minutes (per controller's timer).
- The timer value is communicated during OMP session setup and determines how long cached routes are considered valid during peer loss.



Note When you change OMP graceful restart configuration, the OMP session between Cisco SD-WAN Controller and the device is intentionally reset (flapped). This action withdraws and relearns OMP routes for all address families (such as TLOC, IPv4/IPv6 unicast, IPv4 multicast, etc.) within a few seconds. During this period, Bidirectional Forwarding Detection (BFD) sessions will also flap momentarily. This is the expected behavior.

OMP vRoute advertisement optimization using system path

Traditionally, OMP on the Cisco SD-WAN Controller creates a vRoute advertisement (RIB-Out) for each TLOC (Transport Locators) for every vRoute NLRI prefix it learns from a device. If a device has multiple TLOCs, the number of advertisements increases. When all TLOC paths are treated equally in the control plane, the advertisements become duplicates. This duplication consumes excessive memory on the Cisco SD-WAN Controller. As a result, devices can experience out-of-memory issues, reboots, slower network convergence, and reduced SD-WAN scalability.

The OMP RIB-Out optimization using the TLOC path feature changes OMP behavior. Instead of creating vRoute RIB-Outs on a per-TLOC basis, the system now generates them based on each SysIP (System IP). This process significantly reduces memory consumption on Cisco SD-WAN Controllers, increases scalability, and improves convergence performance in large-scale SD-WAN deployments.

Benefits

1. **Flattened RIB-Outs:** Instead of sending multiple RIB-Outs for each TLOC associated with a prefix, only one flat RIB-Out per system IP is sent to peers.
2. **Memory Reduction:** This change drastically reduces the number of RIB-Outs generated. The memory consumption per prefix decreases by a factor of N, where N is the number of source TLOCs. For example, when a prefix with two TLOCs sent to two Cisco SD-WAN Controller peers, the system generates two RIB-Outs instead of four.
3. **Scope:** This optimization applies only to vRoute IPv4 and IPv6 address families in Cisco IOS XE Catalyst SD-WAN Release 17.18.2.
4. **Backward Compatibility:**
 - The feature introduces a new OMP protocol version (version 4) for peers understanding the optimization, while older peers (version 3) will continue to use the legacy per-TLOC format.

- New Cisco SD-WAN Controllers and Cisco IOS XE Catalyst SD-WAN devices can communicate with old devices by dynamically adjusting the RIB-Out format (sending legacy per-TLOC to old peers).
- Control policies that differentiate traffic based on TLOC, color, or service attributes disable this optimization for the affected routes. The system falls back to the legacy per-TLOC RIB-Outs to ensure policy adherence.

Configure OMP

Configure OMP using a configuration group

Before you begin

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

Procedure

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

Step 2 Create and configure a OMP feature in a System profile.

- a) Configure these fields for **Basic Configuration** section.

Table 3: Basic Configuration

Field	Description
Graceful Restart Enable	Enable graceful restart. By default, the graceful restart for OMP is enabled.
Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. A Cisco IOS XE Catalyst SD-WAN device advertises routes to Cisco Catalyst SD-WAN Controllers, and the controllers redistribute 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
BGP AS Path Auto-Translation	Enables automatic translation of the BGP as path length into the OMP preference value for BGP-learned routes. When this option is enabled, routes with shorter BGP aspaths are assigned a higher OMP preference. Specifically, the OMP preference for a BGP-learned route starts at 255 and is decremented by 1 for each AS in the BGP aspath. As a result, routes with fewer AS hops receive a higher preference within OMP route selection.

Field	Description
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco Catalyst SD-WAN Controller that can be installed in the local route table of the Cisco IOS XE Catalyst SD-WAN device. By default, a 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
Advertisement Interval (In Second)	Specify the time between OMP update packets. Range: 0 through 65535 seconds Default: 1 second We recommend you to configure 5 seconds on edge devices and 20 seconds on vSmart.
Hold Time(In Second)	Specify how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed. Range: 0 through 65535 seconds Defaults, by Cisco Catalyst SD-WAN Control Components release: <ul style="list-style-type: none"> • 20.18.x and later: 300 seconds • 20.16.x: 5400 seconds • 20.12.1 to 20.15.x: 300 seconds • Before 20.12.1: 60 seconds Defaults, by Cisco IOS XE Catalyst SD-WAN release: <ul style="list-style-type: none"> • 17.18.1 and later: 300 seconds • 17.16.x: 5400 seconds
EOR Timer(In Second)	Specify how long to wait after an OMP session has gone down and then come back up to send an end-of-RIB (EOR) marker. After this marker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. Range: 1 through 3600 seconds (1 hour) Default: 300 seconds (5 minutes)
Overlay AS	Specify a BGP AS number that OMP advertises to the BGP neighbors of the router.
Shutdown	Enable this option to disable OMP and disable the Cisco Catalyst SD-WAN overlay network. OMP is enabled by default.
OMP Admin Distance Ipv4	To advertise a route over OMP, configure the OMP administrative distance for the IPv4 address lower than the leaked route administrative distance. Range: 1 through 255

Field	Description
OMP Admin Distance Ipv6	To advertise a route over OMP, configure the OMP administrative distance for the IPv6 address lower than the leaked route administrative distance. Range: 1 through 255

b) Configure Timers.

Table 4: Timers

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

c) Configure these fields for **Advertise** section.

Table 5: Advertise

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

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

d) Configure these fields for Best Path.

Table 6: Best Path

Field	Description
Treat Hierarchical and Direct Paths Equally	<p>(Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.13.1a)</p> <p>In a Multi-Region Fabric scenario, if using secondary regions, enable this option to enable packets to use all available paths rather than only direct paths.</p> <p>By default, when a direct path is available to reach a destination, the overlay management protocol (OMP) enables only the direct path to the routing forwarding layer because the direct path uses fewer hops. This logic is part of route optimization. The result is that the forwarding layer, which includes application-aware routing policy, can only use the direct path.</p> <p>Treat Hierarchical and Direct Paths Equally disables this comparison of the number of hops so that traffic can use either the direct secondary-region path (fewer hops) or the primary-region path (more hops). When you disable the comparison of the number of hops, OMP applies equal-cost multi-path routing (ECMP) to all routes, and packets can use all available paths.</p>
Transport Gateway Path Behavior	<p>(Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.13.1a)</p> <p>Choose one of the following:</p> <ul style="list-style-type: none"> • Prefer Transport Gateway Path: For devices that can connect through a transport gateway, use only the transport gateway paths, even if other paths are available. • Do ECMP Between Direct and Transport Gateway Paths: For devices that can connect through a transport gateway and through direct paths, apply ECMP to all available paths.

Field	Description
Site Type	(Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.13.1a) If you configure a value for Transport Gateway Path Behavior , this field appears. Optionally, choose one or more site types to apply the transport gateway path behavior only to those site types.

What to do next

Also see [Deploy a configuration group](#).

Configure OMP using templates

Use the OMP template to configure OMP parameters for all Cisco 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, SD-WAN Manager NMSs, and Cisco Catalyst SD-WAN Controllers. You do not need to explicitly enable OMP. OMP must be operational for the Cisco SD-WAN overlay network to function. If you disable it, the overlay network also gets disabled.

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 [OMP Advertisements](#).

Cisco IOS XE Catalyst SD-WAN device use VRFs in place of VPNs. However, the steps desciebed in this section are still applicable for configuring Cisco IOS XE Catalyst SD-WAN devices through SD-WAN Manager. When you complete the configuration, the system automatically maps the VPN configurations to VRF configurations.

Procedure

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

Step 2 Click **Device Templates**.

Note

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

Step 3 Click **Create Template**.

Step 4 From the **Create Template** drop-down list, choose **From Feature Template**.

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

Step 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.

Step 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.

Step 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 these:

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

What to do next

See [Configure OMP options, on page 24](#).

Configure OMP options

Procedure

Step 1 To configure basic OMP options, click **Basic Configuration** and configure these parameters. All parameters are optional.

Table 7: Basic OMP options

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.

Parameter name	Description
Graceful Restart Timer	Specify how often the OMP information cache is flushed and refreshed. A timer value of 0 disables OMP graceful restart. Range: 0 to 604800 seconds (168 hours, or 7 days) Default: 43200 seconds (12 hours)
Number of Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. A Cisco IOS XE Catalyst SD-WAN device can have up to eight TLOCs, and by default advertises each route-TLOC tuple to the Cisco Catalyst SD-WAN Controller. If a local site has two Cisco 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(s) are advertised. Range: 1 to 16 Default: 4
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco Catalyst SD-WAN Controller that can be installed in the Cisco IOS XE Catalyst SD-WAN device 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 to 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 device. 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.

Step 2 Click **Save**.

Step 3 To configure OMP timers, click **Timers** and configure these parameters:

Table 8: OMP Timers

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

Parameter name	Description
Hold Time	<p>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.</p> <p>Range: 0 to 65535 seconds</p> <p>Defaults, by Cisco Catalyst SD-WAN Control Components release:</p> <ul style="list-style-type: none"> • 20.18.x and later: 300 seconds • 20.16.x: 5400 seconds • 20.12.1 to 20.15.x: 300 seconds • Before 20.12.1: 60 seconds <p>Defaults, by Cisco IOS XE Catalyst SD-WAN release:</p> <ul style="list-style-type: none"> • 17.18.1 and later: 300 seconds • 17.16.x: 5400 seconds
EOR Timer	<p>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.</p> <p>Range: 1 to 3600 seconds (1 hour)</p> <p>Default: 300 seconds (5 minutes)</p>

Step 4 Click **Save**.

Step 5 To advertise routes learned locally by the Cisco IOS XE Catalyst SD-WAN device to OMP, click **Advertise** and configure these parameters:

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

Table 9: OMP Advertisements

Parameter name	Description
Advertise	<p>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:</p> <ul style="list-style-type: none"> • BGP: Click On to advertise BGP routes to OMP. By default, BGP routes are not advertised to OMP. • Connected: Click Off to disable advertising connected routes to OMP. By default, connected routes are advertised to OMP. • OSPF: Click On and click On again in the External field that appears to advertise external OSPF routes to OMP. OSPF inter-area and intra-area routes are always advertised to OMP. By default, external OSPF routes aren't advertised to OMP. • Static: Click Off to disable advertising static routes to OMP. By default static routes are advertised to OMP. <p>To configure per-VPN route advertisements to OMP, use the VPN feature template.</p>

Step 6 Click **Save**.

Configure OMP using the CLI commands

Follow these procedures to configure OMP using CLI commands.

Configure OMP graceful restart using CLI commands

OMP graceful restart is enabled on all Cisco IOS XE Catalyst SD-WAN devices and Cisco Catalyst SD-WAN Controllers. OMP graceful restart timer 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). You can modify this default timer value using CLI commands.

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:

- 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

The graceful restart timer is set up independently on each OMP peer that is, Cisco IOS XE Catalyst SD-WAN device and Cisco Catalyst SD-WAN Controller. Consider a Cisco Catalyst 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, the 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 the SD-WAN Controller during the setup of the OMP session. The SD-WAN device retains the routes it learns from the SD-WAN Controller for 5 minutes, which is the default graceful restart

time value that is used on the SD-WAN Controller and that the controller sent to the device, also during the setup of the OMP session.

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

Procedure

Step 1 To modify the default timer value, enter the global configuration mode:

Example:

```
Device# config-transaction
Device(config)# sdwan
```

Step 2 Enter the **timers graceful-restart-timer** command and specify the time in seconds.

Example:

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

Step 3 To disable OMP graceful restart, use this command:

Example:

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

Configure OMP route advertisement using CLI commands

Enable protocol route advertisements to OMP for all VRFs on a Cisco IOS XE Catalyst SD-WAN device.

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 to OMP for Cisco Catalyst SD-WAN Controller, that is responsible for the device's domain. You can use the **advertise** command to have the device advertise these routes to OMP, consequently to SD-WAN Controller.



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.

Procedure

Step 1 To enable protocol route advertisements for OMP protocol for all VRFs, add the configuration at the global level.

Example:

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

Step 2 To enable protocol route advertisements for a few VRFs, remove the global-level configuration using **no advertise bgp** command.

Example:

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

Step 3 Then, add a per-VRF-level configuration:

Example:

```
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 neither at the global level nor at the VRF level.

Step 4 Next, configure the routes the device advertises to OMP for all VRFs configured on the device:

Example:

```
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.

Step 5 To configure the routes that the device advertises to OMP for a specific VRF on the device, use these command:

Example:

```
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

```

Step 6 For individual VRFs, routes from the specified prefix can be aggregated after advertising them into OMP using the `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:

Example:

```

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.

Configure BGP AS Path propagation into OMP using CLI commands

You can enable Cisco IOS XE Catalyst SD-WAN device to advertise BGP AS path information into OMP, ensuring that devices in the service-side network can receive and utilize this information for loop prevention. Propagating BGP AS path information helps to prevent BGP routing loops by allowing routers to identify and avoid routes that contain their own AS number in the path. It also provides greater visibility into the routing path.

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.

Procedure

Step 1 Enter the global configuration mode and add the BGP address-family configuration for the relevant VRF.

Example:

```

config-transaction
router bgp 200
address-family ipv4 vrf 11
  neighbor 10.20.1.0 remote-as 200

```

```
propagate-aspath
exit
```

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

Step 2

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

Example:

```
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, we recommend you to configure the overlay AS number as a unique AS number within both the overlay and the underlay networks.

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 using CLI commands

You can control and configure the number of route–TLOC tuples that Cisco IOS XE Catalyst SD-WAN devices and Cisco Catalyst SD-WAN Controllers advertise, enabling you to optimize route advertisement and path selection based on your network requirements. You can execute the commands using CLI Add-on template.

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 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, an SD-WAN Controller could potentially learn eight route–TLOC tuples for the same route. By default, SD-WAN devices and SD-WAN Controllers advertise 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.

Procedure

Execute this command:

Example:

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

From Cisco Catalyst SD-WAN Control Components Release 20.8.x, you can configure an 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.

From Cisco Catalyst SD-WAN Control Components Release 20.9.x, you can configure an 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 configured limit is lower than the number of route-TLOC tuples, the SD-WAN device or SD-WAN Controller advertises only the best routes.

Configure the number of installed OMP paths using CLI commands

Cisco IOS XE Catalyst SD-WAN devices install OMP paths received from the SD-WAN Controller into their local route table. By default, Cisco IOS XE Catalyst SD-WAN devices installs a maximum of four unique OMP paths into its route table. You can modify this number using the CLI add-on template.

Procedure

Execute the **ecmp-limit** command:

Example:

```
Device(config-omp) # ecmp-limit <number-of-paths>
```

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

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

Configure the OMP hold time using CLI commands

You can modify the OMP hold time interval using CLI commands.

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 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.

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

Procedure

To modify the OMP hold time interval, use the **timers holdtime** command:

Example:

```
Device(config-omp) # timers holdtime <seconds>
```

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

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

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

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

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

Note

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

Note

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.

Configure the OMP advertisement interval and end-of-RIB timer using CLI commands

By default, OMP sends Update packets once per second. You can modify the interval using a CLI command.

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

Procedure

Step 1 To modify the interval, use the **timers advertisement-interval** command:

Example:

```
Device(config-omp)# timers advertisement-interval <interval>
```

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

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

Step 2 To modify the EOR timer, use the **timers eor-timer** command.

Example:

```
Device(config-omp)# timers eor-timer<eor-timer>
```

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

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

■ Configure the OMP advertisement interval and end-of-RIB timer using CLI commands