Routing Configuration Guide, Cisco IOS XE Catalyst SD-WAN Release 17.x

First Published: 2021-01-01

Last Modified: 2026-03-09

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

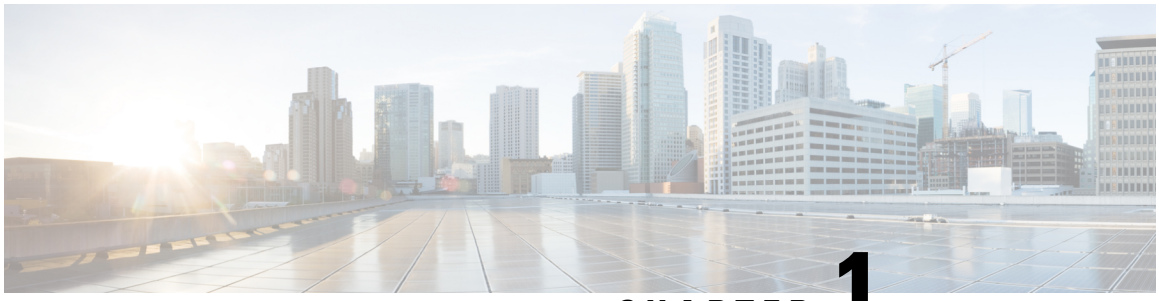

CHAPTER 1

Read Me First

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.

Related References

- Cisco Catalyst SD-WAN Control Components Compatibility Matrix and Server Recommendations
- Cisco Catalyst SD-WAN Device Compatibility

User Documentation

- User Documentation for Cisco IOS XE Catalyst SD-WAN Release 17

Communications, Services, and Additional Information

- Sign up for Cisco email newsletters and other communications at: Cisco Profile Manager.
- For information on the latest technical, advanced, and remote services to increase the operational reliability of your network visit Cisco Services.
- To browse and discover secure, validated enterprise-class apps, products, solutions, and services, visit Cisco Devnet.
- To obtain general networking, training, and certification titles from Cisco Press Publishers, visit Cisco Press.
- To find warranty information for a specific product or product family, visit Cisco Warranty Finder.
- To view open and resolved bugs for a release, access the Cisco Bug Search Tool.
- To submit a service request, visit Cisco Support.

Documentation Feedback

To provide feedback about Cisco technical documentation use the feedback form available in the right pane of every online document.

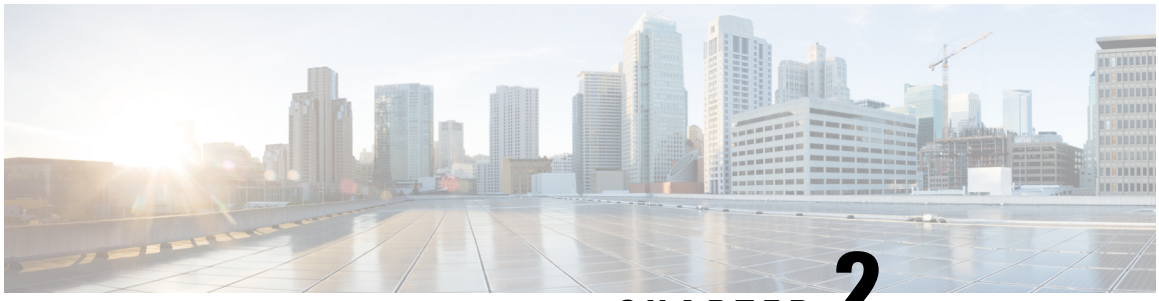

CHAPTER 2

What's New in Cisco IOS XE (Catalyst SD-WAN)

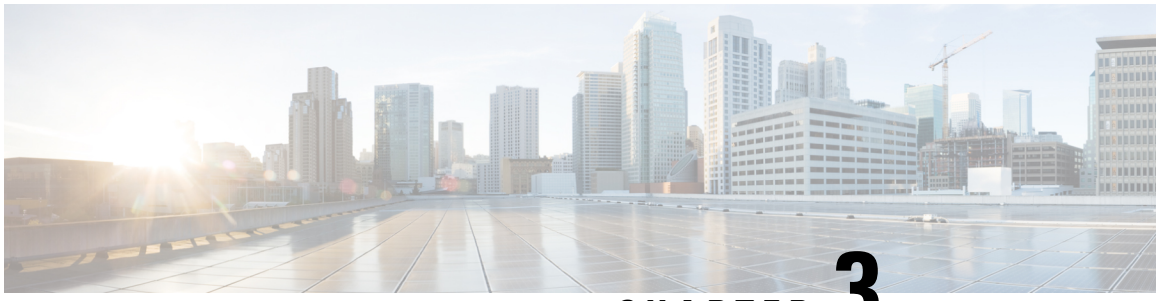

CHAPTER 3

Overlay Management Protocol

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 redestribution 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 learns 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

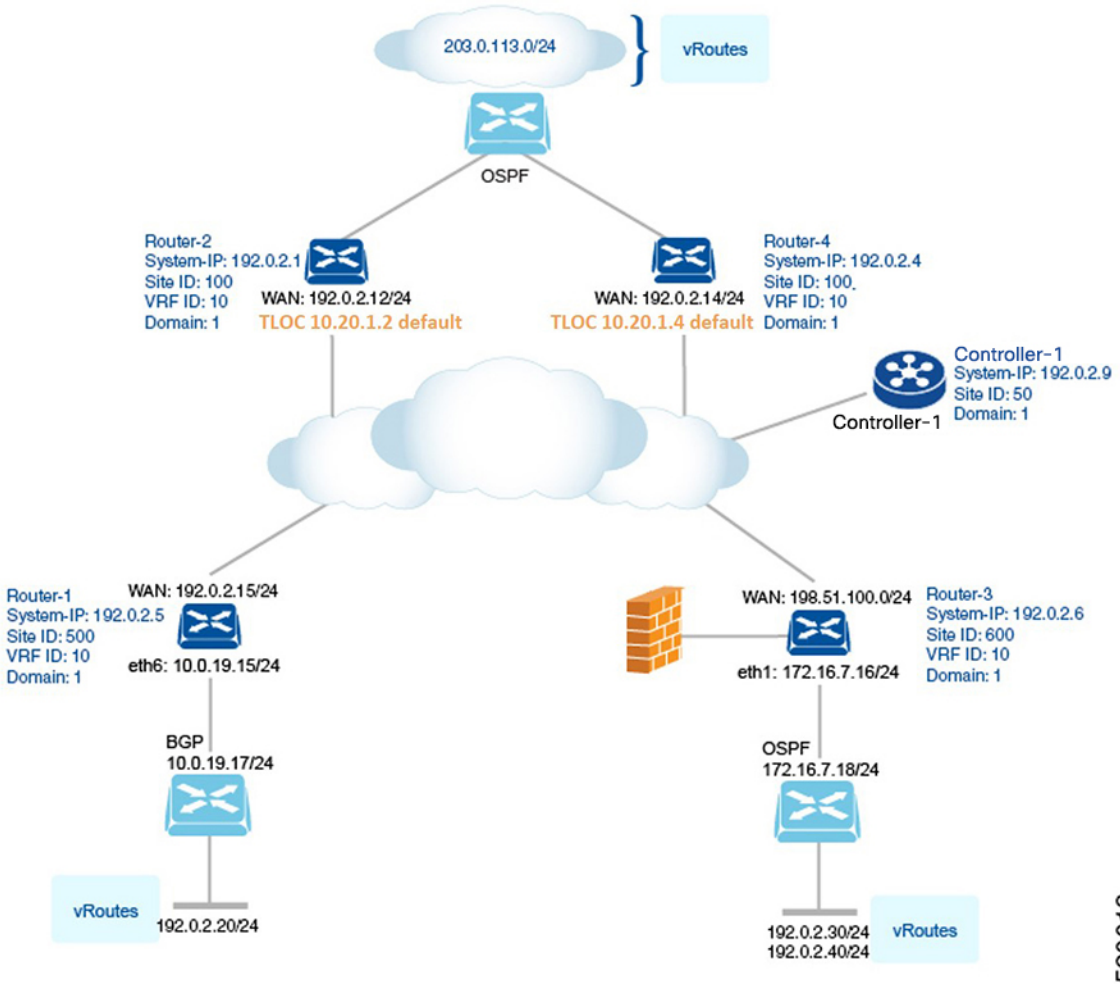

OMP routes

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. Fo rmore information, see OMP vRoute advertisement optimization using system path, on page 22.
- 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

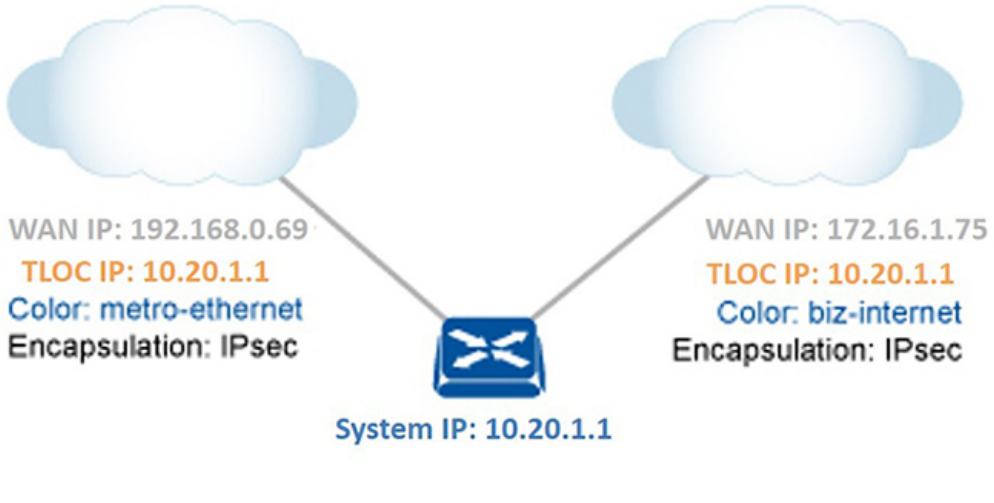

The diagram to the right shows a device that has two WAN connections and hence two TLOCs. The system IP address of the router is 10.20.1.1. The TLOC on the left is uniquely identified by the system IP address 10.20.1.1, the color metro-ethernet, and the encapsulation IPsec. It maps to the physical WAN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

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.

```
# omp
```

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	• 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 15.
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	Direct vs. transport gateway path Prefers 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 • 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	Border router preference Prefers a path with a higher border router preference value.
12	Edge devices	Derived affinity Prefers 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, 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. **Note** When comparing two paths with similar reorigination types, one with an affinity value and one without, prefers the path with an affinity value.
14	Edge devices	TLOC preference Select an OMP path with a higher TLOC preference value. **Note** With TLOC preference and AAR policy configured, outbound and inbound traffic may follow different paths when the preferred TLOC goes out of SLA. For outbound traffic, tunnels in SLA will be preferred regardless of TLOC preference; however, TLOC preference still dictates inbound path selection.

Step	Applies to	Description
15	Edge devices Cisco Catalyst SD-WAN Controller	Origin type and subtype Compares the origin type and subtype, and selects the first match from this 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 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 Controllerto 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: • 20.18.x and later: 300 seconds • 20.16.x: 5400 seconds • 20.12.1 to 20.15.x: 300 seconds • Before 20.12.1: 60 seconds Defaults, by Cisco IOS XE Catalyst SD-WAN release: • 17.18.1 and later: 300 seconds • 17.16.x: 5400 seconds
EOR Timer(In Second)	Specify how long to wait after an OMP session has gone down and then come back up to send an end-of-RIB (EOR) marker. After this marker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. Range: 1 through 3600 seconds (1 hour) Default: 300 seconds (5 minutes)
Overlay AS	Specify a BGP AS number that OMP advertises to the BGP neighbors of the router.
Shutdown	Enable this option to disable OMP and disable the Cisco Catalyst SD-WAN overlay network. OMP is enabled by default.
OMP Admin Distance Ipv4	To advertise a route over OMP, configure the OMP administrative distance for the IPv4 address lower than the leaked route administrative distance. 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	(Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.13.1a) In a Multi-Region Fabric scenario, if using secondary regions, enable this option to enable packets to use all available paths rather than only direct paths. By default, when a direct path is available to reach a destination, the overlay management protocol (OMP) enables only the direct path to the routing forwarding layer because the direct path uses fewer hops. This logic is part of route optimization. The result is that the forwarding layer, which includes application-aware routing policy, can only use the direct path. **Treat Hierarchical and Direct Paths Equally** disables this comparison of the number of hops so that traffic can use either the direct secondary-region path (fewer hops) or the primary-region path (more hops). When you disable the comparison of the number of hops, OMP applies equal-cost multi-path routing (ECMP) to all routes, and packets can use all available paths.
Transport Gateway Path Behavior	(Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.13.1a) Choose one of the following: • **Prefer Transport Gateway Path**: For devices that can connect through a transport gateway, use only the transport gateway paths, even if other paths are available. • **Do ECMP Between Direct and Transport Gateway Paths**: For devices that can connect through a transport gateway and through direct paths, apply ECMP to all available paths.

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)	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.

What to do next

See Configure OMP options, on page 28.

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

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	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.

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 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:

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 recommended 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 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 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.

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 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. 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).

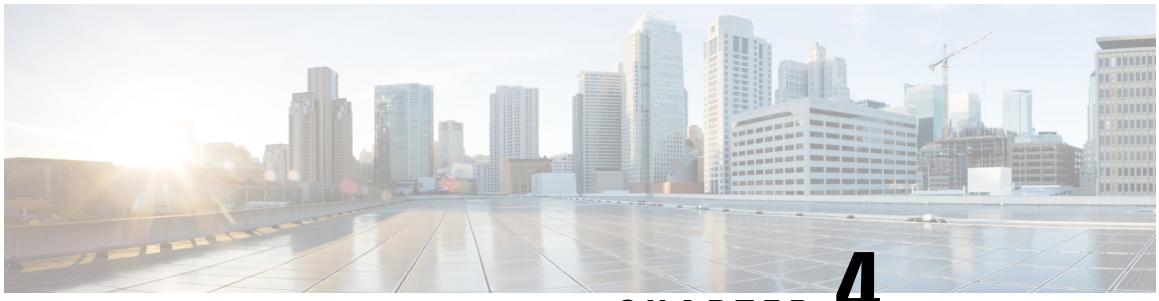

CHAPTER 4

Border Gateway Protocol

Feature history for BGP

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

Table 10: 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).
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.

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`

BGP for Cisco SD-WAN overlay networks

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

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

BGP topologies

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

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

BGP Community Propagation

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

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

Configure BGP

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

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

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

Procedure

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

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

Create a BGP template

Procedure

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

Step 2 Click **Device Templates**.

Step 3 Click **Create Template**.

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

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

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

a. Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.

b. Under **Additional VPN 0 Templates**, click **BGP**.

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

Step 7 To 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.

Step 8 In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.

Step 9 In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Configure basic BGP parameters

Procedure

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

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

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

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

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

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

Configure BGP address family

Procedure

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

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

Parameter	Option	Sub-Option	Description
Redistribute	Click **Redistribute** > **New Redistribute**.		
	Mark as Optional Row	Check **Mark as Optional Row** to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.	
	Protocol	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are:	
		static	Redistribute static routes into BGP.
		connected	Redistribute connected routes into BGP.
		ospf	Redistribute Open Shortest Path First routes into BGP.
		omp	Redistribute Overlay Management Protocol routes into BGP.
		nat	Redistribute Network Address Translation routes into BGP.
		natpool-outside	Redistribute outside NAT routes into BGP.
		At a minimum, choose the following: • 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.		

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

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

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

Configure BGP neighbors

Before you begin

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

Procedure

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

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

Parameter Name	Options	Sub-Options	Description
Address Family	Click **On** and select the address family. Enter the address family information. The software supports only the BGP IPv4 unicast address family.		
	Address Family	Select the address family. The software supports only the BGP IPv4 unicast address family.	
	Maximum Number of Prefixes	Specify the maximum number of prefixes that can be received from the neighbor. Range: 1 through 4294967295 Default: 0	
		Threshold	Specify the threshold at which to generate a warning message or restart the BGP connection. The threshold is a percentage of the maximum number of prefixes. You can specify either a restart interval or a warning only.
		Restart Interval	Specify the duration to wait for restarting the BGP connection.*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.		

Step 2 Click **Save**.

Configure MPLS interface

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

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

To configure an MPLS interface, do the following:

Procedure

Step 1 Click **MPLS Interface**.

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

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

Configure label range

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

Sample configuration:

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

Configure route targets

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

Procedure

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

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

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

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

Configure advanced neighbor parameter

Procedure

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

Parameter Name	Description
Next-Hop Self	Click **On** to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Click **On** to send the local router's BGP community attribute to the BGP neighbor.
Send Extended Community	Click **On** to send the local router's BGP extended community attribute to the BGP neighbor.
Negotiate Capability	Click **On** to allow the BGP session to learn about the BGP extensions that are supported by the neighbor.
Source Interface Address	Enter the IP address of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor.
Source Interface Name	Enter the name of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor, in the format **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
Advertisement Interval	For the BGP neighbor, set the minimum route advertisement interval (MRAI) between when BGP routing update packets are sent to that neighbor. Range: 0 through 600 seconds Default: 5 seconds for IBGP route advertisements; 30 seconds for EBGP route advertisements

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

Change the scope of a parameter value

Before you begin

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

Procedure

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

Parameter Name	Description
Device Specific	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

Procedure

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

Step 2 Click **Device Templates**.

Step 3 Click **Create Template**.

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

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

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

a. Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.

b. Under **Additional VPN 0 Templates**, click **BGP**.

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

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

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

Step 8 Click **Save**.

Configure BGP routing in a service profile using configuration group

Before you begin

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

Procedure

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

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

a. Configure Basic Configuration fields.

Table 11: Basic Configuration

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

b. Configure Unicast Address Family fields.

Table 12: Unicast Address Family

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

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

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

c. Configure Neighor fields.

Table 13: Neighbor

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

Field	Description
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Send Label	Enable this option to allow the routers advertise to each other so that they can send MPLS labels with the routes. If the routers successfully negotiate their ability to send MPLS labels, the routers add MPLS labels to all the outgoing BGP updates.
Add Neighbor Address Family	
Family Type*	Choose the BGP IPv4 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Maximum Prefix Reach Policy*	Choose one of the following options: • **Policy Off**: Policy is off. • **Policy On - Restart**: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: • **Maximum Number of Prefixes***: Enter the maximum prefix limit. Range: 1 to 4294967295 • **Threshold (percentage)**: Enter the threshold value: Range: 1 to 100 Default: 75 • **Restart Interval (minutes)***: Enter the time interval. Range: 1 to 65535 minutes • **Policy On - Warning message**: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • **Policy On - Disable Peer Neighbor**: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

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

Field	Description
Hold Time (seconds)	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor, to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive time)
Add IPv6 Neighbor Address Family	
Family Type*	Choose the BGP IPv6 unicast address family.
In Route Policy	Specify the name of a route policy to apply to prefixes received from the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Out Route Policy	Specify the name of a route policy to apply to prefixes sent to the neighbor. Route policy is not supported in Cisco vManage Release 20.9.1.
Maximum Prefix Reach Policy*	Choose one of the following options: • **Policy Off**: Policy is off. • **Policy On - Restart**: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: • **Maximum Number of Prefixes***: Enter the maximum prefix limit. Range: 1 to 4294967295 • **Threshold (percentage)**: Enter the threshold value: Range: 1 to 100 Default: 75 • **Restart Interval (minutes)***: Enter the time interval. Range: 1 to 65535 minutes • **Policy On - Warning message**: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • **Policy On - Disable Peer Neighbor**: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

Configure BGP routing in a transport profile using a configuration group

Before you begin

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

Procedure

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

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

a. Configure Basic Configuration fields.

Table 14: Basic Configuration

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

b. Configure Unicast Address Family.

Table 15: Unicast Address Family

Field	Description
IPv4 Settings	

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

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

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

c. Configure MPLS Interface.

Table 16: MPLS Interface

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

d. Configure Neighbor.

Table 17: Neighbor

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

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

Field	Description
Maximum Prefix Reach Policy*	Choose one of the following options: • **Policy Off**: Policy is off. • **Policy On - Restart**: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: • **Maximum Number of Prefixes***: Enter the maximum prefix limit. Range: 1 to 4294967295 • **Threshold (percentage)**: Enter the threshold value: Range: 1 to 100 Default: 75 • **Restart Interval (minutes)***: Enter the time interval. Range: 1 to 65535 minutes • **Policy On - Warning message**: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • **Policy On - Disable Peer Neighbor**: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.
IPv6 Settings	
Address*	Specify the IP address of the BGP neighbor.
Description	Enter a description of the BGP neighbor.
Remote AS*	Enter the AS number of the remote BGP peer.
Interface Name	Enter the interface name. This interface is used as the source of the TCP session when establishing neighborship. We recommend that you use a loopback interface.
Allowas in Number	Enter the number of times to allow the advertisement of the autonomous system number (ASN) of a provider edge (PE) device. The range is 1 to 10. If no number is specified, the default value of three times is used.
AS Override	Enable this option to replace the AS number of the originating router with the AS number of the sending BGP router.
Shutdown	Disable this option to enable BGP for the VPN.
Advanced Options	
Next-Hop Self	Enable this option to configure the router to be the next hop for routes advertised to the BGP neighbor.

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

Field	Description
Maximum Prefix Reach Policy*	Choose one of the following options: • **Policy Off**: Policy is off. • **Policy On - Restart**: Configure the time interval at which a peering session is re-established by a device when the number of prefixes that have been received from a peer has exceeded the maximum prefix limit. When you choose this option, the following fields appear: • **Maximum Number of Prefixes***: Enter the maximum prefix limit. Range: 1 to 4294967295 • **Threshold (percentage)**: Enter the threshold value: Range: 1 to 100 Default: 75 • **Restart Interval (minutes)***: Enter the time interval. Range: 1 to 65535 minutes • **Policy On - Warning message**: Configure the device to disable the restart capability to allow you to adjust a peer that is sending too many prefixes. • **Policy On - Disable Peer Neighbor**: When the device receives too many prefixes from a peer, and the maximum prefix limit is exceeded, the peering session is disabled or brought down.

e. Configure Advanced fields.

Table 18: Advanced

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

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

Configure BGP using CLI

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

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


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


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

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

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

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

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

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

Example of configuring Service-Side routing using CLI

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

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

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

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

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

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

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

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

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

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

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

!
address-family ipv4 unicast
neighbor 10.20.25.16 activate
exit-address-family
!
address-family 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
```

Verify BGP redistribute route in OMP

To verify BGP redtribute route in OMP:

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

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

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

```
peer    192.168.0.1
```

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

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

Generating output, this might take time, please wait ...
Code:
C -> chosen
I -> installed
Red -> redistributed
Rej -> rejected
L -> looped
R -> resolved
S -> stale
Ext -> extranet
Inv -> invalid
Stg -> staged
IA -> On-demand inactive
U -> TLOC unresolved
BR-R -> Border-Router reoriginated
TGW-R -> Transport-Gateway reoriginated
R-TGW-R -> Reoriginated Transport-Gateway reoriginated
DERIVED
AFFINITY AFFINITY
PATH PSEUDO
GROUP GROUP
FROM PEER ID LABEL STATUS KEY TLOC IP COLOR
ENCAP PREFERENCE NUMBER NUMBER
------------------------------------------------------------------------------
-------------------------------------------------
12.16.255.20 1 1003 C,I,R 1 10.2.1.25 publicinternet
ipsec - None None
12.16.255.20 2 1003 C,I,R 1 10.2.1.25 lte
ipsec - None None
12.16.255.20 3 1003 C,I,R 1 10.2.1.25 gold
ipsec - None None
12.16.255.20 4 1003 C,I,R 1 10.2.1.25 silver
ipsec - None None
12.16.255.20 5 1003 C,I,R 1 10.2.1.25
```

Redistribute BGP routes and AS path information

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

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

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

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

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

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

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

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

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

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

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

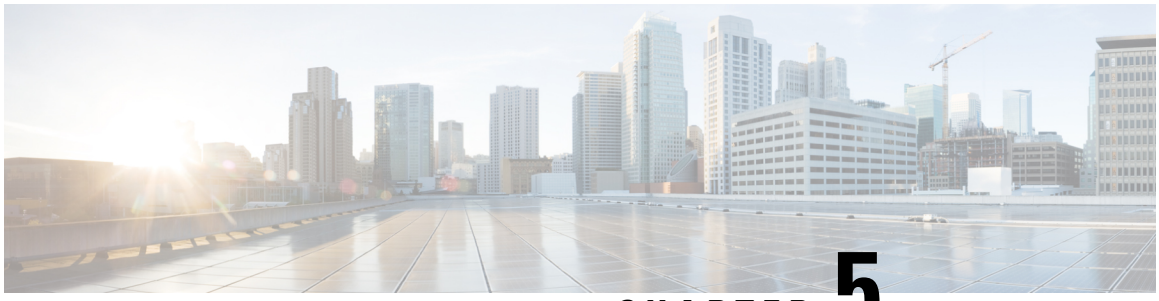

CHAPTER 5

Open Shortest Path First Routing Protocol

Feature history for OSPF routing protocol

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

Table 19: Feature History

Feature	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 IPSec Authentication	Cisco IOS XE Catalyst SD-WAN Release 17.18.1a Cisco Catalyst SD-WAN Manager Release 20.18.1	OSPFv3 authentication protects OSPFv3 routing payload with IPSec encryption and hashing. The solution is based on statically configured cryptographic keys that build on top of crypto map solutions. OSPFv3 is supported on service side and transport side.

OSPF and OSPFv3 protocol

OSPF

OSPF is a widely used link-state routing protocol designed for Internet Protocol (IP) networks. As an Interior Gateway Protocol (IGP), it operates within a single autonomous system (AS). OSPF gathers link-state information from routers to construct a topology map of the network.

The Cisco Catalyst SD-WAN overlay network supports OSPF unicast routing protocols. You can configure these protocols on Cisco IOS XE Catalyst SD-WAN devices in any Virtual Routing and Forwarding (VRF) except for transport and management VRFs to provide reachability to networks at their local site Cisco IOS XE Catalyst SD-WAN devices can redistribute route information learned from OSPF into Overlay Management Protocol (OMP) so that OMP can better choose paths within the overlay network. 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 DTLS connections of the devices so that they can reach the WAN cloud. OSPF can be the routing protocol.

The OSPF sessions run over a DTLS connection created on the loopback interface in VRF 0, the transport VRF responsible for carrying control traffic in the overlay network. The Cisco SD-WAN Validator learns about this DTLS connection via the loopback interface and conveys this information to theCisco 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.

OSPFv3

OSPFv3 is an enhanced version of the OSPF routing protocol specifically designed for IPv6 networks. A significant change in OSPFv3 is the decoupling of IP addressing from the routing topology information. 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 more. 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.

OSPFv3 IPSec authentication

OSPFv3 IPSec authentication refers to a security mechanism used in IPv6 networks where Open Shortest Path First version 3 (OSPFv3) routing protocol packets are authenticated and optionally encrypted through the IP Security (IPSec) protocol. In order to ensure that OSPFv3 packets are not altered and re-sent to the device, OSPFv3 packets must be authenticated. OSPFv3 uses the IPSec secure socket API to add authentication to OSPFv3 packets. This API supports IPv6.

Since OSPFv3 does not provide built-in authentication or encryption, IPSec is employed to ensure the integrity and confidentiality of OSPFv3 routing messages exchanged between routers. This approach helps prevent unauthorized devices from modifying or injecting routing information and protects against unauthorized packet capture and replay attacks, thereby enhancing the security of routing updates in enterprise or service provider IPv6 networks. To implement OSPFv3 IPSec authentication, you must configure IPSec security associations between participating routers, specifying appropriate authentication (such as HMAC with SHA or MD5).

To configure IPsec, you configure a security policy, which is a combination of the security policy index (SPI) and the key (the key is used to create and validate the hash value). IPsec for OSPFv3 can be configured on an interface or on an OSPFv3 area. For higher security, you should configure a different policy on each interface configured with IPsec. If you configure IPsec for an OSPFv3 area, the policy is applied to all of the interfaces in that area, except for the interfaces that have IPsec configured directly. Once IPsec is configured for OSPFv3, IPsec is invisible to you.

The secure socket API is used by applications to secure traffic. The API needs to allow the application to open, listen, and close secure sockets. The binding between the application and the secure socket layer also allows the secure socket layer to inform the application of changes to the socket, such as connection open and close events. The secure socket API is able to identify the socket; that is, it can identify the local and remote addresses, masks, ports, and protocol that carry the traffic requiring security.

Each interface has a secure socket state:

- NULL: Do not create a secure socket for the interface if authentication is configured for the area.
- DOWN: IPsec has been configured for the interface (or the area that contains the interface), but OSPFv3 either has not requested IPsec to create a secure socket for this interface, or there is an error condition.
- GOING UP: OSPFv3 has requested a secure socket from IPsec and is waiting for a CRYPTO_SS_SOCKET_UP message from IPsec.
- UP: OSPFv3 has received a CRYPTO_SS_SOCKET_UP message from IPsec.
- CLOSING: The secure socket for the interface has been closed. A new socket may be opened for the interface, in which case the current secure socket makes the transition to the DOWN state. Otherwise, the interface will become UNCONFIGURED.
- UNCONFIGURED: Authentication is not configured on the interface.

OSPFv3 will not send or accept packets while in the DOWN state.

Restrictions for OSPF and OSPFv3 protocols

- The OSPF for IPv6 (OSPFv3) authentication with IPSec feature is not supported on the IP BASE license package. The Advanced Enterprise Services package license must be used.
- OSPFv3 encryption is not supported.

- The OSPFv3 configuration is supported only on interface-level using configuration groups. However, both interface-level and area-level configuration is supported using CLI add-on templates.

Prerequisites for OSPFv3 IPSec authentication

Configure the IP Security (IPsec) secure socket application program interface (API) on OSPFv3 to enable authentication.

Configure OSPF

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

Procedure

Step 1 Create an OSPF feature template to configure OSPF parameters. OSPF can be used for transport-side routing to enable communication between the Cisco IOS XE Catalyst SD-WAN devices when the router is not directly connected to the WAN cloud.

Step 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 OSPF template

Procedure

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

Step 2 Click **Device Templates**.

In Cisco vManage Release 20.7.x and earlier releases, **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, 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.

Step 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.

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 choose one of the following:

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.

Create basic OSPF

Procedure

To configure basic OSPF, select **Basic Configuration** and then configure the following parameters. All these parameters are optional.

Table 20:

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

Redistribute routes into OSPF

To redistribute routes learned from other protocols into OSPF:

Procedure

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

Step 2 Click **Device Templates**.

Step 3 Click **Create Template**.

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

Step 5 From the **Device Model** drop-down list, 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.

Step 6 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 21:

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.

What to do next

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 Interfaces in an OSPF area

Procedure

Step 1 To configure an OSPF area within a VPN on a Cisco IOS XE Catalyst SD-WAN device, choose **Area** > **Add New Area**. For OSPF to function, you must configure area 0.

Table 22:

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 IOS XE 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.

Step 2 Click **Add** to save the new area.

Step 3 To configure the properties of an interface in an OSPF area, choose **Add Interface**. In the **Add Interface** popup, configure the following parameters:

Table 23:

Parameter Name	Description
Interface Name	Enter the name of the interface, in the format **ge** *slot*/*port* or **loopback** *number*.
Hello Interval	Specify how often the router sends OSPF hello packets. *Range:* 1 through 65535 seconds *Default:* 10 seconds
Dead Interval	Specify how often the Cisco 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

Step 4 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 24:

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.

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

Step 5 Click **Save** to save the interface configuration.

Configure an interface range for summary LSAs

Procedure

Step 1 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 25:

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.

Step 2 Click **Save** to save the area range.

Configure other OSPF properties

Procedure

Step 1 To configure other OSPF properties, click **Advanced** and configure the following properties:

Table 26:

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.

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

Configure OSPFv3 IPSec authentication

Once you have configured OSPFv3 and decided on your authentication, you must define the security policy on each of the devices within the group. The security policy consists of the combination of the key and the SPI. To define a security policy, you must define an SPI and a key.

You can configure an authentication or encryption policy either on an interface or for an OSPFv3 area. When you configure for an area, the security policy is applied to all of the interfaces in the area. For higher security, use a different policy on each interface.

You can configure authentication on virtual links.

- Defining Authentication on an Interface
- Defining Authentication in an OSPFv3 Area

Configure OSPFv3 IPSec authentication at an interface-level using CLI

Before you begin

Follow these steps to configure OSPFv3 IPSec authentication at an interface-level using CLI.

Procedure

Step 1 Enter global configuration mode.

Example:

```
Device# config-transaction
```

Step 2 In the configuration mode, configure an interface type such as, Gigabit Ethernet.

Specifies an interface type and number, and places the device in interface configuration mode.

Example:

```
Device(config)# interface GigabitEthernet3
```

Step 3 Configure OSPFv3 authentication on the interface.

Specifies the authentication type for an interface.

Example:

```
Device(config-if)# ospfv3 authentication ipsec spi 256 sha1 0 09876543210987654321098765432109876543210
```

Configure OSPFv3 IPSec authentication at an area-level using CLI

Before you begin

Follow these steps to configure OSPFv3 IPSec authentication at an area-level.

Procedure

Step 1 Enter global configuration mode.

Example:

```
Device# config-transaction
```

Step 2 Enable OSPFv3 router configuration mode.

Example:

```
Device(config)# router ospfv3 <process-id>
```

Step 3 Configure OSPFv3 authentication on the interface.

Enables authentication in an OSPFv3 area.

Example:

```
Device(config-rtr)# area <area-id> authentication ipsec spi <spi> authentication-algorithm
```

```
Device(config)# router ospfv3 <process-id>
Device(config-rtr)# area <> authentication ipsec spi <> [md5/sha1] <>
interface GigabitEthernetY/Y
ospfv3 <> [ipv4/ipv6] area <>
```

Configure OSPF using CLI

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.

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
```

Configuration examples for IPv6 OSPFv3 IPSec authentication

You can configure OSPFv3 IPSec authentication on an interface or in an area.

Note Interface-level authentication takes priority over area-level authentication

Configuring IPv6 OSPFv3 authentication on an interface

This example shows how to define authentication on Ethernet interface 0/0:

```
interface Ethernet0/0
 ospfv3 1 ipv6 area 0
 ospfv3 1 ipv6 dead-interval 40
 ospfv3 1 ipv6 hello-interval 10
 ospfv3 1 ipv6 retransmit-interval 5
 ospfv3 authentication ipsec spi 256 sha1 0 09876543210987654321098765432109876543210
```

Configuring IPv6 OSPFv3 authentication in an area

This example shows how to define authentication on OSPFv3 area 0:

```
router ospfv3 1
 router-id 10.11.11.1
 area 0 authentication ipsec spi 256 sha1 0 09876543210987654321098765432109876543210
```

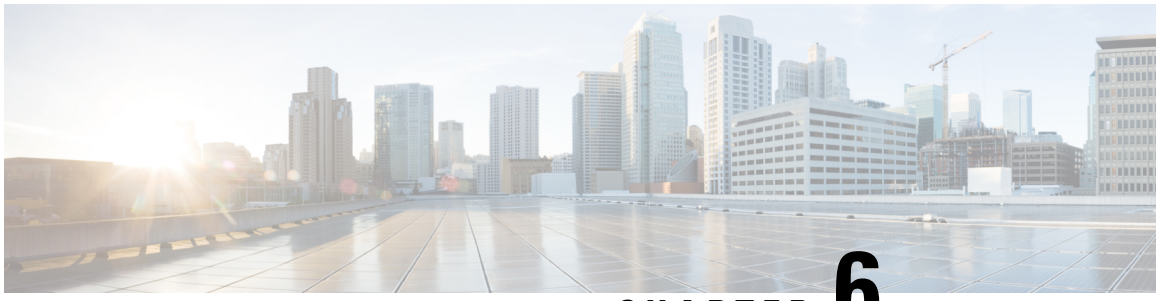

CHAPTER 6

Enhanced Interior Gateway Routing Protocol

EIGRP in Cisco IOS XE Catalyst SD-WAN devices

Enhanced Interior Gateway Routing Protocol (EIGRP) is a Cisco proprietary routing protocol that:

- is an open-standard Interior Gateway Protocol (IGP),
- is an enhancement to the original Interior Gateway Routing Protocol (IGRP) developed by Cisco,
- does not fully update if there are no changes in the network, which reduces flooding activities in other IGPs.

EIGRP is supported only on Cisco IOS XE Catalyst SD-WAN devices. For more information, see Introduction to EIGRP.

Benefits of EIGRP

EIGRP provides several advantages that enhance network performance and simplify management for network administrators.

These are the key benefits of using EIGRP:

- Increased network width: EIGRP supports an increased network width from 15 to 100 hops.
- Fast convergence: EIGRP provides fast convergence.
- Incremental updates: EIGRP uses incremental updates, which minimizes bandwidth consumption.
- Protocol-independent neighbor discovery: EIGRP supports protocol-independent neighbor discovery.
- Easy scaling: EIGRP is designed for easy scaling.

EIGRP restrictions

When implementing EIGRP, note these 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 Catalyst SD-WAN Controller centralized control policy.

Configure EIGRP

Use this task to set up EIGRP for routing in Cisco Catalyst SD-WAN, enabling efficient and scalable routing across your overlay.

You can EIGRP using Cisco SD-WAN Manager templates, configuration groups, or directly via the CLI on Cisco IOS XE Catalyst SD-WAN devices.

Before you begin

Follow one of these steps to configure EIGRP:

Procedure

Step 1 Configure EIGRP using Cisco SD-WAN Manager templates.

This method involves creating a reusable EIGRP feature template in Cisco SD-WAN Manager, which defines the basic EIGRP parameters such as the autonomous system ID, and specifies how routes are redistributed. This template can then be applied to multiple devices. For more information, see Configure EIGRP using Cisco SD-WAN Manager, on page 88.

Step 2 Configure EIGRP using a configuration group.

This method applies EIGRP configurations within a service profile using configuration groups in Cisco SD-WAN Manager, suitable for service-side routing within specific VPNs. For more information, see Configure EIGRP using a configuration group, on page 90.

Step 3 Configure EIGRP using CLI.

This method involves directly entering EIGRP configuration commands on the Cisco IOS XE Catalyst SD-WAN device's command-line interface, providing granular control over the configuration.

What to do next

After configuring EIGRP, verify the configuration using the appropriate `show` commands on the device. For example, `show ip route vrf <vpn-id>` for IPv4 EIGRP routes. For more information, see Verification commands for EIGRP configuration, on page 93.

Configure EIGRP using Cisco SD-WAN Manager

Use this task to configure the EIGRP routing protocol using the EIGRP feature template in Cisco SD-WAN Manager.

This template defines the basic EIGRP parameters, such as the autonomous system ID. It also specifies how routes are redistributed. You can apply this template to multiple devices.

Before you begin

Follow these steps to create an EIGRP template:

Procedure

Step 1 From the Cisco SD-WAN Managermenu, choose **Configuration** > **Templates** > **Feature Templates**.

Step 2 Click **Add Template** and select a device from the list.

Step 3 From the **Other Templates** section, choose **EIGRP** and enter a name and a description for the template.

Step 4 Click **Basic Configuration** to configure the local autonomous system number for the template.

a) In the **Autonomous System ID** field, enter the local AS number (Range: 1-65,535).

Step 5 To redistribute routes from one protocol into an EIGRP routing domain, click **New Redistribute** under **IP4 Unicast Address Family** and enter the following parameter values:

a) Select **Mark as Optional Row**.

This action marks the 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.

b) Choose the **Protocol** from which to redistribute routes into EIGRP (e.g., **omp**, **connected**, **static**, **ospf**, **bgp**, **nat-route**).

c) Enter the name of the **Route Policy** to apply to redistributed routes.

d) Click **Add** to save the redistribution information.

Step 6 To advertise a prefix into the EIGRP routing domain, click **Network**, then click **New Network** and enter the following parameter values:

a) Select **Mark as Optional Row**.

This action marks the 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.

b) In the **Network Prefix** field, enter the network prefix you want EIGRP to advertise in the format of prefix/mask.

c) Click **Add** to save the network prefix.

Step 7 To configure advanced parameters for EIGRP, click **Advanced** and configure the following parameter values:

a) In the **Hold Time (seconds)** field, set the interval after which EIGRP considers a neighbor to be down (Range: 0-65,535, Default: 15 seconds).

b) In the **Hello Interval (seconds)** field, set the interval at which the router sends EIGRP hello packets (Range: 0-65,535, Default: 5 seconds).

c) In the **Route Policy Name** field, enter the name of an EIGRP route policy.

Step 8 To configure authentication for EIGRP routes click **Authentication**:

a) Click **Authentication**

b) Choose **global** in the **Authentication Key** drop-down list, and then choose **md5** or **hmac-sha-256**.

- For **MD5**: Enter an **MD5 Key ID** and **MD5 Authentication Key**.
- For **HMAC-SHA-256:** Enter an Authentication Key.

c) 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.

Step 9 To configure interface parameters for EIGRP routes, click **Interface**, and enter the following parameter values:

a) Select **Mark as Optional Row**.

This action marks the 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.

b) In the **Interface name** field, enter the interface name(s) on which EIGRP should run.
c) Set **Shutdown** to **No** (default) to enable the interface to run EIGRP, or **Yes** to disable it.
d) Click **Add** to save the interfaces.

An EIGRP feature template is created, ready to be applied to devices.

What to do next

When the configuration has been applied to the devices, verify the EIGRP configuration using the appropriate `show` commands on the device. For example, `show ip route vrf <vpn-id>` for IPv4 EIGRP routes. For more information, see Verification commands for EIGRP configuration, on page 93.

Configure EIGRP using a configuration group

Use this task to apply EIGRP configurations within a service profile using configuration groups in Cisco SD-WAN Manager.

This method is suitable for service-side routing within specific VPNs and leverages configuration groups for simplified management.

Before you begin

Follow these steps to configure EIGRP routing in a service profile:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Configuration Groups**.
Step 2 Configure an EIGRP Routing feature in a Service Profile.
Step 3 Configure basic settings.

a) In the **Autonomous System ID** field, enter the local autonomous system (AS) number (Range: 1-65535).
b) For Network: Enter the IP address and subnet mask.
c) For Interface:

1. Select one of the **Protocol** options from which to redistribute routes into EIGRP (e.g., bgp, connected, nat-route, omp, ospf, ospfv3, static).

Note
From Cisco IOS XE Catalyst SD-WAN Release 16.12.1b and later, you can set metric values for redistribution by using the CLI add-on feature template. Use the following command:

```
redistribute ospf 1 metric 1000000 1 1 1 1500
```

For more information, see *CLI Add-on Feature Templates*. .

2. Enter the name of the Route Policy to apply to redistributed routes.

d) Configure IPv4 unicast address family:

1. Click **Add Interface**.
2. Provide a value for the **AF Interface**.
3. Set **Shutdown** to OFF (default) to enable the interface, or ON to disable it.
4. (Optional) Click **Add Summary Address** and enter an IPv4 address and choose a subnet mask.

e) Configure authentication:

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

f) Configure advanced settings:

1. In the **Hold Time (seconds)** field, set the interval after which EIGRP considers a neighbor to be down. Specify a value in the range of 0-65535. The default is 15 seconds.
2. In the **Hello Interval (seconds)** field, set the interval at which the router sends EIGRP hello packets. Specify a value between 0 and 65535. The default is 5 seconds.
3. In the **Route Policy** field, enter the name of an EIGRP route policy.
4. Toggle **Filter** to **ON**. This filters routes that do not match the policy.

EIGRP routing is configured within the specified service profile, ready for deployment.

What to do next

Verify the EIGRP configuration using the appropriate `show` commands on the device. For example, `show ip route vrf <vpn-id>` for IPv4 EIGRP routes. For more information, see Verification commands for EIGRP configuration, on page 93.

Configure EIGRP using CLI

You can use this task to configure EIGRP parameters directly on Cisco IOS XE Catalyst SD-WAN devices via the command-line interface.

Before you begin

To configure EIGRP using the CLI, complete these steps:

Procedure

Step 1 Configure EIGRP on Cisco IOS XE Catalyst SD-WAN devices using these commands:

```
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
!
```

Step 2 (Optional) To advertise EIGRP routes to OMP, use these commands:

```
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
  !
 !
```

EIGRP is configured on the Cisco IOS XE Catalyst SD-WAN device via CLI.

Verification commands for EIGRP configuration

This section details the essential verification commands used to verify the EIGRP configuration on Cisco IOS XE Catalyst SD-WAN devices.

View IPv4 EIGRP routes

Use the `show ip route vrf <vpn-id>` command to view IPv4 EIGRP routes for a specific VRF.

```
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

Use the `show ip route vrf <vpn-id>` command to view IPv6 EIGRP routes for a specific VRF.

```
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

Use the `show eigrp address-family ipv4 vrf <vpn-id> topology <ipv4-prefix>` command to view OMP routes that have been redistributed into EIGRP for a specific VRF and IPv4 prefix.

```
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)
```

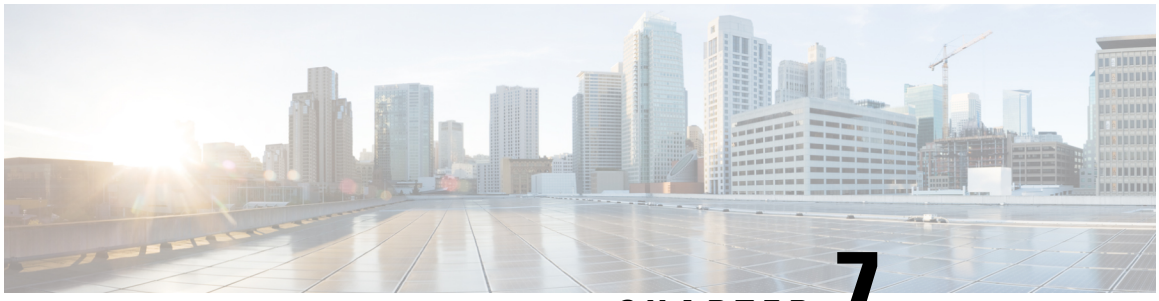

CHAPTER 7

Routing Information Protocol

The Routing Information Protocol (RIP) is a distance-vector routing protocol that uses broadcast or multicast User Datagram Protocol (UDP) data packets to exchange routing information. Commonly deployed in small to medium TCP/IP networks, RIP calculates routes based on hop count and sends routing-update messages at regular intervals and upon network topology changes.

Feature history for RIP

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

Table 27: 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	RIPv2 support on Cisco IOS XE Catalyst SD-WAN devices enables you to configure RIPv2 on these 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

Feature Name	Release Information	Description
RIPng (IPv6) support on Cisco IOS XE Catalyst SD-WAN devices	Cisco IOS XE Catalyst SD-WAN Release 17.10.1a Cisco vManage Release 20.10.1	RIPng (IPv6) Support on Cisco IOS XE Catalyst SD-WAN devices adds support for IPv6 addresses and prefixes. This feature facilitates redistribution of connect, static, Overlay Management Protocol (OMP), and Open Shortest Path First (OSPF) routes into Routing Information Protocol next generation (RIPng).

RIP in Cisco Catalyst SD-WAN

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

RIPv2 (RIP for IPv4)

RIP Version 2 (RIPv2) is the Cisco IOS software implementation of RIP for IPv4 networks. Each RIP process maintains a local database containing best-cost RIP routes learned from neighboring RIP-enabled routers. Route redistribution allows routes to be specified by a prefix, using a route map and prefix list.

RIPng (RIP for IPv6)

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

Key Features of RIPv2

The Cisco implementation of RIPv2 supports:

- Plain text and Message Digest Algorithm 5 (MD5) authentication.
- Route summarization.
- Classless Inter-Domain Routing (CIDR).
- 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.

Default Behavior and Configuration (RIPv2)

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, or only RIPv2 packets. The RIP

version that an interface sends can be configured to override this default behavior, and the processing of received packets can also be controlled. RIPv2 is supported on both the service side and transport side.

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

RIPng Redistribution Support

As an Interior Gateway Protocol (IGP), RIPng supports the redistribution of:

- 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: The 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.

RIPng VRF-aware Support

When not in Virtual Routing and Forwarding (VRF) mode, every IPv6 RIPng process and its associated configuration maintain all routes in the same routing table. IPv6 RIPng VRF-aware support enhances isolation, modularity, and potential performance by reducing the number of routes stored in a single routing table. It also allows network administrators to create different RIP routing tables and share a single protocol configuration block.

RIPng Loop Avoidance

RIPng in large networks is prone to routing loops, which can result in traffic taking longer paths. To avoid route looping, RIP and RIPng routes are identified using the well-known OMP RIP tag.

RIPv2 and RIPng OMP Route Tagging Process

The following process illustrates how OMP route tagging helps avoid routing loops in RIPv2 and RIPng:

Figure 3: RIPv2 and RIPng Topology

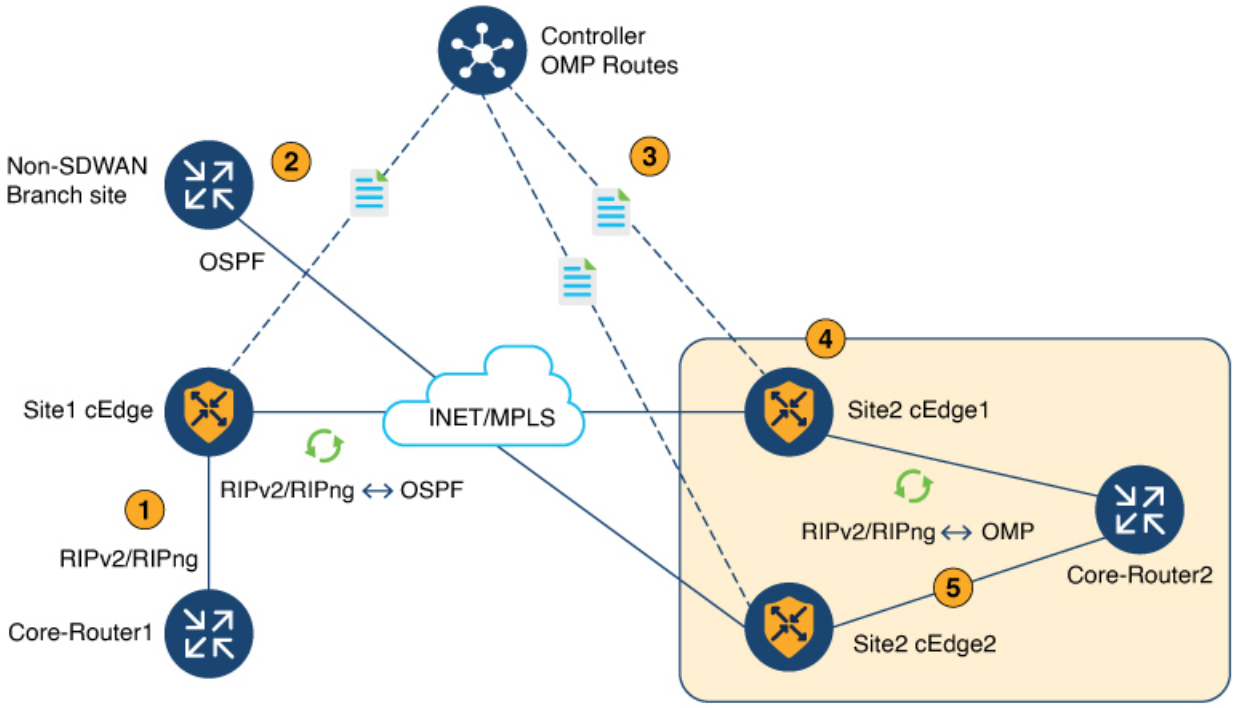

1. Core-Router1 advertises RIPv2 and RIPng routes to Site1.

   RIPv2 and RIPng routes have a default administrative distance of 120, while OMP routes have a default administrative distance of 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 (e.g., 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 an administrative distance (AD) of 251 (preferred over 120).

   If the OMP route is withdrawn, the Site-2 Edge2 router installs the route learned through the RIPv2 and RIPng protocol via a 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, on page 104 for more details on RIPng configurations using the CLI.

Prerequisites for using RIP

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

Restrictions for using RIP

When implementing RIP or RIPng (IPv6), be aware of the following limitations and restrictions:

RIPv2 (IPv4) restrictions

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

RIPng (IPv6) restrictions

- 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 RIPv2 using the CLI

Use this task to configure RIPv2 on Cisco IOS XE Catalyst SD-WAN devices via the Command Line Interface (CLI).

This section provides detailed instructions for configuring RIPv2, including process setup, VRF-aware support, route summarization, and various other RIP parameters. You can apply the configuration using CLI device templates or CLI Add-on feature templates.

Before you begin

You must complete the initial VRF routing table and address family submode configurations on the device. You can run these commands in any order.

Follow these steps to configure RIPv2 using the CLI:

Procedure

Step 1 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)#
```

Step 2 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
```

Step 3 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}
```

Step 4 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
```

Step 5 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
```

Step 6 Configure interpacket delay.

Configure interpacket delay for outbound RIP updates, in milliseconds:

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

Step 7 Redistribute the routes into the RIP routing process.

Redistribute the specified routes into the IPv4 RIP routing process. Only configure protocol redistribution after you configure the source router protocols. The protocol argument can be one of these keywords.

- **bgp**
- **connected**
- **isis**
- **eigrp**
- **omp**
- **ospf**
- **ospfv3**
- **static**
- **static**

In Cisco IOS XE Catalyst SD-WAN Release 17.7.1aa, 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]
```

Step 8 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]
```

Step 9 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. Use classful (Class A, Class B, Class C) IP network ID addressing for network configurations.

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

a) 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
```

b) Apply an offset list to routing metrics:

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

c) Adjust routing protocol timers:

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

Step 10 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
```

Step 11 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
```

Step 12 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
```

RIPv2 is configured on the Cisco IOS XE Catalyst SD-WAN device.

What to do next

Verify the RIPv2 configurations using the CLI. For more details, see the Verification commands for RIPv2 configuration, on page 102.

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
!
```

Verification commands for RIPv2 configuration

This section details the essential verification commands used to confirm the RIPv2 configurations and operational status on Cisco IOS XE Catalyst SD-WAN devices. These commands are crucial for inspecting configured settings, routing table entries, RIP database details, and neighbor information to ensure proper RIP functionality within the network.

View RIP routing configurations in the running configuration

```
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
```

Display 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
```

Display RIP routes under a specific 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
```

Display the contents of the 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
```

Display 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
```

Display RIP protocol configurations

```
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

Use this task to configure Routing Information Protocol next generation (RIPng) on Cisco IOS XE Catalyst SD-WAN devices via the Command Line Interface (CLI).

This section explains how to configure RIPng. It covers VRF-aware support, routing process setup, route tagging, and interface-specific parameters. You can use CLI device templates and CLI Add-on feature templates to set up RIPng.

Before you begin

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

Use these steps to configure RIPng using the CLI::

Procedure

Step 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
```

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

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

For **ripng-instance**, use **sdwan**.

Step 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)#
```

Step 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
```

Step 5 Configure a route tag.

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

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

Step 6 Create an entry in the IPv6 prefix list:

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

Step 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]
```

Step 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]
```

Step 9 Configure the interface.

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

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:

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.

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.

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}
```

RIPng is configured on the Cisco IOS XE Catalyst SD-WAN device

What to do next

Verify the RIPng configurations using the CLI. For more details, see the Verification commands for RIPng configuration, on page 106.

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
!
```

Verification commands for RIPng configuration

This section details the essential verification commands used to confirm the Routing Information Protocol next generation (RIPng) configurations and operational status on Cisco IOS XE Catalyst SD-WAN devices. These commands are crucial for inspecting configured settings and routing table entries to ensure proper RIPng functionality within the network.

Display RIPng routes in the VRF routing table

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
```

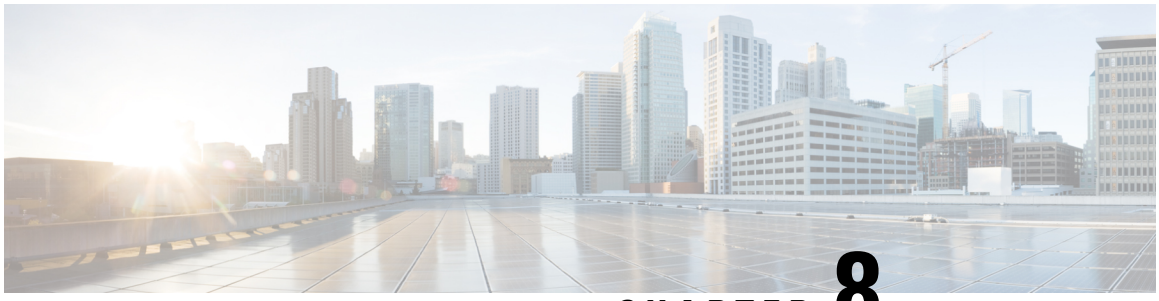

CHAPTER 8

Multicast Overlay Routing

Feature history for multicast overlay

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

Table 28: Feature history

Feature Name	Release Information	Description
Support for multicast overlay routing Protocols	Cisco IOS XE Catalyst SD-WAN Release 17.2.1r	This feature enables efficient distribution of one-to-many traffic. Multicast routing protocols such as IPv4 Multicast, IGMPv3, PIM SSM, PIM ASM, Auto RP, and Static RP distribute data, such as audio or video streaming broadcasts, to multiple recipients. Using multicast overlay protocols, a source can send a single packet of data to a single multicast address, which is then distributed to an entire group of recipients.
Multicast over L3 TLOC extension	Cisco IOS XE Catalyst SD-WAN Release 17.3.2 Cisco SD-WAN Release 20.3.1	This feature enables support for transport location (TLOC), which allows addition of the peer's transport to avoid the extra cost of additional IP addresses. It also allows dynamic load balancing across multiple transports.
Dynamic rendezvous point (RP) selection by a PIM BSR	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco SD-WAN Release 20.5.1	This feature adds support for automatic selection of an RP candidate using a PIM BSR in an IPv4 multicast overlay. There is no single point of failure because every site has a local RP. A Cisco IOS XE Catalyst SD-WAN device device is selected as the RP, not a service-side device.

Feature Name	Release Information	Description
Support for MSDP to interconnect Cisco SD-WAN and non-SD-WAN domains	Cisco IOS XE Catalyst SD-WAN Release 17.11.1a Cisco Catalyst SD-WAN Control Components Release 20.11.1	This feature enables Multicast Source Discovery Protocol (MSDP) interoperability between Cisco IOS XE Catalyst SD-WAN devices in Cisco Catalyst SD-WAN and the devices in a non-SD-WAN setup. **Note** This feature does not provide support for MSDP peers formed between Cisco IOS XE Catalyst SD-WAN devices in the overlay network.
Multicast support for hub and spoke topologies	Cisco IOS XE Catalyst SD-WAN Release 17.15.1a Cisco Catalyst SD-WAN Manager Release 20.15.1	This feature enables efficient distribution of traffic on edge devices using hub-and-spoke network topology. Multicast routing protocols such as IPv4 Multicast, IGMPv3, PIM SSM, PIM ASM, Auto RP, and Static RP distribute data to multiple recipients.

Multicast overlay routing for Cisco IOS XE Catalyst SD-WAN

A Cisco IOS XE Catalyst SD-WAN multicast overlay is a routing protocol that:

- extends Protocol Independent Multicast Source-Specific Multicast (PIM-SSM) over the Cisco Catalyst SD-WAN overlay using Overlay Management Protocol (OMP),
- integrates PIM-SM in customer VPNs with OMP in the overlay, leverages Cisco IOS XE MVPN, and uses OMP replicators to optimize the multicast distribution tree across the overlay topology, and
- supports IGMPv2 and IGMPv3 reports, advertising receiver multicast interest to remote Cisco Catalyst SD-WAN routers using OMP and enabling dynamic join or prune actions for optimized and secure multicast delivery over the overlay network.

The Cisco IOS XE Catalyst SD-WAN multicast overlay implementation extends native multicast by creating a secure optimized multicast tree that runs on top of the overlay network.

Protocol Independent Multicast Sparse-Mode (PIM-SM) is deployed in the customer VPNs, and the OMP replicator is used in overlay multicast to optimize the multicast distribution tree.

The Cisco IOS XE Catalyst SD-WAN router advertises receiver's multicast interest using OMP and participates in join or prune actions with replicators, which use OMP to relay these actions to routers providing overlay connectivity to the PIM-RP or source.

Multicast overlay supported features

Note From Cisco IOS XE Catalyst SD-WAN Release 17.3.2, TLOC extension with multicast and multicast application-aware route policy features are supported.

- IPv4 Overlay Multicast (PIM SSM), IPv4 Overlay Multicast (PIM ASM)

- PIM-RP on IOS XE VPN
- Replicator with geo-location (GPS)
- Static RP and Auto-RP
- PIM Bootstrap Router (BSR)
- IGMP v2, IGMP v3, and PIM on service side
- IPSec and GRE Encapsulation
- vEdge and IOS XE Catalyst SD-WAN Interop
- Overlay Multicast Signaling using OMP

Multicast overlay supported protocols

The Cisco IOS XE Catalyst SD-WAN overlay multicast network supports the Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) and multicast template configurations on all the platforms.

- Protocol Independent Multicast
- Internet Group Management Protocol
- Multicast Source Discovery Protocol

Restrictions for multicast overlay routing

Multicast overlay routing does not support these features:

- MSDP/Anycast-RP on Cisco Catalyst SD-WAN routers.
- IPv6 overlay and IPv6 underlay.
- Dynamic BFD tunnel for multicast.
- Multicast with asymmetric unicast routing.
- Multicast overlay working does not support Data Policy. If a data policy is configured, only the required traffic is matched and multicast traffic is not matched.
- The Cisco vEdge device is used only as the Last Hop Router (LHR), whereas Cisco IOS XE Catalyst SD-WAN devices can be used in all multicast roles (FHR, LHR, RP and Replicator roles).
- Bidirectional PIM is not supported with hub-and-spoke or full-mesh deployments.
- On Cisco 1000 Series Integrated Services Routers, when IGMP snooping is enabled and there are no local receivers for multicast traffic in the VLAN, the multicast traffic floods to all ports in the VLAN.

Restrictions for multicast routing with hub-and-spoke topology

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.15.1a and Cisco Catalyst SD-WAN Manager Release 20.15.1.

- You can configure multicast rendezvous point and replicator node on hub-site devices only. Replicator cannot be configured on spoke-site devices.
- MSDP interconnect feature is not supported with hub-and-spoke multicast deployment.
- You can configure multicast routing on hub-and-spoke using CLI add-on template only.
- On-demand tunnel between spoke sites is not supported with multicast.
- Multicast is supported only with centralized control policy–based hub-and-spoke deployment. Intent-based configuration, as described in the Hub-and-Spoke chapter, is not supported.

Configure multicast overlay routing

Prerequisites:

1. If you want to limit rendezvous point (RP) selection, configure an IPv4 ACL using a CLI add-on template. Configure an IPv4 ACL using a standard or extended access list and attach it to your device before enabling PIM.

   You must have created a valid standard or extended ACL prior to using the ACL in your multicast configuration.
2. If you want to limit rendezvous point (RP) selection, configure an IPv4 ACL using a CLI add-on template. Configure an IPv4 ACL using a standard or extended access list and attach it to your device before enabling PIM.
3. You cannot configure an ACL for a PIM feature template using Cisco SD-WAN Manager. You must configure the ACL using a CLI add-on template. The Cisco IOS XE Catalyst SD-WAN multicast overlay implementation supports IOS XE standard or extended access lists.
4. At least one replicator is mandatory for overlay multicast configuration.
5. You can optionally configure IGMP to allow individual hosts on the service side to join multicast groups within a particular VPN.

- Configure multicast using configuration groups
- Configure multicast using device templates
- Configure multicast using the CLI
- Configure an ACL for multicast using a CLI Add-On Template
- Configure PIM
- Rendezvous point selection process by a PIM BSR
- Configure IGMP
- Configure PIM and IGMP using the CLI
- Configure MSDP using a CLI template

Configure multicast using configuration groups

From Cisco IOS XE Catalyst SD-WAN Release 17.11.1a you have the option to configure multicast using Configuration Groups.

Procedure

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

Step 2 Click **…** adjacent to the configuration group name and choose **Edit**.

Step 3 Click **Service Profile**.

Step 4 Click **Add Feature**.

Step 5 From the feature drop-down list, choose **Multicast**.

The Cisco IOS XE Catalyst SD-WAN overlay multicast network supports the following protocols:

- PIM
- IGMP
- MSDP

The following tables describe the options for configuring the Multicast feature.

Field	Description
Type	Choose a feature from the drop-down list.
Feature Name*	Enter a name for the feature.
Description	Enter a description of the feature. The description can contain any characters and spaces.

Table 29: Basic Configuration

Field	Description
SPT Only	Enable this option to ensure that the Rendezvous Points (RPs) can communicate with each other using the shortest-path tree.
Local Replicator	Enable this option to configure the Cisco IOS XE Catalyst SD-WAN device as a multicast replicator.
Threshold	Specify a value. Optional, keep it set to the default value if you are not configuring a replicator.

Table 30: PIM

Field	Description
Source Specific Multicast (SSM)	Enable this option to configure SSM.

Field	Description
ACL	Specify an access control list value. An access control list allows you to filter multicast traffic streams using the group and sometimes source IPv4 or IPv6 addresses. Configure an IPv4 access control list using a standard or extended access list and attach it to your device before enabling PIM. You must have created a valid standard or extended ACL before using the ACL in your multicast configuration. **Note** You cannot configure an ACL for a PIM feature template using Cisco SD-WAN Manager. You must configure the ACL using a CLI add-on template. For information on configuring ACL using the CLI add-on template, see the section **Configure an ACL for Multicast Using a CLI Add-On Template** in chapter **Multicast Overlay Routing** of the Cisco SD-WAN Routing Configuration Guide.
SPT Threshold	Specify the traffic rate, in kbps, at which to switch from the shared tree to the shortest-path tree (SPT). Configuring this value forces traffic to remain on the shared tree and travel via the RP instead of via the SPT.
Add Interface	
Interface Name	Enter the name of an interface that participates in the PIM domain, in the format **ge** *slot /port*.
Query Interval(sec)	Specify how often the interface sends PIM query messages. Query messages advertise that PIM is enabled on the router.
Join/Prune Interval(sec)	Specify how often PIM multicast traffic can join or be removed from a rendezvous point tree (RPT) or shortest-path tree (SPT). Cisco IOS XE Catalyst SD-WAN device send join and prune messages to their upstream RPF neighbor.
How do you want to configure your Rendezvous Point (RP) Cisco IOS XE SD-WAN supports the following modes:	
Static	Click this check box to a specify the static IP address of a rendezvous point (RP).
Add Static RP	
IP Address	Specify the static IP address of a rendezvous point (RP).
ACL	Specify an ACL value.
Override	Enable this option for cases when dynamic and static group-to-RP mappings are used together and there is an RP address conflict. In this case, the RP address configured for a static group-to-RP mapping takes precedence. If you do not enable this option, and there is RP address conflict, dynamic group-to-RP mappings will take precedence over static group-to-RP mappings.
Auto RP	Click this check box to enable reception of PIM group-to-RP mapping updates. This enables reception on the Auto-RP multicast groups, 224.0.1.39 and 224.0.1.40.
RP Announce	Click this check box to enable transmission of Auto-RP multicast messages.

Field	Description
RP Discovery	Click this check box to enable Auto-RP automatic discovery of rendezvous points (RPs) in the PIM network so that the router can serve as an Auto-RP mapping agent. An Auto-RP mapping receives all the RPs and their respective multicast groups and advertise consistent group-to-RP mapping updates.
Interface	Specify the source interface for Auto-RP RP Announcements or RP Discovery messages.
Scope	Specify the IP header Time-to-Live (TTL) for Auto-RP RP Announcements or RP Discovery messages.
PIM-BSR	Configure a PIM BSR.
RP Candidate	
Interface Name	Choose the interface that you used for configuring the PIM feature template.
Access List	Add an access list value if you have configured the access list with a value.
Interval	Add an interval value if you have configured the interval with a value.
Priority	Specify a higher priority on the Cisco IOS XE SD-WAN device than on the service-side device.
BSR Candidate (Maximum: 1)	
Interface Name	Chose the same interface from the drop-down list that you used for configuring the PIM feature template.
Hash Mask Length	Specify the hash mask length. Valid values for hash mask length are 0–32.
Priority	Specify a higher priority on the Cisco IOS XE Catalyst SD-WAN device than on the service-side device.
RP Candidate Access List	Add a value if you have configured the RP candidate access list with a value. An RP candidate uses a standard ACL where you can enter the name for the access list.

Table 31: IGMP

Field	Description
Add IGMP	
Interface	Enter the name of the interface to use for IGMP. To add another interface, click **Add**.
Version	Specify a version number. Optional, keep it set to the default version number.
Group Address	Enter a group address to join a multicast group.

Field	Description
Source Address	Enter a source address to join a multicast group.
Add	Click **Add** to add the IGMP for the group.

Table 32: MSDP

Field	Description
Originator-ID	Specify the ID of the originating device. This ID is the IP address of the interface that is used as the RP address.
Connection Retry Interval	Configure an interval at which MSDP peers will wait after peering sessions are reset before attempting to re-establish the peering sessions.
Mesh Group	
Mesh Group Name	Enter a mesh group name. This configures an MSDP mesh group and indicates that an MSDP peer belongs to that mesh group. **Note** All MSDP peers present on a device that participate in a mesh group must be in a full mesh with all other MSDP peers in the group. Each MSDP peer on each device must be configured as a peer using the **ip msdp peer** command, and as a member of the mesh group using the **ip msdp mesh-group** command.
Peer-IP	Configure an MSDP peer specified by an IP address.
Advanced Settings	
Connect-Source Interface	Enter the primary address of a specified local interface that is used as the source IP address for the TCP connection.
Peer Authentication Password	Enables MD5 password encryption for a TCP connection between two MSDP peers. **Note** MD5 authentication must be configured with the same password on both MSDP peers. Otherwise, a connection between them cannot be established.
Keep Alive	Configure an interval at which an MSDP peer will send keepalive messages.
Hold-Time	Configure an interval at which the MSDP peer will wait for keepalive messages from other peers before declaring them as down.
Remote AS	Specifies the autonomous system number of the MSDP peer. This keyword and argument are used for display purposes only.
SA Limit	Limits the number of SA messages allowed in the SA cache from the specified MSDP.

Field	Description
Default Peer	Configure a default peer from which to accept all MSDP SA messages.

Configure multicast using the CLI

To configure multicast, excute this command:

Procedure

sdwan multicast address-family ipv4 vrf *1* **replicator** [**threshold** <*num*>]

Example:

```
Device(config)# sdwan
Device(config)# multicast
  Device(config)# address-family ipv4 vrf 1
  Device(config)# replicator threshold 7500
  Device(config)# !
 !
```

Configure multicast using Cisco SD-WAN Manager

When a Cisco IOS XE Catalyst SD-WAN router is used as a replicator, follow these steps to configure multicast:

Procedure

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

Step 2 Click **Device Templates.**

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

Step 4 Click **Service VPN** in the **Service VPN** section. Click the **Service VPN** drop-down list.

Step 5 Under **Additional VPN Templates**, click **Multicast**.

Step 6 To enable **Local Replicator** on the device, choose **On** (otherwise keep it **Off**).

Step 7 To configure replicator, choose the **Threshold**. (Optional, keep it default if you are not configuring replicator).

Step 8 Save feature template. Attach feature template to device template.

Step 9 In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

What to do next

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 the value.

Configure an ACL for multicast using a CLI Add-On template

You can configure an ACL to limit RP and Bootstrap Router (BSR) selection using a CLI add-on template. An ACL allows you to filter multicast traffic streams using the group and sometimes source IPv4 or IPv6 addresses.

Before you begin

Once you create the CLI add-on template, you attach it to the device.

Procedure

Step 1 To configure an ACL for multicast, *create a CLI add-on feature template and attach it to the device template*.

Example:

```
ip access-list standard 27
1 permit 225.0.0.0 0.255.255.255
2 permit 226.0.0.0 0.255.255.255
3 permit 227.0.0.0 0.255.255.255
4 permit 228.0.0.0 0.255.255.255
5 deny 229.0.0.0 0.255.255.255
6 permit any
ip access-list extended 101
1 permit pim 172.16.10.0 0.0.0.255 any
2 permit pim 10.1.1.0 0.0.0.255 any
```

Step 2 From the **Configuration** > **Templates** window, choose **Feature**.

Step 3 Edit the **Cisco PIM** feature template that you configured for the RP or the BSR candidate by clicking **…** and then clicking **Edit**.

For more information, see Configure a PIM BSR.

Step 4 (Optional) In the **Access List** field for the configured RP candidate, enter the same ACL value as you configured in the CLI add-on template.

Step 5 (Optional) In the **RP Candidate Access List** field for the configured BSR candidate, enter the same ACL value as you configured in the CLI add-on template.

Step 6 Update the feature template and attach the feature template to the device template.

Configure PIM

Use the PIM template for all Cisco IOS XE Catalyst SD-WAN devices.

Configure the PIM Sparse Mode (PIM-SM) protocol using Cisco SD-WAN Manager templates so that a router can participate in the Cisco IOS XE Catalyst SD-WAN multicast overlay network:

Procedure

Create a PIM feature template to configure PIM parameters.

a) From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
b) Click **Device Templates**.

In Cisco Catalyst SD-WAN Control Components Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

c) Click **Create Template**. From the **Create Template** drop-down list, choose **From Feature Template**.
d) From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
e) Click **Service VPN**. Under **Additional VPN Templates**, click **PIM**.
f) From the **PIM** drop-down list, click **Create Template**. The PIM template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining PIM parameters.
g) In the **Template Name** field, enter a name containing up to 128 alphanumeric characters. In the **Template Description** field, enter a description containing 2048 alphanumeric characters.
h) Click **Basic Configuration** and configure SSM – On/Off.
i) Configure access list (if already defined), RP option (Auto-RP or static RP), RP Announce settings and configure the interface name on the service side.

Save feature template and attach feature template to a device template.

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 the value.

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 IOS XE Catalyst 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. You upload the CSV file when you attach a Cisco IOS XE Catalyst 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	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.

To configure PIM, click **Basic Configuration** and configure the following parameters. Parameters marked with an asterisk are required to configure PIM.

Table 33: Basic Configuration

Parameter Name	Description
Auto-RP	Click **On** to enable Auto-RP to enable reception of PIM group-to-RP mapping updates. This will enable reception on the Auto-RP multicast group, 224.0.1.39 and 224.0.1.40. By default, Auto-RP is disabled.
Auto-RP RP Announce	Click **On** to enable transmission of Auto-RP multicast messages. By default, RP Announce is disabled.
Auto-RP RP Discovery	Click **On** to enable Auto-RP automatic discovery of rendezvous points (RPs) in the PIM network so that the router can serve as an Auto-RP mapping agent. An Auto-RP mapping will receive all the RPs and their respective multicast groups and advertise consistent group-to-RP mapping updates. By default, RP Discovery is disabled.
Static-RP	Specify the IP address of a rendezvous point (RP).
SPT Threshold	Specify the traffic rate, in kbps, at which to switch from the shared tree to the shortest-path tree (SPT). Configuring this value forces traffic to remain on the shared tree and travel via the RP instead of via the SPT.
Interface	Specify the source interface for Auto-RP RP Announcements or RP Discovery messages.
Scope	Specify the IP header Time-to-Live (TTL) for Auto-RP RP Announcements or RP Discovery messages.

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

If the router is just a multicast replicator and is not part of a local network that contains either multicast sources or receivers, you do not need to configure any PIM interfaces. The replicator learns the locations of multicast sources and receivers from the OMP messages it exchanges with the Cisco Catalyst SD-WAN Controller. These control plane messages are exchanged in the transport VPN (VPN 0). Similarly, other Cisco IOS XE Catalyst SD-WAN devices discover replicators dynamically, through OMP messages from the Cisco Catalyst SD-WAN Controller.

To configure PIM interfaces, click **Interface**. Then click **Add New Interface** and configure the following parameters:

Parameter Name	Description
Name	Enter the name of an interface that participates in the PIM domain, in the format **ge** *slot /port*.
Hello Interval	Specify how often the interface sends PIM hello messages. Hello messages advertise that PIM is enabled on the router. Range: 1 through 3600 seconds Default: 30 seconds
Join/Prune Interval	Specify how often PIM multicast traffic can join or be removed from a rendezvous point tree (RPT) or shortest-path tree (SPT). Cisco IOS XE Catalyst SD-WAN send join and prune messages to their upstream RPF neighbor. Range: 0 through 600 seconds Default: 60 seconds

To edit an interface, click the pencil icon to the right of the entry.

To delete an interface, click the trash icon to the right of the entry.

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

Configure IGMP

Use the IGMP template for all Cisco IOS XE Catalyst SD-WAN devices. Internet Group Management Protocol (IGMP) allows routers to join multicast groups within a particular VPN.

Before you begin

To configure IGMP using Cisco SD-WAN Manager templates:

1. Create an IGMP feature template to configure IGMP parameters.
2. Create the interface in the VPN to use for IGMP. See the VPN-Interface-Ethernet help topic.
3. Create a VPN feature template to configure VPN parameters. See the VPN help topic

Procedure

Step 1 Navigate to the **Template Window** and Name the **Template**.

a) From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates** .
b) Click **Device Templates**.

In Cisco SD-WAN Release 20.7.x and earlier releases, Device Templates is titled Device.

c) Click **Create Template**. From the **Create Template** drop-down list, choose **From Feature Template**.
d) From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
e) Click **Service VPN**. Click the **Service VPN** drop-down list.
f) Under **Additional VPN Templates**, click **IGMP**.
g) From the **IGMP** drop-down list, click **Create Template**. The IGMP template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining IGMP parameters.
h) Add interface name on the service side to enable IGMP

(Optional) In the **Join Group And Source Address** field, click on **Add Join Group and Source Address**. The **Join Group and Source Address** window displays.

(Optional) Enter group address to join and source address.

i) In the **Template Name** field, enter a name containing up to 128 characters and only alphanumeric characters. In the **Template Description** field, enter a description containing up to 2048 characters and 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 to the left of the parameter field and select the value.

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 IOS XE Catalyst 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. You upload the CSV file when you attach a Cisco IOS XE Catalyst 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	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.

Step 2 Configure Basic IGMP Parameters

To configure IGMP, click **Basic Configuration** to enable IGMP. Then, click **Interface**, and click **Add New Interface** to configure IGMP interfaces. All parameters listed below are required to configure IGMP.

Parameter Name	Description
Interface Name	Enter the name of the interface to use for IGMP. To add another interface, click the plus sign (+).
Join Group Address	Optionally, click **Add Join Group Address** to enter a multicast group. Click **Add** to add the IGMP for the group.

Step 3 To save the feature template, click **Save**.

Configure PIM and IGMP using the CLI

For a Cisco IOS XE Catalyst SD-WAN router located at a site that contains one or more multicast sources, enable PIM on the service-side interface or interfaces. These are the interfaces that connect to the service-side network.

To enable PIM or IGMP per VPN, you must configure PIM or IGMP and its interfaces for all VPNs support multicast services.

PIM configuration is not required in VPN 0 (the transport VPN facing the overlay network) or in VPN 512 (the management VPN).

If a source interface is specified in the `send-rp-discovery` container, ensure that the interface already has an IP address and PIM configured.

Sample configuration:

```
vrf definition 1
 rd 1:1
 address-family ipv4
  exit-address-family
 !
!
ip pim vrf 1 autorp listener
ip pim vrf 1 send-rp-announce Loopback1 scope 12 group-list 10
ip pim vrf 1 send-rp-discovery Loopback1 scope 12
ip pim vrf 1 ssm default
ip access-list standard 10
 10 permit 10.0.0.1 0.255.255.255
!
ip multicast-routing vrf 1 distributed
interface GigabitEthernet0/0/0.1
 no shutdown
 encapsulation dot1Q 1
 vrf forwarding 1
 ip address 172.16.0.0 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
 ip ospf 1 area 0
exit
interface GigabitEthernet0/0/2
 no shutdown
 vrf forwarding 1
 ip address 172.16.0.1 255.255.255.0
 ip pim sparse-mode
 ip ospf 1 area 0
exit
interface Loopback1
 no shutdown
 vrf forwarding 1
 ip address 192.0.2.255 255.255.255.255
 ip pim sparse-mode
 ip ospf 1 area 0
exit
sdwan
 multicast
  address-family ipv4 vrf 1
   replicator threshold 7500

   exit
```

Configure MSDP using a CLI template

Before you begin

- By enabling an MSDP peer, you implicitly enable MSDP.
- IP multicast routing must be enabled and PIM-SM must be configured. For more information, see Configure PIM.
- By default, CLI templates execute commands in global config mode. For more information about using CLI templates, see *CLI Add-On Feature Templates* and *CLI Templates*.

Procedure

Step 1 Enable MSDP and configure an MSDP peer as specified by the DNS name or IP address.

If you specify the **connect-source** keyword, the primary address of the specified local interface type and number values are used as the source IP address for the TCP connection. The **connect-source** keyword is recommended, especially for MSDP peers on a border that peer with a device inside of a remote domain.

Example:

```
ip msdp peer peer ip address connect-source
```

Step 2 Configure an originating address.

Perform this optional task to allow an MSDP speaker that originates an SA message to use the IP address of its interface as the RP address in the SA message.

You can also change the originator ID for any one of the following reasons:

- If you configure multiple devices in an MSDP mesh group for Anycast RP.
- If you have a device that borders a PIM-SM domain and a PIM-DM domain. If a device borders a PIM-SM domain and a PIM-DM domain and you want to advertise active sources within the PIM-DM domain, configure the RP address in SA messages to be the address of the originating device's interface.

Example:

```
ip msdp originator-id type number
```

Step 3 Configure an MSDP Mesh Group to indicate that an MSDP peer belongs to that mesh group.

You can configure multiple mesh groups per device.

Note
All MSDP peers on a device that participate in a mesh group must be fully meshed with all other MSDP peers in the group. Each MSDP peer on each device must be configured as a peer using the **ip msdp peer** command and also as a member of the mesh group using the **ip msdp mesh-group** command.

```
ip msdp mesh-group mesh name{peer-ip address | peer name}
```

How traffic flows in multicast overlay routing

Summary

The following illustration represents the example topology for multicast overlay routing on Cisco IOS XE Catalyst SD-WAN devices:

Workflow

Figure 4: Multicast Overlay Routing Topology

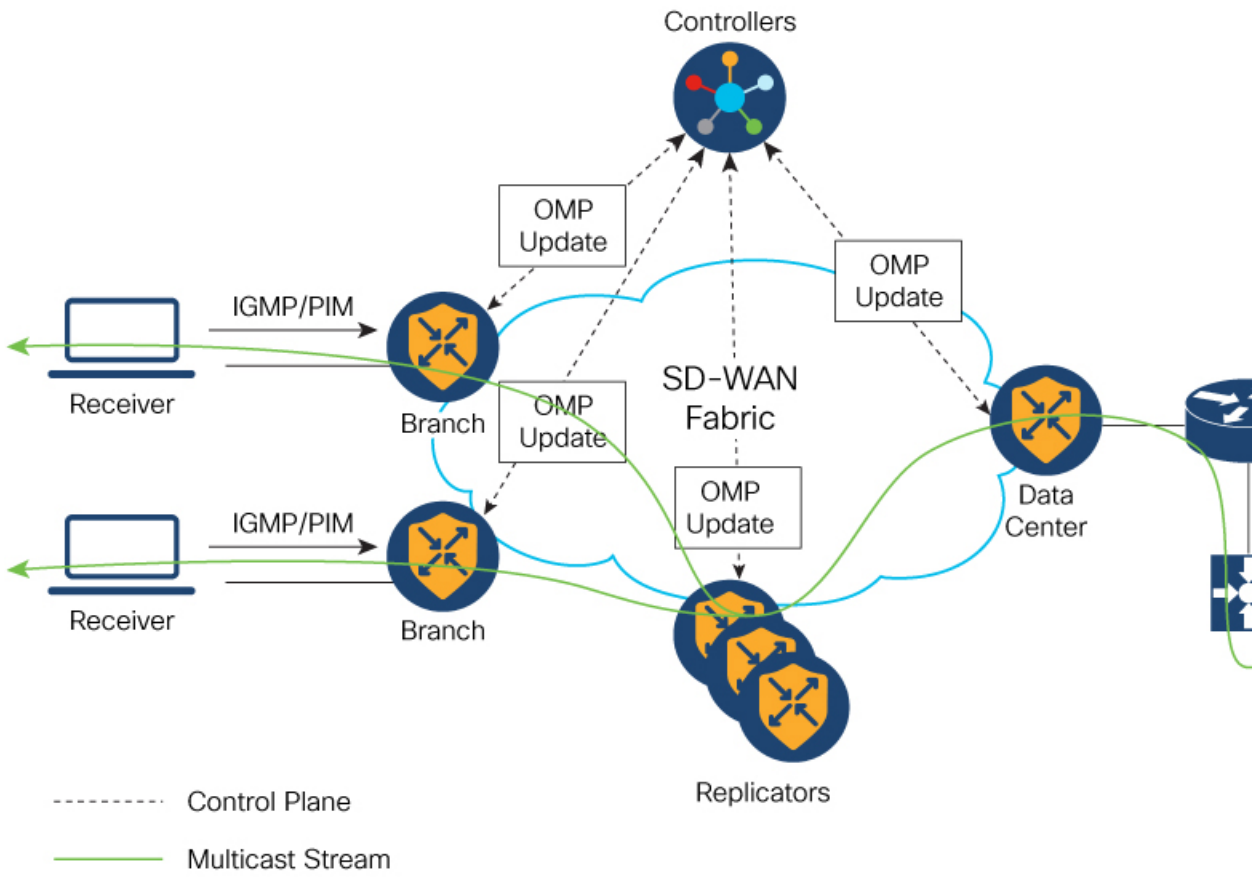

How multicast overlay protocols work in a hub-and-spoke topology

From Cisco IOS XE Catalyst SD-WAN Release 17.15.1a and Cisco Catalyst SD-WAN Manager Release 20.15.1, a source can send a single packet of data to a single multicast address using multicast overlay protocols in a hub-and-spoke topology. This single packet is then distributed to an entire group of recipients.

Use cases for multicast routing on hub-and-spoke

- A sender in a hub site sending multicast traffic to receivers in same or other hub sites.
- A sender in a hub site sending multicast traffic to receivers in spoke sites.
- A sender in a spoke site sending multicast traffic to receivers in hub sites.
- A sender in a spoke site sending multicast traffic to receivers in same/other spoke sites.

Summary

The illustration has the following configurations:

Workflow

Figure 5: Multicast Configuration

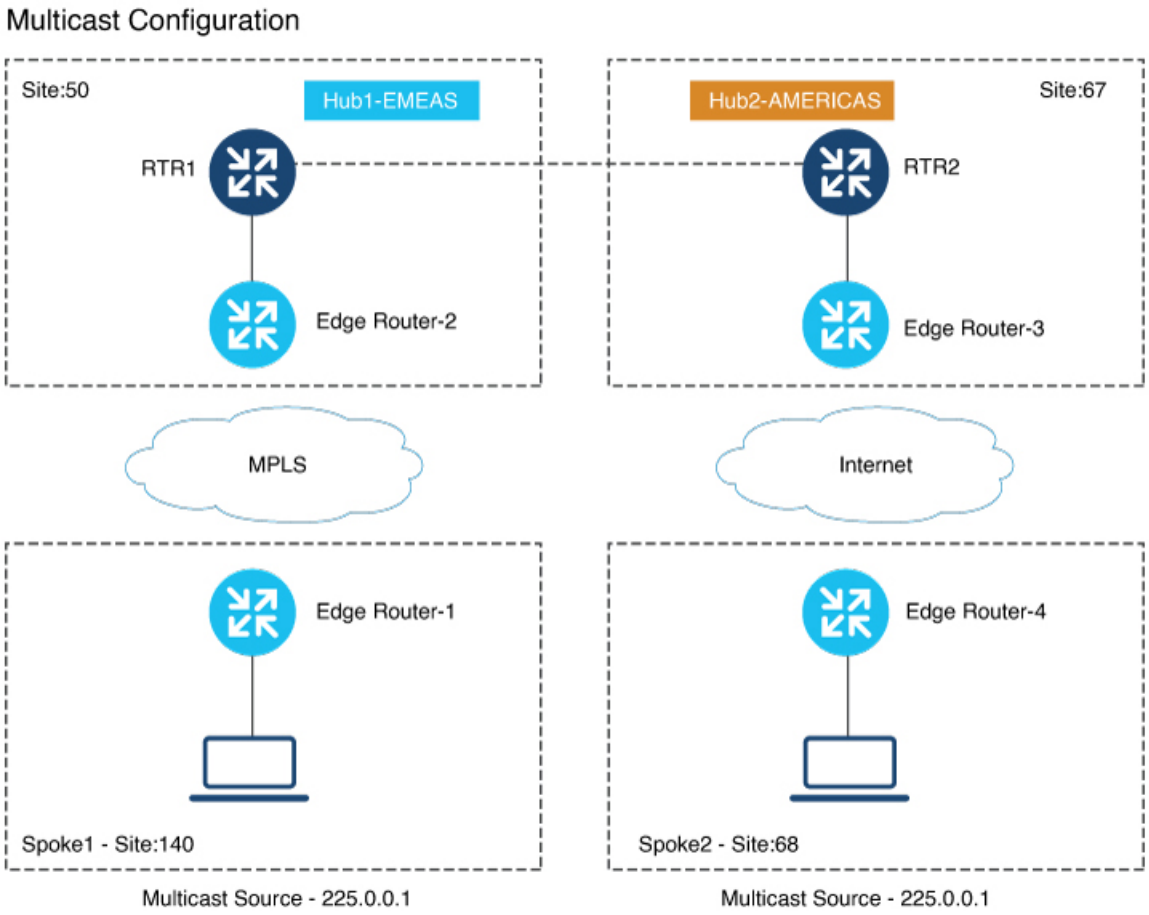

- Any Source Multicast (ASM) with static or AutoRP
- No BFD session between hub sites across different regions
- No BFD sessions between spoke sites
- BFD session must be present between hub sites and all the spoke sites across all regions
- For every site (both hub and spoke) define a control policy. The site-list of the policy specifies all hub and spoke sites excluding the site on which the policy is applied.
- There should be at least one unicast subnet with NextHop as Cisco IOS XE Catalyst SD-WAN device in the route table to forward multicast traffic. This is applicable to Hub-Spoke, Full Mesh and Dual Border scenarios too.
- Configuration Example of Hub-and-spoke Multicast Using the CLI

Configuration example of hub-and-spoke multicast using the CLI

The following example shows the configuration of centralized control policy for hub-and-spoke deployment:

```
policy
 lists
  tloc-list Hub-TLOCs
   tloc 10.10.10.2 color biz-internet encap ipsec
   tloc 192.0.2.1 color biz-internet encap ipsec
  !
  site-list Branches
   site-id 140
   site-id 68
  !
  site-list DCs
   site-id 50
   site-id 67
  !
 !
 control-policy Hub-Control-Policy
  sequence 11
   match tloc
    site-list DCs
   !
   action accept
   !
  !
  sequence 31
   match route
    site-list DCs
   !
   action accept
   !
  !
  default-action reject
 !
 control-policy Spoke-Control-Policy
  sequence 1
   match tloc
    site-list Branches
   !
   action reject
   !
  !
  sequence 11
   match tloc
    site-list DCs
   !
   action accept
   !
  !
  default-action reject
 !
!
apply-policy
 site-list Branches
  control-policy Spoke-Control-Policy out
 !
 site-list DCs
  control-policy Hub-Control-Policy out
 !
!
```

The following example shows the spoke configuration for hub-and-spoke multicast deployment:

```
sdwan
 multicast
  address-family ipv4 vrf 1
   spoke
  !
 !
!
```

Verify multicast routing on hub-and-spoke

Procedure

Use the command **show platform software sdwan multicast active-sources vrf** on spokes to verify multicast source active route next-hop pointing to the selected replicator.

Example:

```
Device# show platform software sdwan multicast active-sources vrf 1

Multicast SDWAN Overlay Received Source-Active Routes:
  (10.0.0.0, 255.0.0.0) next-hop: 192.168.255.254
      src-orig-count: 1, rp-addr: 10.0.0.1
```

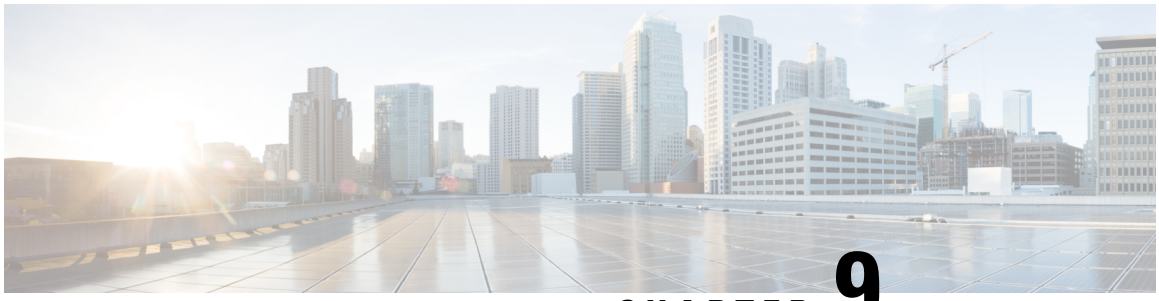

CHAPTER 9

Radio Aware Routing

Feature history for RAR

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

Table 34: Feature History

Feature name	Release information	Description
Radio-Aware Routing Support	Cisco IOS XE Catalyst SD-WAN Release 17.6.1a Cisco vManage Release 20.6.1	This feature enables Radio-Aware Routing (RAR) support on Cisco IOS XE Catalyst SD-WAN devices. RAR uses radio signals to interact with the OSPFv3 routing protocol. Through these signals it can indicate the appearance, disappearance, and link conditions of one-hop routing neighbors. In large mobile networks, connections to routing neighbors may be interrupted by distance and radio obstructions. RAR addresses the challenges that arise when integrating IP routing with radio communications in mobile networks.

RAR for Cisco IOS XE Catalyst SD-WAN devices

Radio-Aware Routing (RAR) is a mechanism that leverages radio interfaces to communicate directly with the Open Shortest Path First version 3 (OSPFv3) protocol, enabling real-time signaling of the presence and link conditions of one-hop routing neighbors.

- RAR improves routing responsiveness in mobile environments by allowing immediate link state updates.
- Standard protocol timers may be too slow for rapidly changing mobile network conditions.
- PPPoE provides the underlying connectivity, with support for OSPFv3 and EIGRP.
- An AX license is required to enable this feature.

In large mobile networks, factors such as distance and radio obstructions can frequently disrupt connections to routing neighbors. If these disruptions are not reported directly to the routing protocols, the protocols rely on their built-in timers to update neighbor status. However, these protocol timers are typically lengthy, which is not ideal for the dynamic conditions of mobile networks.

Connectivity between two Cisco IOS XE Catalyst SD-WAN devices is established over a PPPoE connection, which features variable bandwidth and limited buffering capacity. OSPFv3 and EIGRP are the supported routing protocols for this deployment.

Mobile Ad Hoc Networks (MANETs)

MANETs facilitate device-to-radio communications by addressing the challenges of integrating IP routing with mobile radio technologies in ad hoc networking environments. MANET routing protocols enable signaling between MANET routers, supporting features such as scope-limited flooding and point-to-point delivery of routing protocol messages within the network.

System components of RAR

The RAR feature uses the MANET infrastructure and includes several components: PPPoE, Virtual Multipoint Interface (VMI), QoS, routing protocol interface, and RAR protocols.

Point-to-Point Protocol over Ethernet PPPoE or PPPoE

PPPoE is a well-defined communication mechanism between the client and the server. In the RAR implementation, radio takes the role of the PPPoE client and router takes the role of the PPPoE server. This allows a loose coupling of radio and router, while providing a well-defined and predictable communication mechanism.

As PPPoE is a session or a connection-oriented protocol, it extends the point to point radio frequency (RF) link from an external radio to an IOS router.

PPPoE Extensions

In the Cisco IOS implementation of PPPoE, each session is represented by a virtual access interface, which connects to a radio neighbor. QoS can be applied on these interfaces using PPPoE extensions.

RFC5578 provides extensions to PPPoE to support credit-based flow control and session-based real-time link metrics, which are very useful for connections with variable bandwidth and limited buffering capabilities (such as radio links).

Virtual Multipoint Interface (VMI)

Though PPPoE Extensions provides most of the setup to communicate between a router and a radio, VMI addresses the need to manage and translate events that higher layers (for example, routing protocols) consume.

In bypass mode, every Virtual Access Interface (VAI) representing a radio neighbor is exposed to the routing protocols OSPFv3 and EIGRP so that the routing protocol directly communicates with the respective VAI for both unicast and multicast routing protocol traffic.

In Aggregate mode, VMI is exposed to the routing protocols (such as OSPF) so these protocols can optimize efficiency.

When the network neighbors are viewed as a collection of networks on a point-to-multipoint link with broadcast and multicast capability at VMI, VMI aggregates the multiple virtual access interfaces created from PPPoE. VMI provides a single, multi-access Layer 2 interface with broadcast capability. Using a single interface for

the routing protocol reduces the size of the topology database without impacting network integrity. The VMI layer:

- Redirects unicast routing protocol traffic to the appropriate P2P link (Virtual-Access interface)
- Replicates multicast or broadcast traffic as needed

RAR architecture

Summary

The topology shows the RAR deployment on Cisco IOS XE Catalyst SD-WAN devices.

Workflow

Figure 6: RAR architecture

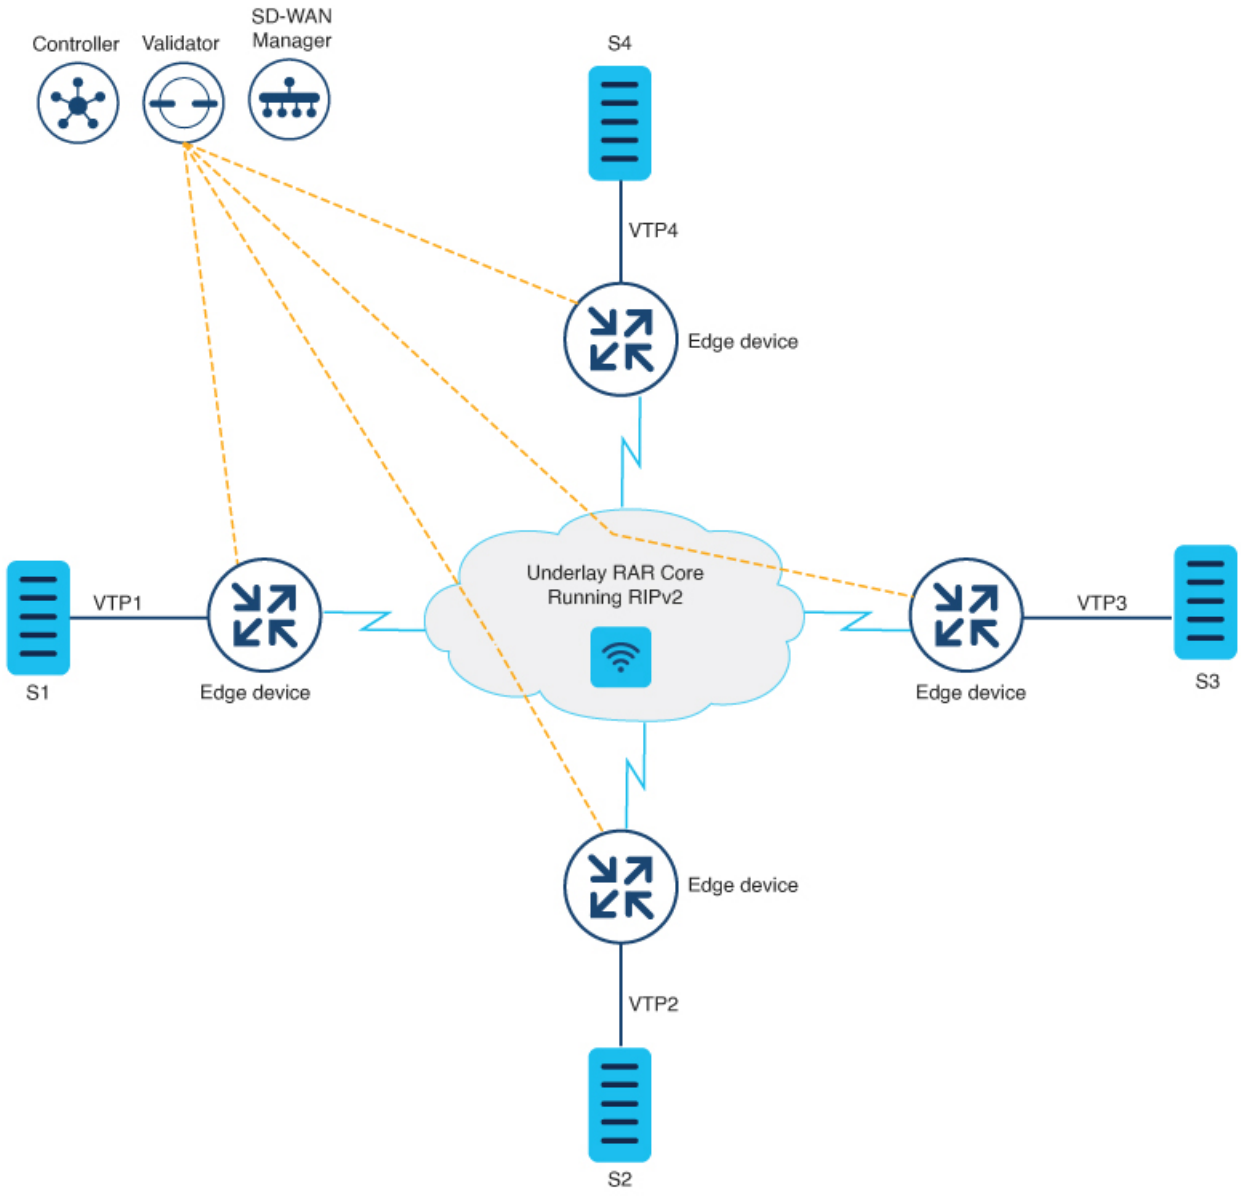

- The four Cisco IOS XE Catalyst SD-WAN devices connect to each other through a radio connected to a physical interface on the device
- PPPoE-RAR configurations happen on all three routers and once the underlay RAR network is established, the Cisco Catalyst SD-WAN tunnels form on the network.
- The loopback interface acts as a WAN interface and binds to the Virtual Multipoint interface (VMI). The VMI interface in turn binds to the physical interface
- The PPP connections between any two devices act as the underlay network.

- The Cisco Catalyst SD-WANCisco SD-WAN Manager tunnels are established over the PPPoE-RAR underlay network.
- Cisco SD-WAN Manager, Cisco SD-WAN Controller, and Cisco SD-WAN Validator connect through a radio connection in the deployment scenario.

Benefits of RAR

The RAR feature offers these benefits:

- Provides faster network convergence through immediate recognition of changes.
- Enables routing for failing or fading radio links.
- Allows easy routing between line-of-sight and non-line-of-sight paths.
- Provides faster convergence and optimal route selection so that delay-sensitive traffic, such as voice and video, is not disrupted
- Provides efficient radio resources and bandwidth usage.
- Reduces impact on the radio links by performing congestion control in the router.
- Allows route selection based on radio power conservation.
- Enables decoupling of the routing and radio functionalities.
- Provides simple Ethernet connection to RFC 5578, R2CP, and DLEP compliant radios.

Supported devices for RAR

These platforms support RAR:

- Cisco 4000 Series Integrated Services Routers
- Cisco 1000 Series Integrated Services Routers
- Cisco ASR 1000 Series Aggregation Services Routers
- Cisco CSR 1000 Series Cloud Service Routers
- Cisco CSR 8000 Series Cloud Service Routers

Restrictions for RAR

The Radio Aware Routing feature has these restrictions:

- The Dynamic Link Exchange Protocol (DLEP) and Router to Radio Control Protocol (R2CP) protocols are not supported.
- Multicast traffic is not supported in aggregate mode.
- Cisco High Availability technology is not supported.

Prerequisites for RAR

The RAR configuration requires Mobile Ad-hoc Networks (MANETs) support. To use the PPP over Ethernet (PPPoE) and virtual multipoint interface (VMI) features for RAR, a unified representation of the MANET to routing protocols (OSPFv3 or EIGRP) is required.

Configure RAR

To configure RAR using SD-WAN Manager, Create a CLI add-on feature template and attach it to the device template.

This section provides examples of RAR configurations that you can add to the CLI add-on template.

Configure a service for RAR:

```
policy-map type service rar-lab
 pppoe service manet_radio //note: Enter the pppoe service policy name as manet_radio
!
```

Configure OSPF routing:

```
router ospfv3 1
 router-id 10.0.0.1
!
 address-family ipv4 unicast
  redistribute connected metric 1 metric-type 1
  log-adjacency-changes
 exit-address-family
 !
 address-family ipv6 unicast
  redistribute connected metric-type 1
  log-adjacency-changes
 exit-address-family
!
ip local pool PPPoEpool2 192.0.2.0 192.0.2.1
```

Configuration of RAR:

```
interface GigabitEthernet0/0/0
no shutdown
no mop enabled
no mop sysid
negotiation auto
pppoe enable group PPPOE_RAR

interface vmi1
ip address 10.0.0.0 255.255.255.0
ipv6 enable
physical-interface GigabitEthernet0/0/0
mode bypass
exit
interface Virtual-Template1
no shutdown
ip unnumbered vmi1
ipv6 enable
ospfv3 1 network manet
ospfv3 1 ipv4 area 0
ospfv3 1 ipv6 area 0
exit
```

```
interface Tunnel100
 no shutdown
 ip unnumbered Loopback100
 tunnel source Loopback100
 tunnel mode sdwan
exit

interface Loopback100
  tunnel-interface
   encapsulation ipsec
   color mpls
   no allow-service bgp
   allow-service dhcp
exit

router ospfv3 1
router-id 10.0.0.1
address-family ipv4 unicast
log-adjacency-changes
redistribute connected
redistribute connected metric 1 metric-type 1
exit-address-family
!
address-family ipv6 unicast
log-adjacency-changes
redistribute connected
redistribute connected metric-type 1
exit-address-family
```

The following example describes QoS provisioning on PPPoE extension session:

```
policy-map rar_policer
 class class-default
  police 10000 2000 1000 conform-action transmit  exceed-action drop  violate-action drop
policy-map rar_shaper
 class class-default
  shape average percent 1

interface Virtual-Template2
 ip address 192.0.2.255 255.255.255.0
 no peer default ip address
 no keepalive
 service-policy input rar_policer
end
```

Configure RAR in bypass mode

Before you begin the RAR configuration, you must first configure the **subscriber authorization enable** command to bring up the RAR session. Without enbaling authorization, the Point-to-Point protocol does not recognize this as a RAR session and may not tag manet_radio in PPPoE protocol. By default, bypass mode does not appears in the configuration. It appears only if the mode is configured as bypass.

Configure a service for RAR:

```
policy-map type service rar-lab
 pppoe service manet_radio //note: Enter the pppoe service policy name as manet_radio
!
```

Configure Broadband:

```
interface pppoe VMI2
 virtual-template 2
```

```
service profile rar-lab
!
interface GigabitEthernet0/0/0
 description Connected to Client1
  negotiation auto
  pppoe enable group VMI2
!

policy-map type service rar-lab
 pppoe service manet_radio //note: Enter the pppoe service policy name as manet_radio
!
```

Configuration in bypass mode:

This exmple shows IP Address configured under virtual template explicitly:

```
interface Virtual-Template2
ip address 192.0.2.255 255.255.255.0
no ip redirects
peer default ip address pool PPPoEpool2
ipv6 enable
ospfv3 1 network manet
ospfv3 1 ipv4 area 0
ospfv3 1 ipv6 area 0
no keepalive
service-policy input rar_policer Or/And
service-policy output rar_shaper
```

This example shows VMI unnumbered configured under virtual template:

```
interface Virtual-Template2
ip unnumbered vmi2
no ip redirects
peer default ip address pool PPPoEpool2
ipv6 enable
ospfv3 1 network manet
ospfv3 1 ipv4 area 0
ospfv3 1 ipv6 area 0
no keepalive
service-policy input rar_policer Or/And
service-policy output rar_shaper
```

Configure the virtual multipoint interface in bypass mode:

```
interface vmi2 //configure the virtual multi interface
ip address 192.0.2.255 255.255.255.0
physical-interface GigabitEthernet0/0/0
mode bypass
interface vmi3//configure the virtual multi interface
 ip address 192.0.2.255 255.255.255.0
 physical-interface GigabitEthernet0/0/1
mode bypass
```

Configure RAR in aggregate mode

Before you configure RAR, you must first configure the **subscriber authorization enable** command to bring up the RAR session. Without enabling authorization, the Point-to-Point protocol does not recognize this as a RAR session and may not tag manet_radio in PPPoE.

The following example is an end-to-end configuration of RAR in the aggregate mode:

Configure a Service for RAR:

```
policy-map type service rar-lab
 pppoe service manet_radio //note: Enter the pppoe service policy name as manet_radio
!
```

Configure Broadband:

```
bba-group pppoe VMI2
 virtual-template 2
service profile rar-lab

!
interface GigabitEthernet0/0/0
 description Connected to Client1
  negotiation auto
  pppoe enable group VMI2

!
policy-map type service rar-lab
 pppoe service manet_radio //note: Enter the pppoe service policy name as manet_radio
!
```

Configuration in Aggregate Mode:

```
interface Virtual-Template2
ip unnumbered vmi2
no ip redirects
no peer default ip address
ipv6 enable
no keepalive
service-policy input rar_policer Or/And
service-policy output rar_shaper
```

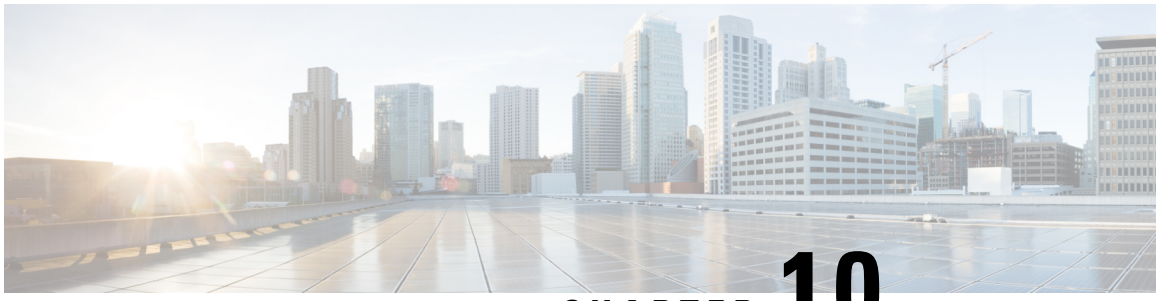

CHAPTER 10

Route Leaking between VPNs

Route leaking is a fundamental mechanism in Cisco Catalyst SD-WAN that facilitates the exchange of routing information between different Virtual Private Networks (VPNs) or Virtual Routing and Forwarding (VRF) instances. This feature enables service sharing, simplifies network migrations, and enhances routing flexibility by allowing routes to be replicated bidirectionally between the global VRF and service VPNs, or directly between service VPNs.

Feature history for route leaking between VPNs

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

Table 35: Feature history

Feature Name	Release Information	Description
Route leaking between global VRF and service VPNs	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	Route leaking enables you to leak routes bidirectionally between the global VRF and service VPNs. The feature allows service sharing and is beneficial in migration use cases because it allows bypassing hubs and provides migrated branches direct access to non-migrated branches.

Feature Name	Release Information	Description
Redistribution of replicated BGP routes to OSPF, EIGRP protocols	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage Release 20.5.1	This feature allows you to leak or replicate BGP routes between the global VRF and service VPNs, then redistribute the leaked BGP routes. The redistribution of the leaked routes to the EIGRP and OSPF protocols occurs after replicating the BGP routes into the corresponding VRF.
Redistribution of replicated routes to BGP, OSPF, and EIGRP Protocols	Cisco IOS XE Catalyst SD-WAN Release 17.6.1a Cisco vManage Release 20.6.1	This feature allows you to configure: • Redistribution of leaked or replicated routes between the global VRF and service VPNs for BGP, OSPF, and EIGRP protocols on Cisco IOS XE Catalyst SD-WANs, • OMP administrative distance option to prefer OMP routes over MPLS routes, • VRRP tracking to track whether a leaked route is reachable.
Route leaking between inter-service VPNs	Cisco IOS XE Catalyst SD-WAN Release 17.9.1a Cisco vManage Release 20.9.1	With this feature, you can leak routes between the service VPNs at the same edge device. Route leaking feature allows redistribution of replicated routes between the inter-service VPN for Connected, Static, BGP, OSPF, and EIGRP on Cisco IOS XE Catalyst SD-WANs.

Supported protocols

To enable users to identify the specific protocols supported for route leaking between global VRF and service VPNs, and for route redistribution between service VPNs and global VRF, including any usage restrictions.

Supported protocols for route leaking

- Connected
- Static

- BGP
- OSPF
- EIGRP

The supported protocols for route redistribution between service VPNs and the global VRF are:

Source Protocols

- Connected
- Static
- BGP
- OSPF
- EIGRP

Destination Protocols

- BGP
- OSPF
- EIGRP

Note EIGRP is supported only on service VPNs, not on the global VRF. As a result, you can leak routes for EIGRP only from service VPNs to the global VRF.

Restrictions for route leaking and redistribution

Observe these restrictions when configuring route leaking and redistribution:

EIGRP

- EIGRP can only be used on service VRFs, not on the global VRF. Therefore, route leaking isn't supported for routes from the global VRF to the service VRFs, and between service VRFs for the EIGRP protocol.
- Redistribution in EIGRP requires bandwidth, load, reliability, delay, and MTU settings to select the best path.

NAT

- Service-side NAT is not supported with route leaking between the global VRF and service VRFs.
- NAT is not supported with transport VRF route leaking.

Unsupported address families and features

- IPv6 address family is not supported for route leaking.
- Inter-service VRF route leaking is not supported on Cisco IOS XE Catalyst SD-WAN devices with multi-tenancy.
- Route leaking using centralized policy is not supported.
- Route leaking across different devices or sites using export policies in Cisco Catalyst SD-WAN is not supported.

Route filtering and capacity

- Each service VRF can leak (import and export) a maximum of 1000 routes.
- Only prefix-lists, tags, and metrics can be matched in route maps that are used to filter leaked routes.

OMP and static routes

- Overlay Management Protocol (OMP) routes do not participate in VRF route leaking to prevent overlay looping.
- Static routes that point to a next-hop resolved through a prefix replicated from a service-side VPN into the global routing table (GRT) are not supported. However, you can configure a static route in a service VPN and replicate it into the GRT.

Redistribution configuration

- Route replicate with **all** keyword is not recommended.
- When configuring route leaking for a VRF, the **route-replicate** command under **global-address-family ipv4** should not use the **all** keyword as the protocol for the unicast option. Instead, specify a particular protocol (e.g., **connected**) to prevent route looping.
- Redistribution of replicated routes into BGP (which were imported into the global routing table from a VRF or into another VRF) is not supported within the same topology. For example, to redistribute a connected route from the BGP global routing table that was originally replicated from VRF 1, use `redistribute connected vrf 1` instead of `redistribute connected`.

Route leaking

Route leaking is a mechanism that enables network segmentation using VPNs and allows sharing of common services that multiple VPNs need to access.

Route Leaking Between Global VRF and Service VPNs

Route leaking between the global or default VRF (transport VPN) and service VPNs allows you to share common services that multiple VPNs need to access. With this feature, routes are replicated through bidirectional route leaking between the global VRF (also known as transport VPN) and service VPNs. Route leaking between VRFs is done using Routing Information Base (RIB).

To leak routes to the routing neighbors, redistribute the leaked routes between the global VRF and service VPNs.

Note In the context of Cisco Catalyst SD-WAN, the terms VRF and VPN are used interchangeably. Although Cisco IOS XE Catalyst SD-WAN devices use VRFs for segmentation and network isolation, the VPN feature template is used to configure them using Cisco SD-WAN Manager. When you use Cisco SD-WAN Manager to configure VPNs for Cisco IOS XE Catalyst SD-WAN devices, Cisco SD-WAN Manager automatically converts the VPN configuration to VRF configuration.

To leak routes to the routing neighbors, redistribute the leaked routes between the global VRF and service VPNs.

OMP administrative distance for leaked routes

You can configure the Cisco SD-WAN Overlay Management Protocol (OMP) administrative distance to a lower value that sets the OMP routes as the preferred and primary route over any leaked routes in a branch-to-branch routing scenario.

Ensure that you configure the OMP administrative distance on Cisco IOS XE Catalyst SD-WAN devices based on the following points:

- If you configure the OMP administrative distance at both the global VRF and service VRF level, the VRF-level configuration overrides the global VRF-level configuration.
- If you configure the service VRF with a lower administrative distance than the global VRF, then except the service VRF, all the remaining VRFs take the value of the administrative distance from the global VRF.

To configure the OMP administrative distance using Cisco SD-WAN Manager, see *Configure Basic VPN Parameters* and *Configure OMP using SD-WAN Manager templates*.

To configure the OMP administrative distance using the CLI, see the Configure OMP Administrative Distance section in Configure OMP Using the CLI.

Inter-Service VPN route leaking

The Inter-Service VPN Route Leaking feature provides the ability to leak selective routes between service VPNs back to the originating device on the same site.

To resolve routing-scalability challenges introduced when you use Cisco Catalyst SD-WAN Control Componentss, you can leak routes between the VPNs at the edge device.

To configure the inter-service VPN route leaking feature using Cisco SD-WAN Manager, see Configure Route Leaking Between Service VPNs.

To configure the inter-service VPNs route leaking feature using the CLI, see Configure Route Leaking Between Service VPNs Using the CLI.

Use VRRP tracker for leaked service VPNs

The Virtual Router Redundancy Protocol (VRRP) can track whether a leaked route is reachable. If tracked route is not reachable, VRRP changes the priority of the VRRP group. It can trigger a new primary router election. The VRRP tracker determines whether a route is reachable based on the existence of the route in the routing table of the routing instance that is included in the VRRP configuration.

To configure the VRRP tracker to track a leaked service VPN using Cisco SD-WAN Manager, see *Configure VRRP for Cisco VPN Interface Ethernet template*.

To configure the VRRP tracker to track any leaked service VPNs using the CLI, see Configure VRRP Tracker for Tracking Leaked Service VPNs Using the CLI.

Features of route leaking

Route leaking offers these features:

- Routes between the global VRF and service VPNs can be leaked directly.
- Multiple service VPNs can be leaked to the global VRF.
- Multiple service VRFs leaking into the same service VRF is supported.
- When routes are leaked or replicated between the global VRF and service VPNs, route properties such as metric, source VPN information, tags, administrative distance, and route origin are retained.
- You can control leaked routes using route maps.
- Route-maps can filter routes using match operations before leaking them.
- The feature can be configured using both—Cisco SD-WAN Manager and CLI.

Use cases for route leaking

Route leaking provides solutions for various network scenarios, offering benefits such as service sharing, simplified migration, and enhanced network management.

Route leaking is applied in these scenarios:

- Service Provider Central Services: This feature allows direct access to SP Central services under MPLS without duplicating them for each VPN. This approach makes accessing central services more efficient.
- Migration: With route leaking, branches that have migrated to Cisco SD-WAN can directly access non-migrated branches bypassing the hub, thus providing improved application SLAs.
- Centralized Network Management: You can manage the control plane and service-side equipment through the underlay.
- Retailer Requirements for PCI compliance: Route leaking for service VPNs is used where the VPN traffic goes through a zone-based firewall on the same branch router while being PCI compliant.

How route preference is determined

When a route is replicated or leaked between the global VRF and service VPNs, a specific set of rules determines which route is preferred. These rules ensure consistent and predictable routing behavior within the network.

Route preference is determined by these rules:

1. For a device that receives routes from two sources that both use the same source VRFs, and one of the routes is replicated, the non-replicated route is preferred.
2. If the first rule does not apply, these rules determine route preference in this order:

a. Prefer the route with smaller administrative distance.

b. Prefer the route with smaller default administrative distance.

c. Prefer a non-replicated route over a replicated route.

d. Compare original VRF names. Prefer the route with the lexicographically smaller VRF name.

e. Compare original subaddress families. Prefer unicast routing over multicast routing.

f. Prefer the oldest route.

Configure route leaking

Use this task to enable route leaking within your Cisco Catalyst SD-WAN. Route leaking allows for service sharing between different VPNs and facilitates network migration by providing direct access between migrated and non-migrated branches.

Before you begin

Before you begin, ensure you understand the concepts of route leaking and have reviewed the associated restrictions. For more information, see Route leaking, on page 142 and Restrictions for route leaking and redistribution, on page 141.

Follow these steps to configure route leaking:

Procedure

Step 1 Configure and enable a localized policy, then attach the route policy.

a) Configure localized route policy. See Configure localized route policy, on page 146.

b) Add the route policy. See Add the route policy, on page 147

c) Attach the localized policy to the device template. See Attach the localized policy to the device template, on page 147.

Step 2 Configure and enable the route leaking feature between global and service VPNs.

a) Configure and enable route leaking between global and service VPNs using a configuration group. See Configure and enable route leaking between global and service VPNs using a configuration group, on page 148.

b) Configure and enable route leaking between global and service VPNs using Cisco SD-WAN Manager. See Configure and enable route leaking between global and service VPNs using Cisco SD-WAN Manager, on page 149.

c) Configure and verify route leaking between global VRF and service VPNs using the CLI. See Configure and verify route leaking between global VRF and service VPNs using the CLI, on page 151.

Step 3 Configure and enable the route leaking feature between service VPNs.

a) Configure route leaking between service VPNs using a configuration group. See Configure route leaking between service VPNs using a configuration group, on page 154.

b) Configure route leaking between service VPNs using Cisco SD-WAN Manager. See Configure route leaking between service VPNs using Cisco SD-WAN Manager, on page 155.

c) Configure route leaking between service VPNs using a CLI template. See Configure route leaking between service VPNs using a CLI template, on page 156.

d) Verify route-leaking configurations between service VPNs using the CLI. See Verify route-leaking configurations between service VPNs using the CLI, on page 161.
e) Configure VRRP tracker for tracking leaked service VPNs using the CLI. See Configure VRRP tracker for tracking leaked service VPNs using the CLI, on page 157.
f) Verify VRRP tracking. See Verify VRRP tracking, on page 162.

Step 4 Attach the service side VPN feature template to the device template. See Attach the service side VPN feature template to the device template.

You have successfully configured route leaking in your Cisco Catalyst SD-WAN, enabling service sharing and optimizing routing paths.

What to do next

To view a configuration example for route leaking, see Configuration example for route leaking, on page 159.

Configure localized route policy

Use this task to create a new localized route policy that defines how routes are handled within your Cisco Catalyst SD-WAN.

.

Before you begin

Follow these steps to configure localized route policy:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies**.
Step 2 Select **Localized Policy**.
Step 3 From the **Custom Options** drop-down, under Localized Policy, select **Route Policy**.
Step 4 Click **Add Route Policy**, and select **Create New**.
Step 5 Enter a name and description for the route policy.
Step 6 In the left pane, click **Add Sequence Type**.
Step 7 In the right pane, click **Add Sequence Rule** to create a single sequence in the policy. Match is selected by default.
Step 8 Select a desired protocol from the **Protocol** drop-down list. The options are: IPv4, IPv6, or both.
Step 9 Click a match condition.
Step 10 On the left, enter the values for the match condition.
Step 11 On the right enter the action or actions to take if the policy matches.
Step 12 Click **Save Match and Actions** to save a sequence rule.
Step 13 If no packets match any of the route policy sequence rules, the default action is to drop the packets. To change the default action:
a) Click **Default Action** in the left pane.
b) Click the **Pencil** icon.
c) Change the default action to **Accept**.

d) Click **Save Match and Actions**.

Step 14 Click **Save Route Policy**.

You have successfully created a localized route policy.

What to do next

Add the route policy. See Add the route policy, on page 147.

Add the route policy

Use this task to import an existing route policy into your localized policy configuration.

Before you begin

Follow these steps to add the route policy:

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Policies**.
Step 2 Choose the **Localized Policy**.
Step 3 Click **Add Policy**.
Step 4 Click **Next** in the Local Policy Wizard until you arrive at the **Configure Route Policy** option.
Step 5 Click **Add Route Policy** and choose **Import Existing**.
Step 6 From the **Policy** drop-down list, choose the route policy that you created previously.
Step 7 Click **Import**.
Step 8 Click **Next**.
Step 9 Enter the **Policy Name** and **Description** for this localized policy.
Step 10 Click **Preview** to view the policy configurations in CLI format.
Step 11 Click **Save Policy**.

You have successfully added the route policy to your localized policy configuration.

What to do next

Attach the localized policy to the device template to apply it to your network devices. For more information, see Attach the localized policy to the device template, on page 147.

Attach the localized policy to the device template

Use this task to apply a previously configured localized policy to a device template.

Before you begin

Follow these steps to attach the localized policy to the device template:

Procedure

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

Step 2 Click **Device Templates** and select the desired template.

Step 3 Click **…**, and click **Edit**.

Step 4 Click **Additional Templates**.

Step 5 From the **Policy** drop-down list, choose the **Localized Policy** that you created.

Step 6 Click **Update**.

Note
Once the localized policy has been added to the device template, selecting the **Update** option immediately pushes a configuration change to all of the devices that are attached to this device template. If more than one device is attached to the device template, you will receive a warning that they are changing multiple devices.

Step 7 Click **Next** and then **Configure Devices**.

Step 8 Wait for the validation process and push configuration from Cisco SD-WAN Manager to the device.

You have successfully attached the localized policy to the device template, and the configuration has been pushed to the associated devices.

What to do next

Now that the localized policy is applied, you can proceed to configure and enable the route leaking feature between global and service VPNs. For more information, see step 2 in Configure route leaking, on page 145.

Configure and enable route leaking between global and service VPNs using a configuration group

Use this task to configure route leaking between the global VRF and service VPNs using a configuration group in Cisco SD-WAN Manager. This method allows for centralized management and deployment of route leaking policies across multiple devices.

Before you begin

Follow these steps to configure and enable route leaking between global and service VPNs using a configuration group:

Procedure

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

Step 2 Create and configure a Service VPN feature from a Service profile.

a) To leak routes from the global VRF:

1. In the **Route Protocol*** field, choose a protocol to configure leak routes from the global VPN to the service VPN that you are configuring. Options include **static**, **connected**, **bgp**, or **ospf**.

2. In the **Select Route Policy** field, choose a route policy from the drop-down list.
3. Under **Redistribution (in service VPN)**, in the **Protocol*** field, choose a protocol from the available options to redistribute the leaked routes. Options include **bgp**, **ospf**(minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.18.1) , or **eigrp**.
4. In the **Select Route Policy** field, choose a route policy from the drop-down list.

b) To leak routes from the service VPNs to the global VRF:

1. In the **Route Protocol*** field, choose a protocol to leak routes from the service VPN that you are configuring to the global VPN. Options include **static**, **connected**, **bgp**, **ospf** or **eigrp**.
2. In the **Select Route Policy** field, choose a route policy from the drop-down list.
3. Under **Redistribution (in global VPN)**, in the **Protocol*** field, choose a protocol from the available options to redistribute the leaked routes. Options include **bgp**, or **ospf**.
4. In the **Select Route Policy** field, choose a route policy from the drop-down list.\

Step 3 Click **Save**.

You have successfully configured route leaking between the global VRF and service VPNs using a configuration group.

What to do next

Deploy the configuration group to apply these settings to your devices. See Deploy a configuration group.

Configure and enable route leaking between global and service VPNs using Cisco SD-WAN Manager

Use this task to configure and enable route leaking between the global VRF and service VPNs using feature templates in Cisco SD-WAN Manager. This method allows you to define route leaking policies and apply them to specific devices.

Before you begin

Note that route leaking can only be configured on service VPNs. The VPN numbers in the Basic Configuration must be within the range 1—511 or 513—65527. VPN 512 is reserved for network management traffic, and VPN 0 is reserved for control traffic.

Procedure

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

Step 2 To configure route leaking, click **Feature Templates**.

Step 3 Do one of the following:

a) To create a feature template:

1. Click **Add Template**.

2. Choose a device from the list of devices. The templates available for the selected device display in the right pane.
3. Choose the **Cisco VPN** template from the right pane.
4. Enter Template Name and Description for the feature template.
5. Click **Global Route Leak** below the **Description** field.
6. To leak routes from the global VRF, click **Add New Route Leak from Global VPN to Service VPN**.
   a. From the **Route Protocol Leak from Global to Service** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
   b. In the **Route Policy Leak from Global to Service** drop-down list, choose **Global**, then select one of the available route policies.
   c. For the **Redistribute to protocol (in Service VPN)** field, click **Add Protocol**.
   d. In the **Protocol** drop-down list, choose **Global** to select a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
   e. In the **Redistribution Policy** drop-down list, choose **Global**, then select one of the available redistribution policies.
   f. Click **Add**.
7. To leak routes from the service VPNs to the global VRF, click **Add New Route Leak from Service VPN to Global VPN**.
   a. In the **Route Protocol Leak from Service to Global** drop-down list, choose **Global** to select a protocol. Otherwise, or choose **Device-Specific** to use a device-specific value.
   b. In the **Route Policy Leak from Service to Global** drop-down list, choose **Global**, then select one of the available route policies.
   c. For the **Redistribute to protocol (in Global VPN)** field, click **Add Protocol**.
   d. From the **Protocol** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
   e. In the **Redistribution Policy** drop-down list, choose **Global**, then select one of the available redistribution policies.
   f. Click **Add**.
8. Click **Save/Update**.
9. To redistribute the leaked routes using Cisco SD-WAN Manager, use the *CLI Add-on Feature templates* to enter the configuration applicable to your environment. Here's an example.

```
Device(config)# router ospf 65535
Device(config-router)# redistribute vrf 1 ospf 103


Device(config)# router eigrp vpn
Device(config-router)# address-family ipv4 vrf 1 autonomous-system 50
Device(config-router-af)# topology base
Device(config-router-af-topology)# redistribute vrf global  ospf 65535
metric  1 2 3 4 5
```

After you create the CLI add-on template, you need to attach it to the protocol template to which you are redistributing routes. In this example, you would attach it to the EIGRP template.

b) To modify an existing feature template:

1. Choose a feature template you wish to modify.
2. Click **...** next to the row in the table, and click **Edit**.
3. Click **Global Route Leak**.
4. To edit information, from the table under **Add New Route Leak from Global VPN to Service VPN** or **Add New Route Leak from Service VPN to Global VPN**, click **Edit**.
5. The update route leak dialog box appears. Perform the necessary modifications.
6. Click **Save Changes**.
7. Click **Update**.

You have configured route leaking between the global VRF and service VPNs within the feature template.

What to do next

Attach the service side VPN feature template to the device template to apply these configurations to your network devices. For more information, see Attach the service side VPN feature template to the device template.

Configure and verify route leaking between global VRF and service VPNs using the CLI

Use this task to directly configure and verify route leaking between VRFs using the command-line interface (CLI).\

Before you begin

Follow these steps to configure and verify route leaking using the CLI:

Procedure

Step 1 Configure and verify route leaking between a global VRF and a service VPN.

- Configure route leaking.

  These examples show how to configure route leaking between a global VRF and a service VPN. In this example, VRF 103 is the service VPN. This example shows that connected routes are leaked into VRF 103 from the global VRF, similarly, the same connected routes are leaked from VRF 103 to the global VRF.

```
vrf definition 103
 !
  address-family ipv4
   route-replicate from vrf global unicast connected
!
global-address-family ipv4
```

```
route-replicate from vrf 103 unicast connected
exit-address-family
```

- Verify route leaking

a. Verify routes leaked from service VRF 103 to the global VRF. Leaked routes are represented by a + sign next to the route. For example, C+ denotes a leaked connected route.

```
Device#show ip route
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

10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
O 10.1.14.0/24 [110/11] via 10.1.15.13, 00:02:22, GigabitEthernet1
C 10.1.15.0/24 is directly connected, GigabitEthernet1
L 10.1.15.15/32 is directly connected, GigabitEthernet1
O 10.1.16.0/24 [110/11] via 10.1.15.13, 00:02:22, GigabitEthernet1
C 10.1.17.0/24 is directly connected, GigabitEthernet2
L 10.1.17.15/32 is directly connected, GigabitEthernet2
172.16.0.0/12 is subnetted, 1 subnets
[170/10880] via 192.168.24.17(103), 01:04:13, GigabitEthernet5.103
192.168.0.0/16 is variably subnetted, 2 subnets, 2 masks
C + 192.0.2.0/24  is directly connected, GigabitEthernet5.103
L & 192.168.24.15/16 is directly connected, GigabitEthernet5.103
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 203.0.113.0/24 is directly connected, GigabitEthernet6
L 203.0.113.15/32 is directly connected, GigabitEthernet6
10.20.0.0/8 is variably subnetted, 2 subnets, 2 masks
C 198.51.100.0/24 is directly connected, GigabitEthernet7
L 198.51.100.15/24 is directly connected, GigabitEthernet7
192.0.2.0/32 is subnetted, 1 subnets
O E2 100.100.100.100 [110/20] via 10.1.15.13, 00:02:22, GigabitEthernet1
172.16.0.0/32 is subnetted, 1 subnets
O E2 172.16.255.14 [110/20] via 10.1.15.13, 00:02:22, GigabitEthernet1
```

b. View Routes Leaked From Global VRF to Service VRF Table

Use the **show ip route vrf** <*vrf id*> command to view the routes leaked from the global VRF to the service VRF table.

```
Device#show ip route vrf 103
Routing Table: 103
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

10.0.0.0/8 is variably subnetted, 14 subnets, 2 masks
C + 10.0.1.0/24 is directly connected, GigabitEthernet9
L & 10.0.1.15/32 is directly connected, GigabitEthernet9
C + 10.0.20.0/24 is directly connected, GigabitEthernet4
L & 10.0.20.15/32 is directly connected, GigabitEthernet4
C + 10.0.100.0/24 is directly connected, GigabitEthernet8
L & 10.0.100.15/32 is directly connected, GigabitEthernet8
C + 10.1.15.0/24 is directly connected, GigabitEthernet1
L & 10.1.15.15/32 is directly connected, GigabitEthernet1
C + 10.1.17.0/24 is directly connected, GigabitEthernet2
L & 10.1.17.15/32 is directly connected, GigabitEthernet2
172.16.0.0/12 is subnetted, 1 subnets
D EX 172.16.20.20
[170/10880] via 192.168.24.17, 01:04:07, GigabitEthernet5.103
192.168.0.0/16 is variably subnetted, 2 subnets, 2 masks
C 192.0.2.0/24 is directly connected, GigabitEthernet5.103
L 192.168.24.15/16 is directly connected, GigabitEthernet5.103
10.0.0.0/8 is variably subnetted, 2 subnets, 2 masks
C + 203.0.113.0/24 is directly connected, GigabitEthernet6
L & 203.0.113.15/32 is directly connected, GigabitEthernet6
10.20.0.0/8 is variably subnetted, 2 subnets, 2 masks
C + 198.51.100.0/24 is directly connected, GigabitEthernet7
L & 198.51.100.15/24 is directly connected, GigabitEthernet7
192.0.2.0/32 is subnetted, 1 subnets
```

Step 2 Configure and verify route filtering before leaking.

To further filter the routes leaked between the global VRF and the service VRF, you can apply a route map as shown in this example.

```
vrf definition 103
 !
  address-family ipv4
   route-replicate from vrf global unicast connected route-map myRouteMap permit 10
    match ip address prefix-list pList seq 5 permit 10.1.17.0/24
!
```

Step 3 Verify route filtering.

```
Device#show ip route vrf 103

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

10.0.0.0/8 is variably subnetted, 8 subnets, 2 masks
C + 10.1.17.0/24 is directly connected, GigabitEthernet2
```

```
L & 10.1.17.15/32 is directly connected, GigabitEthernet2
m 10.1.18.0/24 [251/0] via 172.16.255.14, 19:01:28, Sdwan-system-intf
m 10.2.2.0/24 [251/0] via 172.16.255.11, 17:28:44, Sdwan-system-intf
m 10.2.3.0/24 [251/0] via 172.16.255.11, 17:26:50, Sdwan-system-intf
C 10.20.24.0/24 is directly connected, GigabitEthernet5
L 10.20.24.15/32 is directly connected, GigabitEthernet5
m 10.20.25.0/24 [251/0] via 172.16.255.11, 16:14:18, Sdwan-system-intf
172.16.0.0/32 is subnetted, 3 subnets
m 172.16.255.112 [251/0] via 172.16.255.11, 17:28:44, Sdwan-system-intf
O E2 172.16.255.117 [110/20] via 10.20.24.17, 1d11h, GigabitEthernet5
m 172.16.255.118 [251/0] via 172.16.255.11, 16:14:18, Sdwan-system-intf
```

Step 4 Monitor leaked routes.

```
Device#show ip cef 10.1.17.0 internal
10.1.17.0/24, epoch 2, flags [rcv], refcnt 6, per-destination sharing
[connected cover 10.1.17.0/24 replicated from 1]
sources: I/F
feature space:
Broker: linked, distributed at 4th priority
subblocks:
gsb Connected receive chain(0): 0x7F6B4315DB80
Interface source: GigabitEthernet5 flags: none flags3: none
Dependent covered prefix type cover need deagg, cover 10.20.24.0/24
ifnums: (none)
path list 7F6B47831168, 9 locks, per-destination, flags 0x41 [shble, hwcn]
path 7F6B3D9E7B70, share 1/1, type receive, for IPv4
receive for GigabitEthernet5
output chain:
receive
```

You have configured and verified route leaking and route filtering using the CLI.

Configure route leaking between service VPNs using a configuration group

Use this task to configure route leaking between different service VPNs using a configuration group in Cisco SD-WAN Manager.

Before you begin

Follow these steps to configure route leaking between service VPNs using a configuration group:

Procedure

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

Step 2 Create and configure a Service VPN feature from a Service profile.

a) To leak routes between services:

1. In the **Source VPN** field, enter the value of the source VPN from which routes will be leaked.
2. In the **Route Protocol*** field, choose a protocol from the available options to leak routes from the source service VPN to the service VPN that you are configuring. Options include **static**, **connected**, **bgp**, **ospf** (minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.18.1), or **eigrp**.
3. In the **Select Route Policy** field, choose a route policy from the drop-down list.

4. Under **Redistribution (in Service VPN)**, in the **Protocol*** field, choose a protocol from the available options to redistribute the leaked routes. Options include **bgp**, **ospf** minimum supported release: Cisco Catalyst SD-WAN Manager Release 20.18.1, or **eigrp**.
5. In the **Select Route Policy** field, choose a route policy from the drop-down list.

Step 3 Click **Save**.

You have configured route leaking between service VPNs using a configuration group.

What to do next

Deploy the configuration group to apply these settings to your devices. See Deploy a configuration group.

Configure route leaking between service VPNs using Cisco SD-WAN Manager

Use this task to configure route leaking directly between service VPNs using feature templates in Cisco SD-WAN Manager.

Before you begin

Follow these steps to configure route leaking between service VPNs:

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Navigate to the **Cisco VPN** template for the device.

Note
To create a **VPN** template, see *Create VPN Template*

Step 4 Click **Route Leak**.

Step 5 Click **Route Leak between Service VPN**.

Step 6 Click **Add New Inter Service VPN Route Leak**.

Step 7 From the **Source VPN** drop-down list, choose **Global** to configure the service VPN from where you want to leak the routes, or choose **Device-Specific** to use a device-specific value.

You can configure service VPNs within the range VPNs 1 to 511, and 513 to 65530, for service-side data traffic on Cisco IOS XE Catalyst SD-WAN devices. (VPN 512 is reserved for network management traffic. VPN 0 is reserved for control traffic using the configured WAN transport interfaces.)

Step 8 From the **Route Protocol Leak to Current VPN** drop-down list, choose **Global** to select a route protocol to enable route leaking to the current VPN. Otherwise, choose **Device-Specific** to use a device-specific value.

You can choose **Connected**, **Static**, **OSPF**, **BGP**, and **EIGRP** protocols for route leaking.

Step 9 From the **Route Policy Leak to Current VPN** drop-down list, choose **Global** to select a route policy to enable route leaking to the current VPN. Otherwise, choose **Device-Specific** to use a device-specific value.

Note
This field is disabled if no route policies are available.

Step 10 To configure **Redistribute to protocol (in Service VPN)**, click **Add Protocol**.

a. From the **Protocol** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.

You can choose **Connected**, **Static**, **OSPF**, **BGP**, and **EIGRP** protocols for redistribution.

b. (Optional) From the **Redistribution Policy** drop-down list, choose **Global**. Next, choose one of the available redistribution policies from the drop-down list.

Note
This field is disabled if no route policies are available.

Step 11 Click **Add**.

Step 12 Click **Save**.

You have successfully configured route leaking between service VPNs within the feature template.

What to do next

Attach the service side VPN feature template to the device template to apply these configurations to your network devices. For more information, see Attach the service side VPN feature template to the device template.

Configure route leaking between service VPNs using a CLI template

Use this task to configure route leaking between different service VPNs on the same device using a CLI template.

This topic provides sample CLI configurations to configure interservice VPN route leaking on Cisco IOS XE Catalyst SD-WAN devices.

Before you begin

Ensure your Cisco IOS XE Catalyst SD-WAN device is running Cisco IOS XE Catalyst SD-WAN Release 17.9.1a or later.

Follow these steps to configure inter-service VRF route leaking using a CLI template:

Procedure

Step 1 Replicate routes between interservice VRFs on the same device..

```
vrf definition vrf-number
address-family ipv4
route-replicate from vrf source-vrf-name unicast protocol [route-map map-tag]
```

Step 2 Redistribute the routes that are replicated between the service VPNs:

You can configure the subnets only for bgp, nhrp, ospf, ospfv3, and static protocol types.

```
router ospf process-id vrf vrf-number
redistribute vrf vrf-name protocol subnets[route-map map-tag]
```

The following is a complete configuration example for interservice VRF route replication and redistribution:

```
vrf definition 2
 rd 1:2
 !
 address-family ipv4
  route-replicate from vrf 1 unicast static route-map VRF1_TO_VRF2
 exit-address-family
 !
!
ip prefix-list VRF1_TO_VRF2 seq 5 permit 10.10.10.97/32
!
route-map VRF1_TO_VRF2 permit 1
 match ip address prefix-list VRF1_TO_VRF2
!
router ospf 2 vrf 2
 redistribute vrf 1 static route-map VRF1_TO_VRF2
```

You have successfully configured inter-service VRF route leaking using a CLI template.

What to do next

Verify the route leaking configurations. See Verify route-leaking configurations between service VPNs using the CLI, on page 161.

Configure VRRP tracker for tracking leaked service VPNs using the CLI

Use this task to configure a Virtual Router Redundancy Protocol (VRRP) tracker for monitoring the reachability of leaked routes within service VPNs using the CLI.

Before you begin

Follow these steps to configure a VRRP tracker for tracking leaked service VPNs:

Procedure

Step 1 Configure a track.

a) Enter global configuration mode, and track the state of an IP route and enter tracking configuration mode.

```
Device# config-transaction
Device(config)# track object-number {ip} route address|prefix-length  { reachability  | metric
threshold}
```

b) Configure a VPN routing and forwarding (VRF) table.

```
Device(config-track)# ip vrf vrf-name
```

c) Return to privileged EXEC mode.

```
Device(config-track)# end
```

Step 2 Configure VRRP version 2 (VRRPv2).

a) Configure an interface type such as, Gigabit Ethernet.

```
Device(config)# interface type number [name-tag]
```

b) Associate a VRF instance with the Gigabit Ethernet interface.

```
Device(config-if)# vrf forwarding vrf-name
```

c) Set a primary IP address for the Gigabit Ethernet interface.

```
Device(config-if)# ip address ip-address [mask]
```

d) Enable the autonegotiation protocol to configure the speed, duplex mode, and flow control on a Gigabit Ethernet interface.

```
Device(config-if)# negotiation auto
```

e) Create a VRRP group and enter VRRP configuration mode.

```
Device(config-if)# vrrp group address-family ipv4
```

f) Enable the support of VRRP version 2 simultaneously with VRRP version 3.

```
Device(config-if-vrrp)# vrrpv2
```

g) Configure interface list tracking as a single entity.

```
Device(config-if-vrrp)# track track-list-name [decrement priority]
```

h) Configure the preemption delay so that a device with higher priority waits for a minimum period before taking over.

```
Device(config-if-vrrp)# preempt delay minimum seconds
```

i) Specify a primary IP address for VRRP.

```
Device(config-if-vrrp)# address ip-address primary
```

Step 3 Configure a VRF.

a) Configure a VRF routing table instance and enter the VRF configuration mode.

```
Device(config)# vrf definition vrf-number
```

b) Set an address family IPv4 in vrf configuration mode.

```
Device(config-vrf)# address-family ipv4
```

c) Exit from address-family configuration mode.

```
Device(config-ipv4)# exit-address-family
```

The following is a sample configuration for configuring the VRRP tracking.

Use the following configuration to add a track to a VRF red.

```
config-transaction
track 1 ip route 10.1.15.13 255.255.255.0 reachability
ip vrf red
```

Use the following configuration to configure interface tracking and decrement the device priority.

```
interface GigabitEthernet 1.101
vrf forwarding 100
 ip address 10.1.15.13 255.255.255.0
 negotiation auto
 vrrp 2 address-family ipv4
  vrrpv2
  priority 220
```

```
  track 1 decrement 25
  preempt delay minimum 30
  address 10.1.15.100 primary
exit
```

Use the following configuration to configure the VRF routing table instance for the configured VRF.

```
vrf definition 100
!
 address-family ipv4
 exit-address-family
```

You have successfully configured a VRRP track object to monitor the reachability of a leaked route within a service VPN, and integrated it with VRRP on a specified interface.

What to do next

Verify the VRRP tracking configuration. See Verify VRRP tracking, on page 162.

Configuration example for route leaking

Route leaking is typically used in scenarios requiring the use of shared services. Configuring route replication allows mutual redistribution between VRFs or VPNs. Route replication allows shared services because routes are replicated or leaked between the global VRF and service VPNs, enabling clients in one VPN to reach matching prefixes in another VPN.

Sample Topology

In this section, we'll use an example topology to show route-leaking configuration. Here, Edge routers 1 and 2 are located in two different sites in the overlay network and are connected to each other through MPLS. Both the edge routers have route leaking configured to be able to access services in the underlay network. Router 1 sits behind Edge Router 1 in the service side. The local network at this site runs OSPF. Router 2 sites behind the Edge Router 2 on network that has EIGRP in VRF 1. Router 3 also sites behind Edge Router 2 and has OSPF running in VRF 200.

Figure 7: Route Leaking

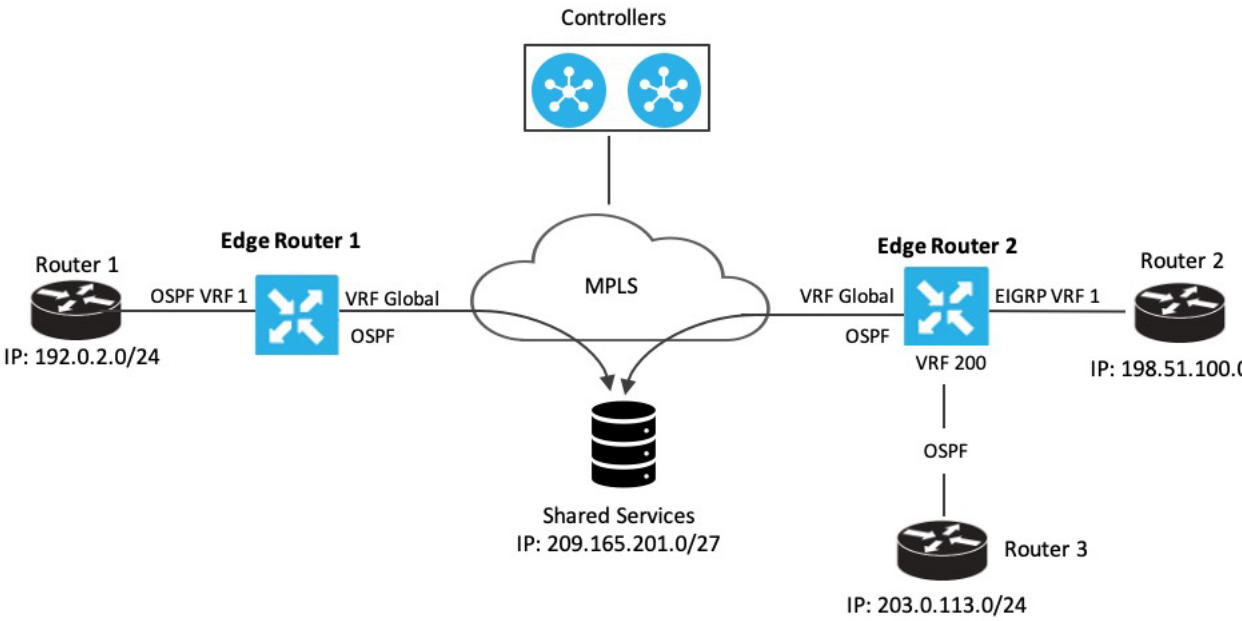

Edge Router 1 imports the source IP address of Router 1, 192.0.2.0/24 to the global VRF on Edge Router 1. Thus 192.0.2.0/24 is a route leaked into the global VRF. Edge Router 2 imports the source IP address of Router 2, 198.51.100.0/24 and the source IP address of Edge Router 3, 203.0.113.0/24 to the global VRF on Edge Router 2.

Shared services in the underlay MPLS network are accessed through a loopback address of 209.21.25.18/27. The IP address of the shared services is advertised to the global VRF on Edge Routers 1 and 2 through OSPF. This shared service IP address is then leaked to VRF 1 in Edge Router 1 and VRF 1 and VRF 200 in Edge Router 2. In terms of route-leaking, the leaked routes are imported into the service VRFs on both the edge routers.

Note OMP doesn't advertise any leaked routes from service VPNs into the overlay network to prevent route looping.

Configuration Examples

This example shows the configuration of BGP and OSPF route leaking between the global VRF and VPN 1 on Edge Router 2.

```
vrf definition 1
 rd 1:1
 !
  address-family ipv4
   route-replicate from vrf global unicast ospf 65535
 !
global-address-family ipv4
 route-replicate from vrf 1 unicast eigrp
 exit-address-family
```

This example shows the configuration of BGP and OSPF route leaking between the global VRF and VPN 200 on Edge Router 2.

```
vrf definition 200
 rd 1:200
 !
  address-family ipv4
   route-replicate from vrf global unicast ospf 65535
 !
global-address-family ipv4
 route-replicate from vrf 200 unicast eigrp
 exit-address-family
```

Verify route-leaking configurations between service VPNs using the CLI

Use this task to verify that routes are being leaked and redistributed correctly between service VPNs on your Cisco IOS XE Catalyst SD-WAN device using the CLI.

Before you begin

Ensure your Cisco IOS XE Catalyst SD-WAN device is running Cisco IOS XE Catalyst SD-WAN Release 17.9.1a or later.

Follow these steps to verify route-leaking configurations between service VPNs:

Procedure

Step 1 View routes replicated for redistribution to a service VRF.

```
Device# show ip route vrf 2
Routing Table: 2
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

      10.0.0.0/8 is variably subnetted, 3 subnets, 2 masks
S   +    10.10.10.97/32 [1/0] via 10.20.1.2 (1)
C        10.20.2.0/24 is directly connected, GigabitEthernet5
L        10.20.2.1/32 is directly connected, GigabitEthernet5
```

Step 2 View replicated routes in the Cisco Express Forwarding (CEF) table for specific replicated routes.

```
Device# show ip cef vrf 2 10.10.10.97 internal
10.10.10.97/32, epoch 0, RIB[S], refcnt 6, per-destination sharing
  sources: RIB
  feature space:
    IPRM: 0x00048000
    Broker: linked, distributed at 3rd priority
```

```
subblocks:
  Replicated from VRF 1
ifnums:
  GigabitEthernet3(9): 10.20.1.2
path list 7F890C8E2F20, 7 locks, per-destination, flags 0x69 [shble, rif, rcrsv, hwcn]
  path 7F890FB18F08, share 1/1, type recursive, for IPv4
    recursive via 10.20.1.2[IPv4:1], fib 7F890B609578, 1 terminal fib, v4:1:10.20.1.2/32
    path list 7F890C8E3148, 2 locks, per-destination, flags 0x49 [shble, rif, hwcn]
        path 7F890FB19178, share 1/1, type adjacency prefix, for IPv4
         attached to GigabitEthernet3, IP adj out of GigabitEthernet3, addr 10.20.1.2 7F890FAE4CD8

output chain:
  IP adj out of GigabitEthernet3, addr 10.20.1.2 7F890FAE4CD8
```

You have successfully verified that routes are being leaked and redistributed between service VPNs as configured. The presence of `+` in the routing table and `Replicated from VRF` in the CEF output confirms the successful operation.

Verify VRRP tracking

Verifying VRRP tracking provides insight into the operational status of VRRP groups and their associated track objects. This information is crucial for confirming that VRRP is actively monitoring the reachability of leaked routes.

VRRP group status details

The following is a sample output from the **show vrrp details** command that shows the status of the configured VRRP groups on a Cisco IOS XE Catalyst SD-WAN device.

```
Device# show vrrp details
GigabitEthernet1/0/1 - Group 1 - Address-Family IPv4
  State is BACKUP         <------ check states
  State duration 2 mins 13.778 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 0000.5E00.0101
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 1 (Configured 100)    <----- shows current and configured priority
    Track object 121 state DOWN decrement 220  Master Router is 10.1.1.3, priority is 200
<---- track object state
  Master Advertisement interval is 1000 msec (learned)
  Master Down interval is 3609 msec (expires in 2737 msec)
  FLAGS: 0/1
  VRRPv3 Advertisements: sent 27 (errors 0) - rcvd 149
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    VRRPv2 incompatibility: 0
    IP Address Owner conflicts: 0
    Invalid address count: 0
    IP address configuration mismatch : 0
    Invalid Advert Interval: 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to master: 0
    Init to backup: 1 (Last change Wed Feb 17 23:02:04.259)
    Backup to master: 1 (Last change Wed Feb 17 23:02:07.869)  <----- check this for flaps
    Master to backup: 1 (Last change Wed Feb 17 23:02:32.008)
```

```
    Master to init: 0
    Backup to init: 0
```

Track object status

The following is a sample output from the **show track** command that displays information about objects that are tracked by the VRRP tracking process.

```
Device# show vrrp details
GigabitEthernet1/0/1 - Group 1 - Address-Family IPv4
  State is BACKUP         <------ check states
  State duration 2 mins 13.778 secs
  Virtual IP address is 10.0.0.1
  Virtual MAC address is 0000.5E00.0101
  Advertisement interval is 1000 msec
  Preemption enabled
  Priority is 1 (Configured 100)    <----- shows current and configured priority
    Track object 121 state DOWN decrement 220  Master Router is 10.1.1.3, priority is 200
<---- track object state
  Master Advertisement interval is 1000 msec (learned)
  Master Down interval is 3609 msec (expires in 2737 msec)
  FLAGS: 0/1
  VRRPv3 Advertisements: sent 27 (errors 0) - rcvd 149
  VRRPv2 Advertisements: sent 0 (errors 0) - rcvd 0
  Group Discarded Packets: 0
    VRRPv2 incompatibility: 0
    IP Address Owner conflicts: 0
    Invalid address count: 0
    IP address configuration mismatch : 0
    Invalid Advert Interval: 0
    Adverts received in Init state: 0
    Invalid group other reason: 0
  Group State transition:
    Init to master: 0
    Init to backup: 1 (Last change Wed Feb 17 23:02:04.259)
    Backup to master: 1 (Last change Wed Feb 17 23:02:07.869)  <----- check this for flaps
    Master to backup: 1 (Last change Wed Feb 17 23:02:32.008)
    Master to init: 0
    Backup to init: 0
```

VRRP interface configuration

The following is a sample output from the show running-config interface command that shows the configuration of a Gigabit Ethernet interface that is tracked by the VRRP tracking process.

```
Device# show running-config interface GigabitEthernet 4
Building configuration...
Current configuration : 234 bytes
!
interface GigabitEthernet4
ip address 172.16.0.1 255.255.255.0
negotiation auto
vrrp 7 address-family ipv4
  priority 200
  vrrpv2
  track 5 decrement 5ß-------priority decreament
  address 172.16.0.0 primary
  exit-vrrp
no mop enabled
no mop sysid
end
```

Route redistribution

Route redistribution is a mechanism that allows routing information learned from one routing protocol to be shared with another routing protocol. This process is essential for maintaining connectivity and enabling communication across different routing domains or segments within a network. It facilitates the exchange of routing information between VRFs (global to service, or service to service) and ensures that routes from disparate routing environments are known throughout the network.

Configure route redistribution between global VRF and service VPNs using the CLI

Use this task to configure route redistribution between the global VRF and service VPNs using the CLI.

Follow these steps to configure route redistribution between global VRF and service VPNs:

Procedure

Step 1 Enter the global configuration mode, and create a BGP routing process.

You can use the router eigrp, or router ospf to configure a routing process for a specific routing protocol. This example shows the syntax for BGP routing protocol. To know about the command syntax for various protocols, see the Cisco IOS XE SD-WAN Qualified Command Reference Guide.

```
Device# config-transaction
Device(config)# router bgp autonomous-system-number
```

Step 2 Configure an IPv4 address family for service VPNs. This example shows the command syntax for the BGP and EIGRP protocols.

- BGP protocol:

  ```
  Device(config-router-af)# address-family ipv4  [unicast][vrf vrf-name]
  ```

- EIGRP protocol:

  ```
  Device(config-router-af)# address-family ipv4 vrf vrf-number
  ```

Step 3 Redistribute routes between the global VRF and service VPNs. Here, we're showing the syntaxes for the BGP, OSPF, and EIGRP protocols. .

- Redistribute routes from service VPNs to the global VRF.

  a. BGP protocol:

  ```
  Device(config-router-af)# redistribute vrf vrf-name src_protocol  [src_protocol_id] [route-map
   route-map-name]
  ```

  b. OSPF protocol:

  ```
  Device(config-router-af)# redistribute vrf vrf-name src_protocol  [src_protocol_id]  [match
  {internal|external 1|external 2}]  [metric {metric-value}]  [subnets] [route-map route-map-name]
  ```

c. EIGRP protocol:

```
Device(config-router-af)# redistribute vrf vrf-name src_protocol [src_protocol_id] [metric
bandwidth-metric delay-metric reliability-metric effective-bandwidth-metric mtu-bytes]
[route-map route-map-name]
```

- Redistribute routes from the global VRF to service VPNs.

  a. BGP protocol:

  ```
  Device(config-router-af)# redistribute vrf global src_protocol  [src_protocol_id] [route-map
   route-map-name]
  ```

  b. OSPF protocol:

  ```
  Device(config-router-af)# redistribute vrf global src_protocol  [src_protocol_id]  [match
  {internal|external 1|external 2}]  [subnets] [route-map route-map-name]
  ```

  c. EIGRP protocol:

  ```
  Device(config-router-af)# redistribute vrf global src_protocol [src_protocol_id] [metric
  bandwidth-metric delay-metric reliability-metric effective-bandwidth-metric mtu-bytes]
  ```

The following is a sample configuration for configuring route redistribution between a global VRF and service VPN. In this example, VRF 103 and VRF 104 are the service VPNs. The example shows that BGP routes are redistributed from the global VRF to VRF 103, VRF 104.

```
config-transaction
router bgp 100

address-family ipv4 vrf 103
redistribute vrf global bgp 100 route-map test2
!
address-family ipv4 vrf 104
redistribute vrf global bgp 100 route-map test2
!
```

The following is a sample configuration for configuring the OSPF internal and external routes that are redistributed from the global VRF 65535 to the service VRF.

In this case, all OSPF routes are redistributed into the service VRF by using both the **internal** and **external** keywords.

```
config-transaction
router ospf 1
redistribute vrf global ospf 65535 match internal external 1 external 2 subnets route-map ospf-route-map
```

The following is a sample configuration for configuring the OSPF internal and external routes that are redistributed from service VPNs to the global VRF.

```
config-transaction
router ospf 101
redistribute vrf 101 ospf 101 match internal external 1 external 2 metric 1 subnets route-map
ospf-route-map
```

The following is a sample configuration for configuring the BGP routes that are redistribution from a service VPN to the global VRF.

```
config-transaction
router bgp 50000
address-family ipv4 unicast
redistribute vrf 102 bgp 50000 route-map BGP-route-map
```

The following is a sample configuration for configuring the BGP routes that are redistribution from the global VRF to a service VPN.

```
config-transaction
router bgp 50000
address-family ipv4 vrf 102
redistribute vrf global bgp 50000
```

The following is a sample configuration for configuring route redistribution of BGP, connected, OSPF, and static protocols from the global VRF to VRF 1 when configuring under EIGRP routing process.

```
config-transaction
router eigrp 101
address-family ipv4 vrf 1
redistribute vrf global bgp 50000 metric 1000000 10 255 1 1500
redistribute vrf global connected metric 1000000 10 255 1 1500
redistribute vrf global ospf 65535 match internal external 1 external 2 metric 1000000 10 255 1 1500
redistribute vrf global static metric 1000000 10 255 1 1500
```

You have successfully configured route redistribution between the global VRF and service VPNs using the CLI.

What to do next

Verify the route redistribution.

Verify route redistribution

Use this task to verify that routes are being successfully redistributed between VRFs.

The following is a sample output from the show ip bgp command using the internal keyword. This example shows that a route from VRF 102 is redistributed successfully to the global VRF after the route is replicated.

```
Device# show ip bgp 10.10.10.10 internal

BGP routing table entry for 10.10.10.10/8, version 515
Paths: (1 available, best #1, table default)
Not advertised to any peer
Refresh Epoch 1
700000 70707
10.10.14.17 from 0.0.0.0 (172.16.255.15)
Origin IGP, aigp-metric 77775522, metric 7777, localpref 100, weight 32768, valid, sourced,
 replicated, best
Community: 0:7227 65535:65535
Extended Community: SoO:721:75 RT:50000:102
rx pathid: 0, tx pathid: 0x0
net: 0x7FB320235DC0, path: 0x7FB320245DF8, pathext: 0x7FB3203A4660
flags: net: 0x0, path: 0x808040003, pathext: 0x81
attribute: 0x7FB38E5B6258, ref: 14
Updated on Jul 1 2021 01:16:36 UTC
```

In this output, the route is redistributed from VRF 102 to the global VRF.

The following is a sample output from the show ip route command that shows the routes replicated for redistribution.

```
Device# show ip route 10.10.10.10

Routing entry for 10.10.10.10/8
Known via "bgp 50000", distance 60, metric 7777
Tag 700000, type external,
replicated from topology(102)
```

```
Redistributing via ospf 65535, bgp 50000
Advertised by ospf 65535
bgp 50000 (self originated)
Last update from 10.10.14.17 5d15h ago
Routing Descriptor Blocks:
* 10.10.14.17 (102), from 10.10.14.17, 5d15h ago
opaque_ptr 0x7FB3202563A8
Route metric is 7777, traffic share count is 1
AS Hops 2
Route tag 700000
MPLS label: none
```

The following is a sample output from the **show ip bgp vpnv4 vrf** command using the **internal** keyword.

```
Device# show ip bgp vpnv4 vrf 102 209.165.201.0 internal

BGP routing table entry for 1:102:10.10.10.10/8, version 679
BGP routing table entry for 1:209.165.201.0/27, version 679
Paths: (1 available, best #1, table 102)
Advertised to update-groups:
4
Refresh Epoch 1
7111 300000
10.1.15.13 (via default) from 0.0.0.0 (172.16.255.15)
Origin IGP, aigp-metric 5755, metric 900, localpref 300, weight 32768, valid, sourced,
replicated, best
Community: 555:666
Large Community: 1:2:3 5:6:7 412789:412780:755
Extended Community: SoO:533:53 RT:50000:102
rx pathid: 0, tx pathid: 0x0
net: 0x7FB38E5C5718, path: 0x7FB3202668D8, pathext: 0x7FB38E69E960
flags: net: 0x0, path: 0x808040007, pathext: 0x181
attribute: 0x7FB320256798, ref: 7
Updated on Jul 6 2021 16:43:04 UTC
```

In this output, the route is redistributed from the global VRF to VRF 102.

The following is a sample output from the show ip route vrf command that shows the routes replicated for redistribution for VRF 102.

```
Device# show ip route vrf 102 209.165.201.0

Routing Table: 102
Routing entry for 209.165.201.0/27
Known via "bgp 50000", distance 20, metric 900
Tag 7111, type external,
replicated from topology(default)
Redistributing via bgp 50000
Advertised by bgp 50000 (self originated)
Last update from 10.1.15.13 00:04:57 ago
Routing Descriptor Blocks:
* 10.1.15.13 (default), from 10.1.15.13, 00:04:57 ago
opaque_ptr 0x7FB38E5B5E98
Route metric is 900, traffic share count is 1
AS Hops 2
Route tag 7111
MPLS label: none
```

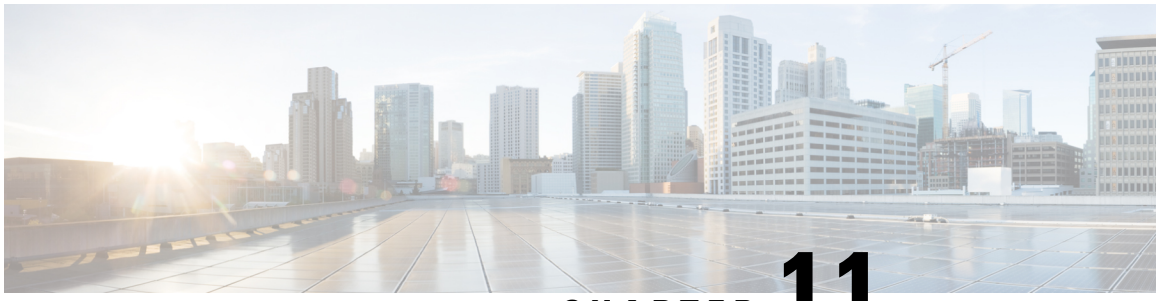

CHAPTER 11

Bi-directional Forwarding Detection Protocol

Feature history for BFD protocol

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

Table 36: Feature history

Feature name	Release information	Description
BFD Troubleshooting for Cisco Catalyst SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 17.15.1aCisco Catalyst SD-WAN Manager Release 20.15.1	This feature enables you to troubleshoot BFD protocols using radioactive tracing, check device logs, and use debugging commands to gather detailed information about BFD operations.
Automatically Suspend Unstable Cisco Catalyst SD-WAN BFD Sessions	Cisco IOS XE Catalyst SD-WAN Release 17.10.1aCisco vManage Release 20.10.1	With this feature you can automatically suspend unstable Cisco Catalyst SD-WAN BFD sessions based on flap-cycle thresholds or SLA parameters. It also allows you to monitor all suspended BFD sessions and manually reset them when needed.
BFD for Routing Protocols in Cisco Catalyst SD-WAN	Cisco IOS XE Catalyst SD-WAN Release 17.3.1a Cisco vManage Release 20.3.1	This feature extends BFD support to BGP, OSPF, and EIGRP in the Cisco Catalyst SD-WAN solution, allowing BFD to provide a consistent, uniform failure-detection method that quickly identifies forwarding-path failures and enables faster reconvergence.

BFD protocol for Cisco SD-WAN

A Bi-directional Forwarding Detection (BFD) is a network protocol that

- detects failures rapidly between forwarding engines,
- operates with low overhead, and
- enables faster reconvergence of business-critical applications.

BFD provides a single, standardized method to detect link, device, or protocol failures across any layer and media.

BFD in enterprise networks

In enterprise networks, organizations increasingly run business-critical applications on a shared IP infrastructure. They design these networks with high redundancy to protect data and ensure reliability. However, redundancy works effectively only when network devices detect failures and switch to alternate paths quickly.

Traditional protocols often take more than a second to identify failures, which delays recovery for time-sensitive applications. BFD solves this problem by detecting failures rapidly and triggering faster recovery, allowing networks to maintain consistent performance and uptime.

How BFD works in Cisco Catalyst SD-WAN

With this feature, the Cisco Catalyst SD-WAN solution includes two independent BFD types that operate separately without conflict.

BFD Support for Cisco Catalyst SD-WAN Routing Protocols (Legacy BFD): This legacy BFD feature already exists in Cisco IOS XE and extends to the Cisco Catalyst SD-WAN solution starting from Cisco IOS XE Catalyst SD-WAN Release 17.3.1a.

Cisco Catalyst SD-WAN BFD: This feature specifically supports overlay BFD, which already exists in the Cisco Catalyst SD-WAN solution.

This type of BFD detects failures in the overlay tunnel and has these characteristics:

- It operates by default and cannot be disabled.
- It typically runs for the Cisco Catalyst SD-WAN Overlay Management Protocol (OMP).
- In addition to detecting link failures, Cisco Catalyst SD-WAN BFD measures latency, loss, jitter, and other link statistics that application-aware routing uses.

Table 37: Differences: BFD for Cisco Catalyst SD-WAN routing protocols versus Cisco Catalyst SD-WAN BFD

BFD for Cisco Catalyst SD-WAN routing protocols	Cisco Catalyst SD-WAN BFD

• Runs on both, transport-side and service-side interfaces
• The following protocols can be registered: BGP, OSPF, and EIGRP
  • BGP (transport and service side
  • EIGRP (service side)
  • OSPF and OSPFv3 (service side)
• Detects link failures for peers in terms of whether a peer is up or down
	• Runs on a Cisco Catalyst SD-WAN tunnel to detect failures in the overlay tunnel
	• Is enabled by default and cannot be disabled
	• Is typically enabled for OMP
	• Besides link failures, it also measures latency, loss, and other link statistics used by application-aware routing

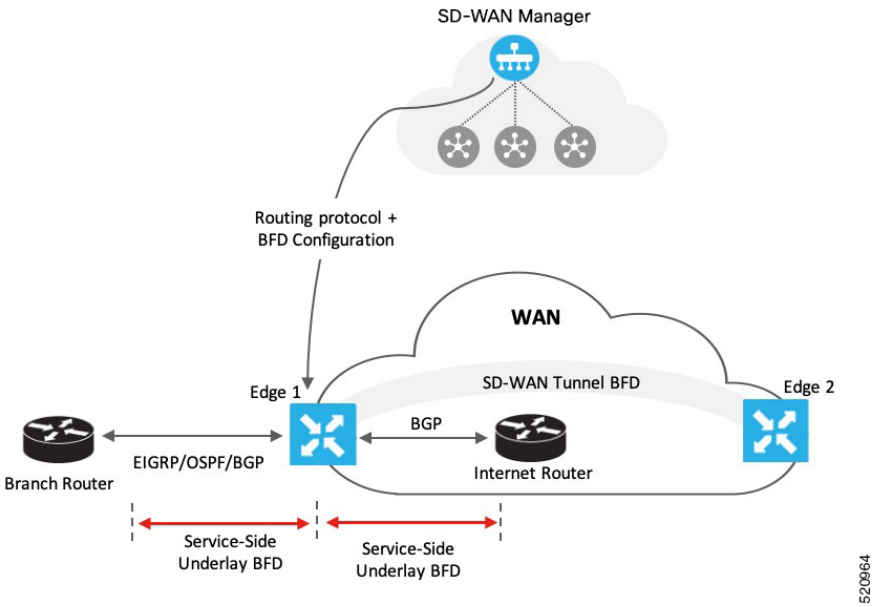

As shown in the image, you configure BFD for a routing protocol through Cisco SD-WAN Manager. Cisco SD-WAN Manager then pushes this configuration to the edge router. In this example, OSPF receives forwarding path detection failure messages from BFD. When a physical link fails, BFD notifies OSPF, prompting it to shut down its neighbors and withdraw or restore any routing information exchanged with remote neighbors.

Similarly, Edge 1 connects to the internet router through its transport interface. You configure BFD for BGP between the transport side of Edge 1 and the internet router. In this setup, BFD monitors the connection's health and reports any detected failures.

Supported protocols and interfaces

Supported protocols

In Cisco Catalyst SD-WAN, the following routing protocols can receive forwarding-path failure notifications from BFD:

- BGP
- EIGRP
- OSPF
- OSPFv3

Supported interfaces

BFD supports the following interface types:

- GigabitEthernet
- TenGigabitEthernet
- FiveGigabitEthernet
- FortyGigabitEthernet
- HundredGigabitEthernet
- SVIs
- Subinterfaces

Restrictions for Cisco IOS XE Catalyst SD-WAN devices in controller mode

- The device supports only single-hop BFD.
- The device does not support BFD for static routes.
- To change the BFD session mode between software and hardware, you must remove all existing BFD configurations and reconfigure them.
- The device supports BFD only for BGP, EIGRP, OSPF, and OSPFv3.
- Cisco SD-WAN Manager cannot monitor BFD for routing protocols in Cisco Catalyst SD-WAN; you must use CLI show commands for monitoring.
- After a BFD session is established, the device does not update BFD session modes (echo-no-echo or software- hardware) immediately when you change BFD template parameters in Cisco SD-WAN Manager; it applies the change only after the session flaps at least once.

Static route support for BFD sessions in transport VPN

For BFD to work correctly, the next-hop IP address must be known and reachable. Static routes configured incorrectly can cause BFD session failures.

Static route configuration guidelines for BFD sessions in transport VPN

- Do not configure static routes pointing directly to Ethernet (non-P2P) interfaces without specifying a next-hop IP. This causes BFD failures because the next-hop IP is unknown.
- For Ethernet interfaces using DHCP, ensure the DHCP server provides the router IP. Then use either the DHCP-learned route or configure a static route with the dhcp keyword to dynamically resolve the next-hop IP. For more information on DHCP server and client configurations, see Configuring the Cisco IOS XE DHCP Server and Configuring the Cisco IOS XE DHCP Client.
- For point-to-point (P2P) interfaces, such as cellular, you can configure static routes that point directly to the interface.
- In P2P networks, the remote BFD endpoint is always a single hop away with no intermediate devices.
- In non-P2P networks, the remote BFD endpoint is typically multiple hops away, with the next hop being only the first of several intermediate hops.

Static route support matrix for BFD sessions

Table 38: Static route support matrix for BFD sessions

Exit interface type	Static route format	BFD support	Description
Ethernet (Non-P2P)	ip route 0.0.0.0 0.0.0.0 GigabitEthernetX	Not supported	Next-hop IP unknown; causes BFD failure. Use next-hop IP or DHCP-based static route.
Ethernet (Non-P2P)	ip route 0.0.0.0 0.0.0.0 <next-hop IP>	Supported	Recommended when next-hop IP is static and known.
Ethernet (Non-P2P)	ip route 0.0.0.0 0.0.0.0 dhcp or ip route 0.0.0.0 0.0.0.0 GigabitEthernetX dhcp	Supported	For DHCP interfaces; dynamically resolves next-hop IP from DHCP server.
Point-to-Point (P2P)	ip route 0.0.0.0 0.0.0.0 CellularX	Supported	Fully supported static route format for BFD.

Verify static routes for BFD sessions

To verify the installation and next-hop resolution of static default routes regardless of whether the route uses a next-hop IP address, DHCP-resolved next hop, or a directly connected interface use the **show ip route** command.

Examples

Ethernet (Non-P2P) with multiple next-hop IPs

For these configurations below:

```
ip route 0.0.0.0 0.0.0.0 10.0.0.1
ip route 0.0.0.0 0.0.0.0 172.16.0.1
```

Verify the routing table for transport VPN:

```
Device# show ip route
S* 0.0.0.0/0 [1/0] via 10.0.0.1
             [1/0] via 172.16.0.1
Device#
```

Ethernet (Non-P2P) with DHCP-resolved Next-Hop

For this configuration below:

```
ip route 0.0.0.0 0.0.0.0 dhcp
```

Verify the routing table for transport VPN:

```
Device# show ip route
S* 0.0.0.0/0 [1/0] via 192.168.0.1
Device#
```

Ethernet (Non-P2P) with DHCP on specific interface

For this configuration below:

```
ip route 0.0.0.0 0.0.0.0 GigabitEthernet0/0/1 dhcp
```

Verify the routing table for transport VPN:

```
Device# show ip route
S* 0.0.0.0/0 [1/0] via 192.168.0.1, GigabitEthernet0/0/1
Device#
```

Point-to-point (P2P) directly connected interface

For this configuration below:

```
ip route 0.0.0.0 0.0.0.0 Cellular0/0/0
```

Verify the routing table for transport VPN:

```
Device# show ip route
S* 0.0.0.0/0 is directly connected, Cellular0/0/0
Device#
```

Combination of P2P and DHCP-based static routes

For these configurations below:

```
ip route 0.0.0.0 0.0.0.0 Cellular0/0/0
ip route 0.0.0.0 0.0.0.0 dhcp
```

Verify the routing table for transport VPN:

```
Device# show ip route
S* 0.0.0.0/0 is directly connected, Cellular0/0/0
             [1/0] via 192.168.0.1
Device#
```

Configure BFD 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 BFD feature in a System profile.

a) Configure basic settings.

Table 39: Basic Configuration

Field	Description
Poll Interval(In Millisecond)	Specify how often BFD polls all data plane tunnels on a router to collect packet latency, loss, and other statistics used by application-aware routing. Range: 1 through 4,294,967,296 (2^{32} – 1) milliseconds Default: 600,000 milliseconds (10 minutes)
Multiplier	Specify the value by which to multiply the poll interval, to set how often application-aware routing acts on the data plane tunnel statistics to figure out the loss and latency and to calculate new tunnels if the loss and latency times do not meet the configured SLAs. Range: 1 through 6 Default: 6
DSCP Values for BFD Packets(decimal)	Specify the Differentiated Services Code Point (DSCP) value of the BFD packets that is used in the DSCP control traffic. Range: 0-63 Default: 48

b) Configure colors.

Table 40: Color

Field	Description
Add Color	

Field	Description
Color*	Choose the color of the transport tunnel for data traffic moving between the devices. The color identifies a specific WAN transport provider. Values: 3g, biz-internet, blue, bronze, custom1, custom2, custom3, default, gold, green, lte, metro-ethernet, mpls, private1 through private6, public-internet, red, silver Default: default
Hello Interval (milliseconds)*	Specify how often BFD sends Hello packets on the transport tunnel. BFD uses these packets to detect the liveness of the tunnel connection and to detect faults on the tunnel. Range: 100 through 300000 milliseconds Default: 1000 milliseconds (1 second)
Multiplier*	Specify how many Hello packet intervals BFD waits before declaring that a tunnel has failed. BFD declares that the tunnel has failed when, during all these intervals, BFD has received no Hello packets on the tunnel. This interval is a multiplier of the Hello packet interval time. Range: 1 through 60 Default: 7
Path MTU Discovery*	Enable or disable path MTU discovery for the transport tunnel. When path MTU discovery is enabled, the path MTU for the tunnel connection is checked periodically, about once per minute, and it is updated dynamically. When path MTU discovery is disabled, the expected tunnel MTU is 1472 bytes, but the effective tunnel MTU is 1468 bytes. Default: Enabled
Default DSCP value for BFD packets*	Specify the Differentiated Services Code Point (DSCP) value of the BFD packets that is used in the DSCP control traffic. Range: 0-63 Default: 48

The **show sdwan bfd session** output displays BFD parameters based on the negotiation to choose greater value of hello interval and multiplier combined. The calculation of multiplier is changed from the show output per the negotiation result.

What to do next

Also see Deploy a configuration group.

Configure BFD for routing protocols

Use one of these methods to configure BFD for routing protocols:

- Templates
- CLI commands

Configure BFD for routing protocols using templates

Cisco SD-WAN Manager does not provide an independent template to configure BFD for routing protocols. However, you can register or deregister supported protocols to receive BFD packets by adding configurations through the CLI add-on template in Cisco SD-WAN Manager. Use the CLI add-on template to:

- Add a single-hop BFD template and specify parameters such as timer, multiplier, and session mode.
- Enable the BFD template under interfaces. You can add only one BFD template per interface.
- Enable or disable BFD for supported routing protocols. The configuration steps differ for each protocol-BGP, EIGRP, OSPF, and OSPFv3.

Starting with release Cisco IOS XE Catalyst SD-WAN Release 17.15.1a, if SD-WAN mode is not configured for the tunnel interface, the BFDs become inactive for the tunnel interface.

Complete the tasks below to configure BFD for routing protocols using CLI commands.

- Enable BFD for routing protocols
- Attach feature template to device template

Enable BFD for routing protocols

- Configure BFD for Service-Side BGP
- Configure BFD for Transport-Side BGP
- Configure BFD for Service-Side EIGRP
- Configure BFD for Service-Side OSPF and OSPFv3

Configure BFD for service-side BGP

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
Step 2 Click **Feature Templates**.
Step 3 Click **Add Template**.
Step 4 Choose a device from the device list.

Step 5 Choose the **CLI Add-on Template** under **Other Templates**.

Step 6 Enter the CLI configuration to create a single-hop BFD template and enable BFD for service-BGP, as shown in the example below.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
 !
interface GigabitEthernet1
   bfd template t1

router bgp 10005
address-family ipv4 vrf 1
    neighbor 10.20.24.17 fall-over bfd
    !
  address-family ipv6 vrf 1
    neighbor 2001::7 fall-over bfd
```

In this example, you create a single-hop BFD template by specifying the minimum interval, maximum interval, and multiplier. These parameters are mandatory. You can also optionally configure other BFD parameters, such as echo mode (enabled by default) and BFD dampening (disabled by default). After creating the template, you enable it under an interface (GigabitEthernet1 in this example).

To modify a BFD template already enabled on an interface, you must first remove the existing template, update it, and then enable it on the interface again.

If you attach the BFD configuration to a device template that already includes a BGP feature template, ensure that you update the BGP configuration in the CLI add-on template to include the **no neighbor ip-address ebgp-multihop** command. This update is required because the **neighbor ip-address ebgp-multihop** command is enabled by default in the BGP feature template.

Step 7 Click **Save**.

Step 8 Attach the CLI Add-on Template with this configuration to the device template.

For the configuration to take effect, the device template must have a BGP feature template attached to it.

Step 9 Attach the device template to the device.

Configure BFD for transport-side BGP

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Click **Add Template**.

Step 4 Choose a device from the device list.

Step 5 Choose the **CLI Add-on Template** under **Other Templates**.

Step 6 Enter the CLI configuration to create a single-hop BFD template and enable BFD for transport-BGP, as shown in the example below.

```
bfd-template single-hop t1
```

```
interval min-tx 500 min-rx 500 multiplier 3
!
interface GigabitEthernet1
bfd template t1
!
router bgp 10005
neighbor 10.1.15.13 fall-over bfd
!
sdwan
interface GigabitEthernet1
tunnel-interface
allow-service bfd
allow-service bgp
```

In this example, you create a single-hop BFD template by specifying the minimum interval, maximum interval, and multiplier. These parameters are mandatory. You can also configure optional BFD parameters, such as echo mode (enabled by default) and BFD dampening (disabled by default). After creating the template, you enable it under an interface (GigabitEthernet1 in this example), because GigabitEthernet1 is also the SD-WAN tunnel source, allowing service under its tunnel interface ensures that both BGP and BFD packets pass over the tunnel.

To modify a BFD template already enabled on an interface, you must remove the existing template, update it, and then enable it on the interface again.

If you attach the BFD configuration to a device template that already includes a BGP feature template, ensure that you update the BGP configuration in the CLI add-on template to include the **no neighbor ip-address ebgp-multihop** command. This update is required because the **neighbor ip-address ebgp-multihop** command is enabled by default in the BGP feature template.

Step 7 Click **Save**.

Step 8 Attach the CLI Add-on Template with this configuration to the device template.

For the configuration to take effect, the device template must have a BGP feature template attached to it.

Step 9 Attach the device template to the device.

Configure BFD for service-side EIGRP

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Click **Add Template**.

Step 4 Choose a device from the device list.

Step 5 Choose the **CLI Add-on Template** under **Other Templates**.

Step 6 Enter the CLI configuration to add a single-hop BFD template and enable BFD for EIGRP as shown in the example below.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
 !
interface GigabitEthernet5
```

```
   bfd template t1

router eigrp myeigrp
address-family ipv4 vrf 1 autonomous-system 1
    af-interface GigabitEthernet5
     bfd
```

In this example, you create a single-hop BFD template by specifying the minimum interval, maximum interval, and multiplier. These parameters are mandatory. You can also configure optional BFD parameters, such as echo mode (enabled by default) and BFD dampening (disabled by default).

After creating the template, you enable it under an interface (GigabitEthernet5 in this example).

To modify a BFD template already enabled on an interface, you must first remove the existing template, update it, and then enable it on the interface again.

Step 7 Click **Save**.

Step 8 Attach the CLI Add-on Template with this configuration to the device template.

For the configuration to take effect, the device template must have a BGP feature template attached to it.

Step 9 Attach the device template to the device.

Configure BFD for service-side OSPF and OSPFv3

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Click **Add Template**.

Step 4 Choose a device from the device list.

Step 5 Choose the **CLI Add-on Template** under **Other Templates**.

Step 6 Enter the CLI configuration to add a single-hop BFD template enable BFD for OSPF and OSPFv3 as shown in the examples below.

OSPF

```
bfd-template single-hop t1
interval min-tx 500 min-rx 500 multiplier 3
!
interface GigabitEthernet5
bfd template t1
!
interface GigabitEthernet1
bfd template t1
!
router ospf 1 vrf 1
 bfd all-interfaces
!
```

OSPFv3

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3

interface GigabitEthernet5
   bfd template t1
router ospfv3 1
 address-family ipv4 vrf 1
  bfd all-interfaces
```

In these examples, you create a single-hop BFD template by specifying the minimum interval, maximum interval, and multiplier. These parameters are mandatory. You can also configure optional BFD parameters, such as echo mode (enabled by default) and BFD dampening (disabled by default).

After creating the template, you enable it under an interface (GigabitEthernet5 in this example).

To modify a BFD template already enabled on an interface, you must first remove the existing template, update it, and then enable it on the interface again.

Step 7 Click **Save**.

Step 8 Attach the CLI Add-on Template with this configuration to the device template.

For the configuration to take effect, the device template must have a BGP feature template attached to it.

Step 9 Attach the device template to the device.

Attach feature template to device template

Use these steps to attach the configuration to a device template.

After you create a CLI add-on template to enable BFD, attach the template to the device template for the configuration to take effect.

In Cisco SD-WAN Release 20.7.x and earlier releases, **Device Templates** is titled **Device**.

Before you begin

Make sure the device template already includes the relevant feature template (BGP, OSPF, or EIGRP) before you attach the CLI add-on template.

Procedure

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

Step 2 Click **Device Templates**.

Step 3 Click **Create Template** and choose **From Feature Template** from the drop-down options.

Step 4 From the **Device Model** drop-down options, choose a device.

Enter a name and description for the template.

Step 5 Click **Create**.

Step 6 Click **Additional Templates**.

Step 7 In the **CLI Add-on Template** field, choose the CLI add-on template you configured to enable BFD for routing protocols.

Step 8 Click **Create**.

What to do next

Attach device template to device

Configure BFD for routing protocols using CLI commands

Follow these steps to configure BFD for BGP, EIGRP, OSPF, and OSPF3 using device CLI.

Procedure

Step 1 Create a BFD template.

The CLI configuration for creating a BFD template remains the same irrespective of the protocol you configure it for.

Create a single-hop BFD template as shown in the example below.

```
bfd-template single-hop t1
 interval min-tx 500 min-rx 500 multiplier 3
```

Step 2 Enable BFD for service-side BGP.

This example shows BGP configured, BFD enabled on the interface under VRF 1, and then enabled for service-side BGP.

```
interface GigabitEthernet5
bfd template t1
!
router bgp 10005
  bgp log-neighbor-changes
  distance bgp 20 200 20
  !
  address-family ipv4 vrf 1
  bgp router-id 10.20.24.15
  redistribute connected
  neighbor 10.20.24.17 remote-as 10007
  neighbor 10.20.24.17 activate
  neighbor 10.20.24.17 send-community both
  neighbor 10.20.24.17 maximum-prefix 2147483647 100
  neighbor 10.20.24.17 fall-over bfd
  exit-address-family
  !
  address-family ipv6 vrf 1
  bgp router-id 10.20.24.15
  neighbor 2001::7 remote-as 10007
  neighbor 2001::7 activate
  neighbor 2001::7 send-community both
  neighbor 2001::7 maximum-prefix 2147483647 100
  neighbor 2001::7 fall-over bfd
  exit-address-family
```

Step 3 Enable BFD for transport-side BGP

```
interface GigabitEthernet1
bfd template t1
!
router bgp 10005
bgp router-id   10.1.15.15
bgp log-neighbor-changes
distance bgp 20 200 20
neighbor 10.1.15.13 remote-as 10003
neighbor 10.1.15.13 fall-over bfd
```

```
address-family ipv4 unicast
neighbor 10.1.15.13 remote-as 10003
neighbor 10.1.15.13 activate
neighbor 10.1.15.13 maximum-prefix 2147483647 100
neighbor 10.1.15.13 send-community both
redistribute connected
exit-address-family
!
timers bgp 60 180

sdwan
interface GigabitEthernet1
tunnel-interface
allow-service bgp
allow-service bfd
```

Step 4 Enable BFD for EIGRP.

This example shows EIGRP configured, BFD enabled on the interface under VRF 1, and then enabled for service-side EIGRP.

```
interface GigabitEthernet5
bfd template t1
!
router eigrp myeigrp
address-family ipv4 vrf 1 autonomous-system 1
    af-interface GigabitEthernet5
     no dampening-change
     no dampening-interval
     hello-interval 5
     hold-time      15
     split-horizon
     bfd
     exit-af-interface
    !
    network 10.20.24.0 0.0.0.255
    topology base
     redistribute connected
     redistribute omp
     exit-af-topology
    !
    exit-address-family
   !
```

Step 5 Enable BFD for OSPFv3.

This example shows OSPFv3 configured, BFD enabled on the interface under VRF 1, and then enabled for service-side OSPFv3.

```
interface GigabitEthernet5
   bfd template t1
   ospfv3 1 ipv4 area 0
   ospfv3 1 ipv4 dead-interval 40
   ospfv3 1 ipv4 hello-interval 10
   ospfv3 1 ipv4 network broadcast
   ospfv3 1 ipv4 priority 1
   ospfv3 1 ipv4 retransmit-interval 5
   ospfv3 1 ipv6 area 0
   ospfv3 1 ipv6 dead-interval 40
   ospfv3 1 ipv6 hello-interval 10
   ospfv3 1 ipv6 network broadcast
   ospfv3 1 ipv6 priority 1
   ospfv3 1 ipv6 retransmit-interval 5
```

```
router ospfv3 1
 address-family ipv4 vrf 1
  area 0 normal
  bfd all-interfaces
  router-id 10.20.24.15
  distance 110
  exit-address-family
 !
 address-family ipv6 vrf 1
  area 0 normal
  bfd all-interfaces
  router-id 10.20.24.15
  distance 110
  exit-address-family
 !
!
exit
```

Automatically suspending BFD sessions

A BFD session flap is a network condition that

- occurs when a BFD session repeatedly transitions between up and down states,
- happens because one device in the session becomes unavailable and then available again, or
- repeatedly recovers and fails due to unstable connections, disrupting applications and causing unnecessary traffic steering between overlay paths.

Automatically suspending BFD sessions

To prevent repeated BFD session flaps, Cisco Catalyst SD-WAN automatically suspends unstable BFD sessions based on the following parameters:

Flap cycle

A flap cycle includes this sequence:

- The BFD session is in the up state.
- The BFD session transitions to the down state.
- The BFD session comes back up.

SLA threshold

An SLA threshold determines when to add a BFD session to the suspended list. It defines a limit for traffic metrics such as loss, latency, or jitter. When any metric exceeds the defined threshold, the system suspends the BFD session. These thresholds represent the traffic performance levels specified in the SLA.

An SLA threshold is optional. If you configure one, set higher values for loss, latency, and jitter to prevent conflicts with SLA parameters defined in SLA classes. For more details on SLA classes, see the Cisco Catalyst SD-WAN Policies Configuration Guide.

Benefits of automatically suspending BFD sessions

- You can manually remove the affected circuit or tunnel interface from the BFD suspended list.
- Provides monitoring of a suspended tunnel.

How BFD session suspension works

- As the BFD suspension feature is for forward data traffic, you should enable BFD suspension on the remote-end node to block the reverse data traffic to avoid dropping data traffic.
- This feature does not manipulate the state of the BFD session.

Summary

The BFD session suspension workflow temporarily halts unstable sessions to prevent repeated flapping, allowing only control traffic while blocking data traffic until stability is restored.

Workflow

If a BFD session exceeds the flap-count value within the configured flapping-window interval, then the BFD session must remain suspended until the configured duration interval.

1. For a BFD session in the suspended state, the following occurs:
   - If a session reflaps or exceeds the threshold parameters defined, the session is moved back to suspended state and the duration is reset again.
   - If the session does not flap and is within the threshold range, the session is automatically removed out of the suspended state after the duration interval expires.
   - You can also manually remove suspended BFD sessions by using the **request platform software sdwan auto-suspend reset** command. For more information, see the Cisco IOS XE SD-WAN Qualified Command Reference Guide.

Result

Data traffic is not sent across the overlay network when a BFD session is in the suspended state.

Only regular SLA measurement, echo response, or path maximum transmission unit (PMTU) control traffic is sent across a suspended BFD session.

Restrictions for automatically suspending BFD sessions

Defines limitations for using BFD automatic suspension in Cisco SD-WAN.

- On a Cisco IOS XE Catalyst SD-WAN device with a single TLOC, automatic suspension may drop BFD sessions.
- The last-resort circuit may not work for a single site unless all BFD sessions are down for a tunnel interface. The last-resort circuit is enabled only if all BFD sessions on the non last-resort circuit are suspended or down.
- SD-WAN Manager feature templates do not support configuring automatic suspension of BFD sessions.

- You can configure BFD automatic suspension only through a device CLI or a CLI add-on template.
- When duplicated traffic is sent through a different BFD session, it may still route through a suspended BFD session.

Configure automatic suspension of BFD sessions using a CLI template

To configure automatic suspension of BFD sessions using a CLI template.

If you enable **color all** and a specific **color** , the specific color takes precedence over the color all parameter. For more information on BFD colors, see bfd color.

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates.

Procedure

Step 1 Enable BFD automatic suspension with or without last resort.

Before enabling last resort for the BFD automatic suspension feature, you must enable the last-resort circuit on a tunnel interface.

For more information on last resort, see last-resort-circuit.

```
auto-suspend
    enable-lr
auto-suspend
    no enable-lr
```

Step 2 Configure the flap parameters.

When you use SLA-based BFD automatic suspension, ensure the duration is greater than the BFD multiplier multiplied by the BFD poll interval. We recommend configuring the BFD automatic suspension duration to more than 30 minutes.

```
duration sec
    flapping-window sec
    flap-count flap-count
```

Step 3 (Optional) Configure SLA parameters.

Before configuring SLA thresholds, ensure to configure BFD session flapping parameters and duration.

```
thresholds
   color
    all
     jitter  jitter-value
     latency latency-value
     loss    loss-value
    !
```

Verify automatic suspension of BFD sessions

Use any of these commands to verify automatic suspension of BFD sessions.

To view the total suspend count and check how many times the BFD session has been suspended, use the **show sdwan bfd sessions suspend** command.

The following columns are added for analyzing BFD session suspension metrics: `RE-SUSPEND COUNT`, `SUSPEND TIME LEFT`, `TOTAL COUNT`, and `SUSPEND DURATION`.

```
Device# show sdwan bfd sessions suspend
                          SOURCE TLOC   REMOTE TLOC                    DST PUBLIC       DST
 PUBLIC          RE-SUSPEND  SUSPEND       TOTAL     SUSPEND
SYSTEM IP         STATE    COLOR          COLOR           SOURCE IP      IP
PORT         ENCAP  COUNT        TIME LEFT     COUNT     DURATION
-------------------------------------------------------------------------------------------
172.16.255.14    up       lte           lte             10.1.15.15     10.1.14.14
12426        ipsec   0          0:00:19:52    18        0:00:00:07
```

To check whether a suspended flag has been added to a BFD session and to view other BFD session metrics, use the **show sdwan bfd sessions alt** command.

The following columns are added for BFD suspension:

- BFD-LD (Local Discriminator) is a unique identifier for all BFD sessions. Its value must be non-zero and is used internally by Cisco TAC for troubleshooting.
- The FLAGS column includes the ‘Sus’ flag, which indicates that a BFD session is suspended.”

```
Device# show sdwan bfd sessions alt
*Sus = Suspend
*NA  = Flag Not Set
                                        SOURCE TLOC    REMOTE TLOC                   DST
PUBLIC      DST PUBLIC
SYSTEM IP        SITE ID   STATE       COLOR          COLOR          SOURCE IP     IP
         PORT        ENCAP  BFD-LD    FLAGS    UPTIME
-------------------------------------------------------------------------------------------
172.16.255.14    400       up          3g             lte          10.0.20.15    10.1.14.14
      12426       ipsec  20004     NA       0:19:30:40
172.16.255.14    400       up          lte            lte          10.1.15.15    10.1.14.14
      12426       ipsec  20003     Sus      0:00:02:46
172.16.255.16    600       up          3g             lte          10.0.20.15    10.0.106.1
      12366       ipsec  20002     NA       0:19:30:40
172.16.255.16    600       up          lte            lte          10.1.15.15    10.0.106.1
      12366       ipsec  20001     NA       0:19:20:14
```

To view the BFD sessions where the ‘Sus’ flag is added, use the **show sdwan bfd history** command.

```
Device# show sdwan bfd history
                                                  DST PUBLIC       DST PUBLIC
                    RX      TX
SYSTEM IP        SITE ID     COLOR     STATE       IP               PORT        ENCAP  TIME
                   PKTS    PKTS    DEL   FLAGS
-------------------------------------------------------------------------------------------
172.16.255.16    600         lte       up          10.0.106.1       12366       ipsec
06/03/22 02:51:06      0       0       0     [ ]
172.16.255.16    600         lte       up          10.0.106.1       12366       ipsec
06/03/22 02:52:04      153     154     0     [Sus]
172.16.255.16    600         lte       down        10.0.106.1       12366       ipsec
06/03/22 03:00:50      1085    1085    0     [Sus]
```

To view a summary of BFD sessions, including sessions that are up, down, have flapped, or have been suspended use the **show sdwan bfd summary** command.

The following fields are added for BFD session suspension: `sessions-flap`, `sessions-up-suspended`, and `sessions-down-suspended`.

```
Device# show sdwan bfd summary
sessions-total           4
sessions-up              4
sessions-max             4
sessions-flap            4
poll-interval            60000
sessions-up-suspended    1
sessions-down-suspended  0
```

Monitor and verify BFD configuration

This sections provides a list of commands that you can run to verify your BFD configuration.

Verify the BFD template

Run the **show bfd interface** command to check the BFD template under an interface.

```
Device# show bfd interface
Interface Name: GigabitEthernet5
 Interface Number: 11
 Configured bfd interval using bfd template: 12383_4T1
 Min Tx Interval: 50000, Min Rx Interval: 50000, Multiplier: 3
```

Verify BFD configuration for BGP

Run the **show bfd neighbors client bgp ipv4** command to check the status of the BFD session.

```
Device# show bfd neighbors client bgp ipv4

IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.20.24.17                             1/1          Up        Up        Gi5
```

Verify BFD configuration for EIGRP

Run the **show bfd neighbors client eigrp** command to check the status of the BFD session.

```
Device# show bfd neighbors client eigrp
IPv4 Sessions
NeighAddr                              LD/RD         RH/RS     State     Int
10.20.24.17                             1/1          Up        Up        Gi5
```

Verify BFD configuration for OSPF

Run the **show bfd neighbors client ospf** command to check the status of the BFD session.

```
Device# show bfd neighbors client ospf
IPv4 Sessions
NeighAddr            LD/RD         RH/RS     State     Int
10.20.24.17          1/1          Up        Up        Gi5
```

Troubleshoot common BFD errors

This section explains how to identify and resolve common BFD issues.

Check control connections

If you experience issues with BFD, start by checking the control connection between Cisco SD-WAN Manager and the edge router by running the **show sdwan control connections** command.

```
Device#show sdwan control connections
                                                                                                                         PEER
                                          PEER
CONTROLLER
PEER    PEER PEER            SITE       DOMAIN PEER                                  PRIV
  PEER                                    PUB
GROUP
TYPE    PROT SYSTEM IP       ID         ID     PRIVATE IP                            PORT
  PUBLIC IP                               PORT  LOCAL COLOR     PROXY STATE UPTIME      ID

------------------------------------------------------------------------------------------------
vsmart  dtls 172.16.255.19   100        1      10.0.5.19                            12355
 10.0.5.19                               12355 lte             No    up     0:12:45:44  0

vsmart  dtls 172.16.255.20   200        1      10.0.12.20                           12356
 10.0.12.20                              12356 lte             No    up     0:15:59:45  0

vmanage dtls 172.16.255.22   200        0      10.0.12.22
```

Issues in pushing device template to device

If you identify issues with pushing the device template to the device, collect debug logs on the edge device as shown below.

```
debug netconf all
request platform soft system shell
tail -f /var/log/confd/cia-netconf-trace.log
```

If Cisco SD-WAN Manager has successfully pushed the configuration to the device and the issue still persists, run the **show sdwan running-config** command to view all details related to BFD.

Issues with transport-side BFD

If the transport-side BFD session is down, check the packet filter data under the Cisco Catalyst SD-WAN tunnel interface to ensure that you have allowed the BFD packets to pass through on the transport side. Look for `allow-service bgp` and `allow-service bfd` in the output.

```
Device#show sdwan running-config  | sec sdwan
 tunnel mode sdwan
sdwan
 interface GigabitEthernet1
  tunnel-interface
   encapsulation ipsec
   color lte
   allow-service bgp
   allow-service bfd
   ..............
```

Troubleshoot BFD using radioactive tracing

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.15.1a and Cisco Catalyst SD-WAN Manager Release 20.15.1

BFD troubleshooting focuses on identifying and fixing issues in the BFD protocol, which detects faults between devices. You can use this feature to check device logs and run debugging commands to collect detailed information about BFD activity.

Radioactive tracing helps in selective debugging of a session. Tracing is enabled across the layers for intended BFD session that is identified by tloc-pair or a local discriminator. It enables debug level traces automatically for all the modules while processing a packet that matches the condition.

The following **show** and **debug** commands are used in BFD troubleshooting:

- **debug platform condition start**
- **debug platform condition feature sdwan controlplane bfd**
- **show platform hardware qfp active feature bfd datapath**
- **show logging profile sdwan internal filter**

For more information on these show commands, see the chapter Troubleshooting Commands in the Cisco IOS XE SD-WAN Qualified Command Reference guide.

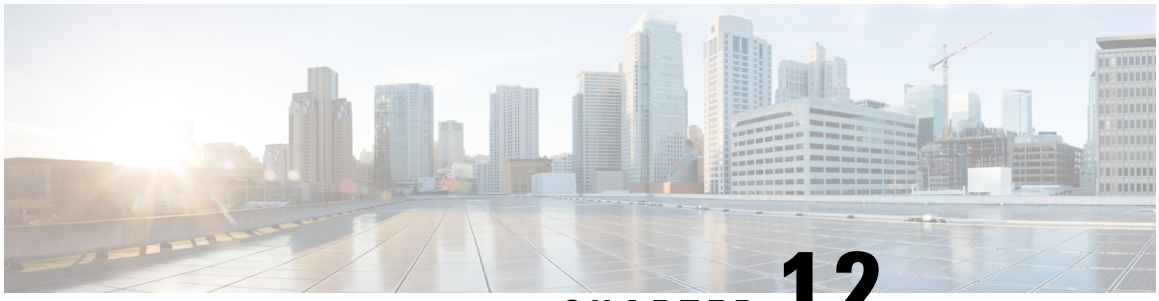

CHAPTER 12

Cisco SD-WAN Controller Route Filtering by TLOC Color

Feature history for route filtering by TLOC color

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

Table 41: Feature History Table

Feature Name	Release Information	Description
Cisco SD-WAN Controller Route Filtering by TLOC Color	Cisco IOS XE Catalyst SD-WAN Release 17.11.1a Cisco Catalyst SD-WAN Control Components Release 20.11.1	Cisco SD-WAN Controller can reduce the number of routes advertised to routers in the network and exclude routes that are not relevant to a particular device. The controller uses the colors of TLOCs on each device to filter and reduce the number of routes. A router that only has private TLOCs does not need routes to public TLOCs. Advertising fewer routes helps to avoid reaching the send path limit for routers in the network.

Route filtering mechanism by TLOC color

A route filtering by TLOC color is a Cisco SD-WAN Controller feature that

- reduces the number of advertised routes to routers by excluding irrelevant routes,
- bases filtering on the colors of TLOCs associated with each device, and
- advertises only routes compatible with one or more of the router's TLOCs.

Route filtering by TLOC color allows Cisco SD-WAN Controllers to selectively advertise routes to individual routers. For each router, only the routes that match the color of one or more of its TLOCs are advertised, ensuring that only relevant routes are sent to each device.

Using route filtering, Cisco SD-WAN Controllers can reduce the number of routes they advertise to routers in the network by excluding routes that are not relevant to a particular device. The filtering is based on the colors of TLOCs on each device: for each individual router, the Cisco SD-WAN Controller advertises only routes that are compatible with one or more of the router's TLOCs.

How Cisco SD-WAN Controller route filtering by TLOC color works

Default behaviour: Cisco SD-WAN Controller route filtering by TLOC color is disabled by default.

Summary

Cisco SD-WAN Controllers apply the following logic when determining whether routes are compatible:

- A TLOC with a public color can resolve a path to a route for a TLOC with a public color on a peer device.
- A TLOC of a particular color can resolve a path to a route for a TLOC of the same color on a peer device.
- A TLOC with a public color cannot resolve a path with a TLOC in a private color set.

Public colors include default, biz-internet, public-internet, lte, 3g, red, green, blue, gold, silver, bronze, custom1, custom2, and so on. Private colors include mpls, metro-ethernet, private1, private2, and so on.

For information about private and public TLOC colors, see Unicast Overlay Routing in the *Cisco SD-WAN Routing Configuration Guide, Cisco IOS XE Release 17.x.*.

For example, if a router only has TLOCs with private colours, Cisco SD-WAN Controllers do not advertise public routes to the device. Similarly, if a router only has TLOCs with public colors, Cisco SD-WAN Controllers do not advertise private routes to the device.

The following illustration provides further detail:

Workflow

Figure 8: Cisco SD-WAN Controller Route Filtering by TLOC Color, With the Feature Enabled

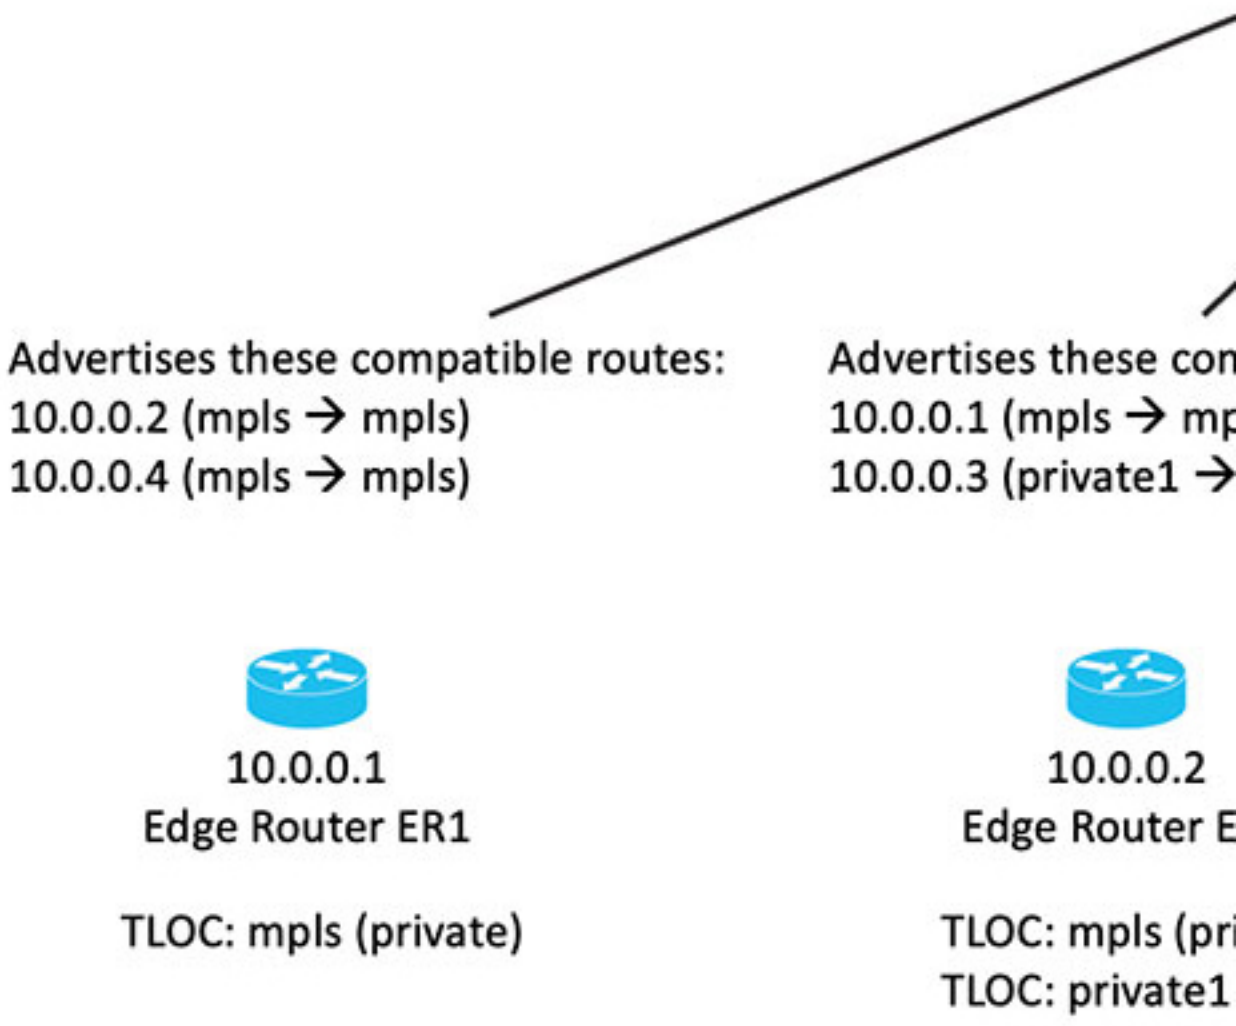

If you change the color assignment of a TLOC, the device updates the Cisco SD-WAN Controllers, enabling them to adjust the Cisco SD-WAN Controller route filtering by TLOC color accordingly.

What's next

You can override the default logic if necessary and do one of the following:

- Configure two TLOC colors to be compatible even if they are incompatible by default.
- Configure two TLOC colors to be incompatible even if they are compatible by default.

This may be helpful in specific unconventional scenarios. See the tloc-color-compatibility command in Override Default TLOC Color Compatibility for Cisco SD-WAN Controller Route Filtering by TLOC Color Using a CLI Template.

The following illustration shows an example of route filtering by TLOC color, with two overrides:

- Configure green and gold to be incompatible.
- Configure mpls and private1 to be compatible.

Figure 9: Cisco SD-WAN Controller Route Filtering by TLOC Color, With the Feature Enabled and Overrides

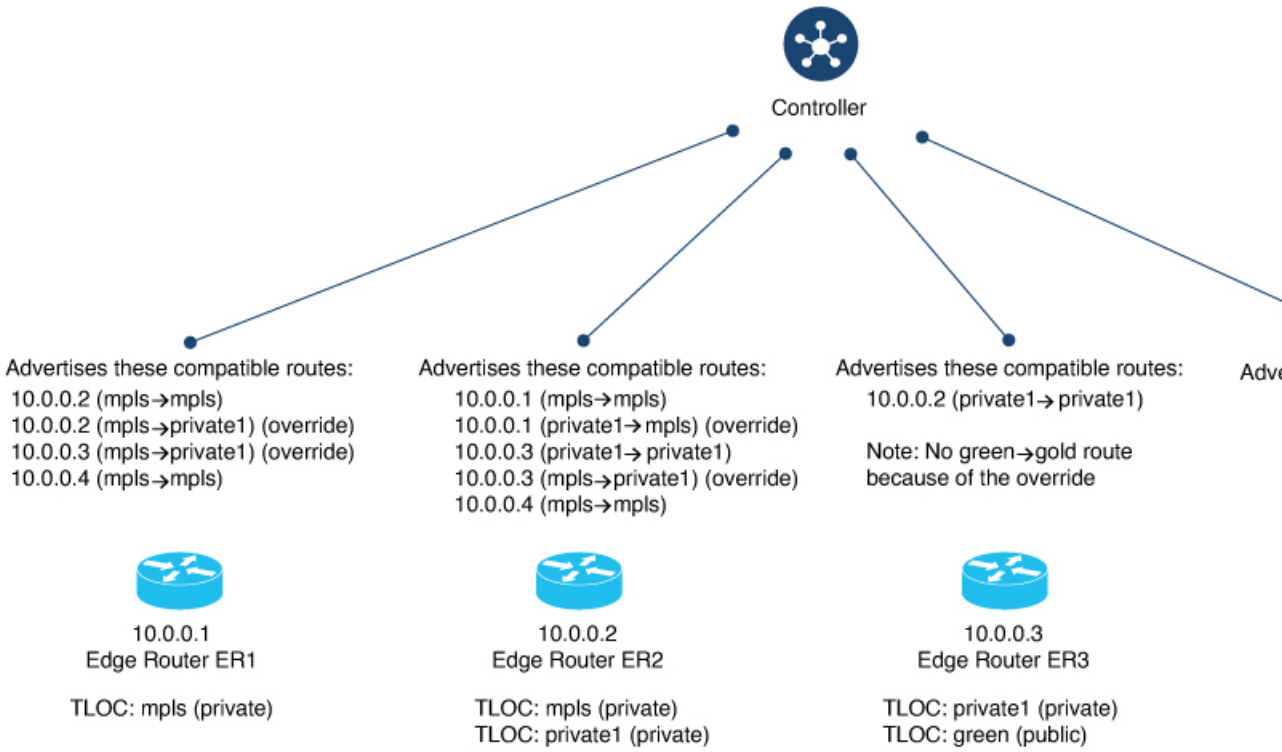

Routers in the network update Cisco SD-WAN Controllers when the status of their TLOCs changes. This may include reconfiguring a TLOC to a different color.

To account for temporary unavailability of a TLOC due to flapping, there is a dampening interval to delay reporting changes of TLOC status. By default, it is 60 seconds, but it can be configured to a value from 60 to 1200 seconds. For information, see Configure the Update Interval for Route Filtering by TLOC Color Using a CLI Template.

Benefits of Cisco SD-WAN Controller route filtering by TLOC color

- **Avoiding the send path limit**: Cisco SD-WAN Controller route filtering by TLOC color helps prevent routers from reaching their send path limit (for example, a limit of 32 routes), even if there are more routes available for a particular prefix.
- **Prioritizing relevant routes**: If the send path limit is set to a low value and many routes are available, filtering ensures that only relevant routes are advertised to the device. This prevents the send path limit from being reached with irrelevant routes and helps avoid routing failures.

Supported devices for Cisco SD-WAN Controller route filtering by TLOC color

Cisco IOS XE Catalyst SD-WAN devices supported for Cisco SD-WAN Controller route filtering by TLOC color.

Prerequisites for Cisco SD-WAN Controller route filtering by TLOC color

For Cisco SD-WAN Controllers to determine the compatibility of paths, the colors of TLOCs must be configured according to convention.

For example, a TLOC handling an MPLS connection must have the color mpls.

Restrictions for Cisco SD-WAN Controller route filtering by TLOC color

- When you enable Cisco SD-WAN Controller route filtering by TLOC color in a network, ensure that all you enable it on all Cisco SD-WAN Controllers in the network.
- Scenarios in which route filtering by TLOC color is enabled on some Cisco SD-WAN Controllers and disabled on others within the same network are not supported.

Configure Cisco SD-WAN Controller route filtering by TLOC color using a CLI template

- Enable route filtering using a CLI template
- Configure the update interval for route filtering by TLOC color using a CLI template
- Override default TLOC color compatibility for Cisco SD-WAN Controller route filtering by TLOC color using a CLI template

Enable route filtering using a CLI template

Before you begin

By default, CLI templates execute commands in global configuration mode. For more information about using CLI templates, see *CLI Add-On Feature Templates* and *CLI Templates*.

This configuration applies to a Cisco SD-WAN Controller.

Procedure

Step 1 Enter OMP mode.

omp

Step 2 Enter filter-route configuration mode.

filter-route

Step 3 Enable route filtering.

outbound tloc-color

```
omp
 filter-route
  outbound tloc-color
!
```

Configure the update interval for route filtering by TLOC color using a CLI template

Before you begin

By default, CLI templates execute commands in global configuration mode.

For more information about using CLI templates, see *CLI Add-On Feature Templates* and *CLI Templates*.

This configuration applies to a Cisco IOS XE Catalyst SD-WAN device.

Procedure

Step 1 Enter OMP configuration mode.

omp

Step 2 Configure the update interval, in seconds, in the range 60 to 1200.

timers
tloc-color-cap-update-interval *interval*

Example:

```
omp
  no shutdown
  ecmp-limit       6
  graceful-restart
  no as-dot-notation
  timers
   holdtime                       15
   tloc-color-cap-update-interval 120
   graceful-restart-timer         120
```

```
  exit
  address-family ipv4
   advertise ospf external
   advertise connected
   advertise static
  !
  address-family ipv6
   advertise ospf external
   advertise connected
   advertise static
  !
 !
!
```

Override default TLOC color compatibility for Cisco SD-WAN Controller route filtering by TLOC color using a CLI template

Before you begin

By default, CLI templates execute commands in global configuration mode.

For more information about using CLI templates, see *CLI Add-On Feature Templates* and *CLI Templates*.

You can override the default logic if necessary and do one of the following:

- Configure two TLOC colors to be compatible even if they are incompatible by default.
- Configure two TLOC colors to be incompatible even if they are compatible by default.

This may be helpful in specific unconventional scenarios.

This configuration applies to a Cisco SD-WAN Controller.

Procedure

Step 1 Enter system mode.

`system`

Step 2 Enter TLOC color compatibility mode.

`tloc-color-compatibility`

Step 3 Enter one or more of the following.

- To configure two TLOC colors to be compatible, do the following:

  `compatible` *`first-color second-color`*

- To configure two TLOC colors to be incompatible, do the following:

  `incompatible` *`first-color second-color`*

This example does the following:

- Configures the lte and private1 TLOC colors to be compatible
- Configures the private1 and private2 TLOC colors to be compatible
- Configures the lte and default TLOC colors to be incompatible
- Configures the lte and 3g TLOC colors to be incompatible

```
system
 host-name vm1
 system-ip 10.0.10.1
 site-id               100
 tloc-color-compatibility
  compatible lte private1
  !
  compatible private1 private2
  !
  incompatible lte default
  !
  incompatible lte 3g
  !
 !
```

Monitor Cisco SD-WAN Controller route filtering by TLOC color

View TLOC colors for a device

When applying route filtering, the controllers use this TLOC color information to determine which routes are relevant to a device.

Procedure

To view the list of the TLOC colors that a device advertises to Cisco SD-WAN Controllers, use the **show support omp peer peer-ip** command on a Cisco SD-WAN Controller.

Example:

The following example shows the TLOC colors that the peer device 10.0.0.15 is advertising—in this case, lte and 3g.

```
device#show support omp peer peer-ip 10.0.0.15 | inc color
ed bitmap: 0xc0, TLOC color supported list: lte 3g
```

Check TLOC color compatibility

Procedure

To check the compatibility of TLOC colors, use the **request support omp tloc-color-compat** command.

Example:

The following example requests information about whether the 3g and lte colors are compatible. These are both public TLOC colors, so they are compatible:

```
vsmart# request support omp tloc-color-compat 3g lte
Checking compatibility for colors:3g and lte
TLOC colors: 3g and lte are compatible
```

The following example requests information about whether the 3g and mpls TLOC colors are compatible. They are incompatible:

```
vsmart# request support omp tloc-color-compat 3g mpls
Checking compatibility for colors:3g and mpls
TLOC colors: 3g and mpls are incompatible
```

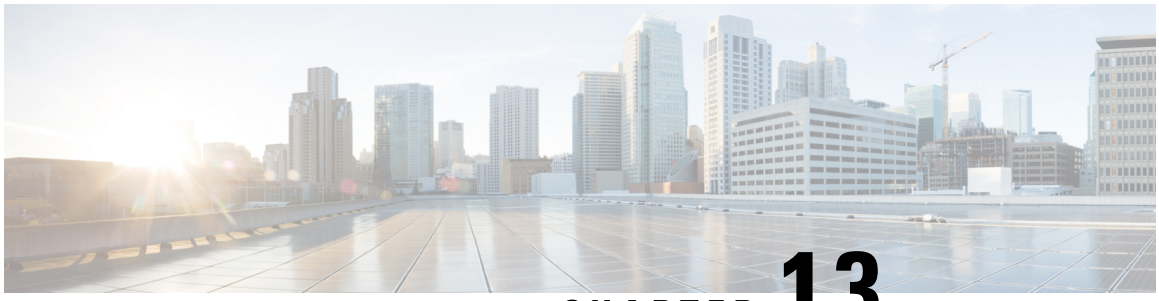

CHAPTER 13

Rendezvous Point Selection Process by a PIM BSR

Feature history for RP selection process by a PIM BSR

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

Table 42: Feature History Table

Feature Name	Release Information	Description
Dynamic Rendezvous Point (RP) Selection by a PIM Boot Strap Router (BSR)	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco SD-WAN Release 20.5.1	This feature adds support for automatic selection of an RP candidate using a PIM BSR in an IPv4 multicast overlay. There is no single point of failure because every site has a local RP. A Cisco IOS XE Catalyst SD-WAN device is selected as the RP, not a service-side device.

RP selection process by a PIM BSR

Summary

PIM uses a BSR to discover and announce RP-set information for each group prefix to all the routers in a PIM domain. This is the same function performed by Auto-RP, but a BSR is part of the PIM version 2 specification.

Cisco Auto-RP cannot co-exist with PIM BSR. Cisco Auto-RP mode must be disabled with spt-only mode.

Workflow

To avoid a single point of failure, you can configure several candidate BSRs in a PIM domain. A BSR is automatically selected from among the candidate BSRs. Bootstrap messages are used to determine which BSR has the highest priority. The router with the highest priority announces to all PIM routers in the domain that it is the BSR. Any router in the network can serve as a BSR candidate.

The elected BSR receives candidate-RP messages from all the candidate RPs in the domain. The bootstrap message sent by a BSR includes information about all of the candidate RPs.

An RP acts as the meeting place for sources and receivers of multicast data. In a PIM-SM network, sources must send their traffic to the RP. This traffic is then forwarded to receivers down a shared distribution tree. By default, when the receiver's first hop router learns about the source, it sends a join message directly to the source and creates a source-based distribution tree between the source and receiver. This source tree includes an RP only if the RP is within the shortest path between the source and receiver.

For a BSR to work for any multicast stream that spans across Cisco Catalyst SD-WAN sites, SPT-only mode is mandatory. For a BSR within a local-site multicast stream within a Cisco Catalyst SD-WAN site, it is not necessary to enable SPT-only mode.

If you have two Cisco IOS XE Catalyst SD-WAN devices in the same site, every Cisco IOS XE Catalyst SD-WAN device needs to be configured as a replicator for traffic to flow.

Benefits of PIM BSR

- IPv4 support.
- Dynamic rather than static selection of an RP.
- Automatic failover if one RP is not available.
- RP discovery is handled by a BSR.
- Configuration of multiple RP candidates for the same group range.
- Selection of a Cisco IOS XE Catalyst SD-WAN device as the RP.

Restrictions for PIM BSR

- IPv6 is not supported.
- Bidirectional PIM is not supported for IPv4.
- BSR is not supported in a hub-and-spoke topology on Cisco IOS XE Catalyst SD-WAN devices.

Topology for RP selection by a PIM BSR

Summary

The following is a sample topology for RP selection by a PIM BSR on Cisco IOS XE Catalyst SD-WAN devices:

Workflow

Figure 10: Topology for PIM BSR Selection

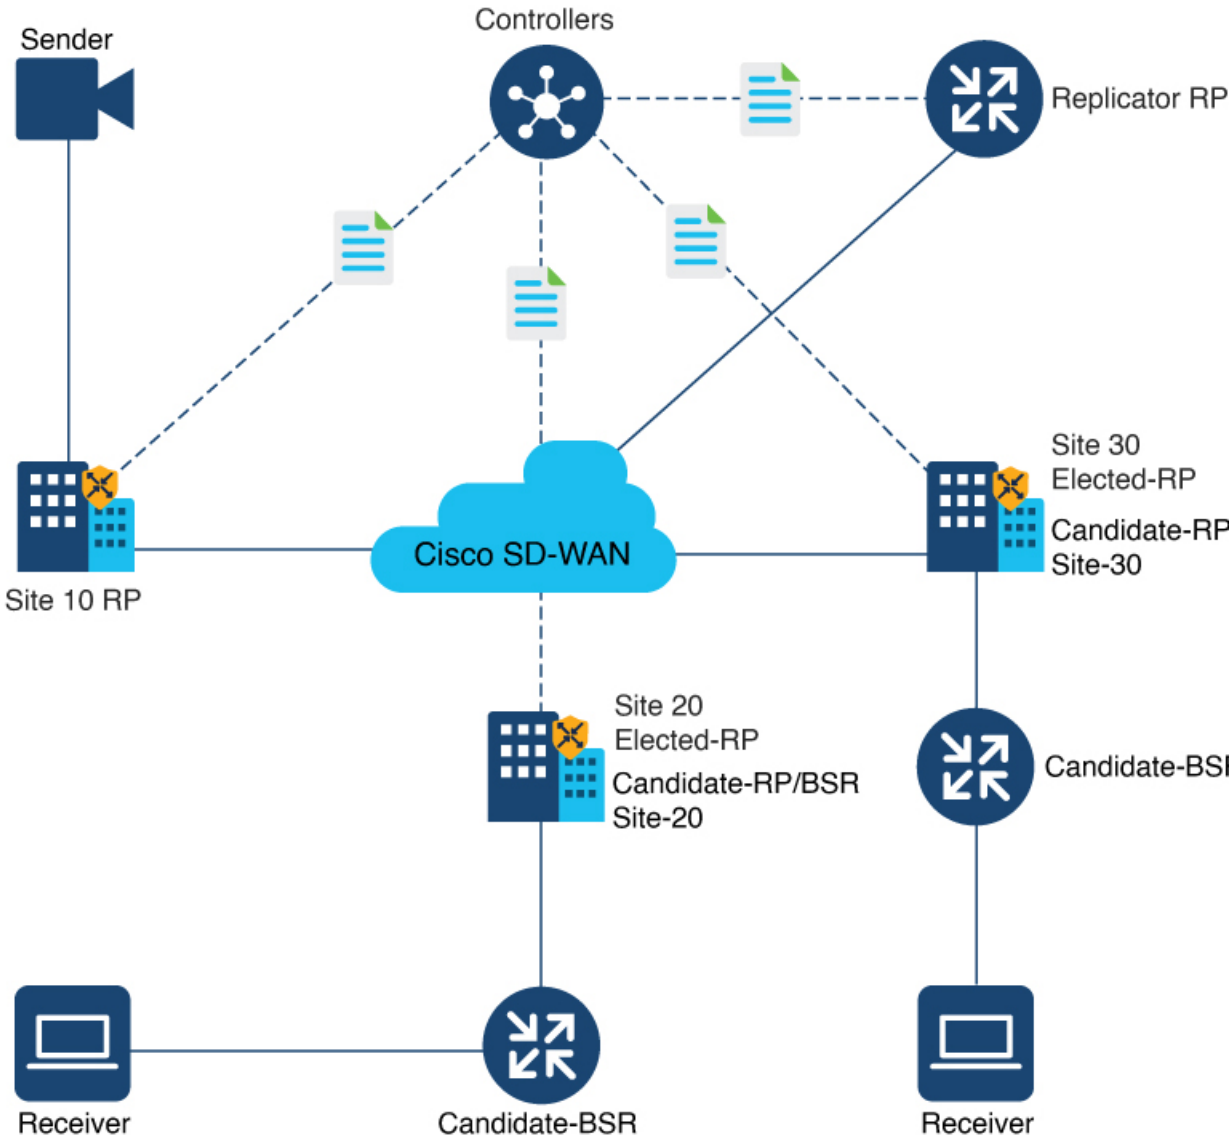

Configure a PIM BSR

Before you begin

Every Cisco Catalyst SD-WAN site must have its own RP.

SPT-only mode must be enabled on all Cisco Catalyst SD-WAN sites.

For a BSR to work for any multicast stream that spans across Cisco Catalyst SD-WAN sites, SPT-only mode is mandatory. For a BSR within a local-site multicast stream within a Cisco Catalyst SD-WAN site, it is not necessary to enable SPT-only mode.

For a PIM BSR to elect the RP, configure the following in Cisco SD-WAN Manager:

Procedure

Step 1 Multicast feature template with **SPT Only** set to **On** for the selected Cisco IOS XE Catalyst SD-WAN device.

Step 2 PIM feature template with an interface.

Step 3 RP candidate.

Step 4 BSR candidate.

Configure Shortest-Path Tree mode for a multicast feature template

Before you begin

In Cisco SD-WAN Manager, configure **SPT Only** mode to ensure that the RPs can communicate with each other using the shortest-path tree.

When configuring a BSR, configuration of **SPT Only** mode is mandatory.

Procedure

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

Step 2 Click **Feature Templates**.

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Click **Add Template**.

Step 4 From the **Select Devices** drop-down list, choose a Cisco IOS XE Catalyst SD-WAN device.

Step 5 In the **Template Name** field, enter a name containing up to 128 alphanumeric characters. In the **Description** field, enter a description containing up to 2048 alphanumeric characters.

Step 6 Under the **Basic Configuration** section for **SPT Only**, choose **On**.

Step 7 To enable the **Local Replicator** on the device, choose **On** (otherwise keep it set to **Off**).

Step 8 To configure a replicator, choose **Threshold**, and specify a value. (Optional, keep it set to the default value if you are not configuring a replicator).

Step 9 Click **Save**.

Configure a Cisco IOS XE Catalyst SD-WAN device as SPT-only using a CLI template

Procedure

Step 1 Configure a Cisco IOS XE Catalyst SD-WAN device as spt-only

```
Device(config)# sdwan multicast address-family ipv4 vrf 1
spt-only
```

Step 2 Use the **show platform software sdwan multicast remote-nodes vrf** command to verify that system IP addresses are configured with spt-only mode:

```
Device# show platform software sdwan multicast remote-nodes vrf 1

Multicast SDWAN Overlay Remote Nodes (* - Replicator):
                                          Received              Sent
                 SPT-Only              (X,G)      (S,G)      (X,G)      (S,G)
 System IP       Mode         Label   Join/Prune Join/Prune Join/Prune Join/Prune
 172.16.255.11   Yes          1003       0/0        0/0        0/0        0/0
 172.16.255.14   Yes          1003       0/0        0/0        1/0       10/10
 172.16.255.16   Yes          1003       0/0        0/0        0/0        0/0
 172.16.255.21   Yes          1003       0/0        0/0        0/0        0/0
```

Example:

Sample Multicast Configuration With SPT-Only

```
Device(config)# sdwan
Device(config)# multicast
Device(config)# address-family ipv4 vrf 1
Device(config)# spt-only
!
```

Configure a PIM feature template and add an interface

Configure a PIM feature template and add an interface for an RP and the BSR candidate.

Procedure

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

Step 2 Click **Feature Templates**.

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Click **Add Template**.

Step 4 From the **Select Devices** drop-down list, choose a Cisco IOS XE Catalyst SD-WAN device.

Step 5 Under **Other Templates**, choose **Cisco PIM**.

Step 6 In the **Template Name** field, enter a name containing up to 128 alphanumeric characters. In the **Description** field, enter a description containing up to 2048 alphanumeric characters.

Step 7 Click **Interface** > **New Interface**. In the **Interface Name** field, specify an interface with a value.

Step 8 In the **Query Interval (seconds)** field, the field auto-populates. In the **Join/Prune Interval (seconds) field**, the field auto-populates.

Step 9 Click **Add** and then click **Save**.

Configure the RP candidate

Configure the same Cisco IOS XE Catalyst SD-WAN device as the candidate RP for all multicast groups or selective groups.

Procedure

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

Step 2 Click **Feature Templates**.

In Cisco SD-WAN Release 20.7.x and earlier releases, **Feature Templates** is titled **Feature**.

Step 3 Edit the PIM feature template that you created by clicking **...** and then clicking **Edit**.

Step 4 Click **Basic Configuration**.

Step 5 Click **RP Candidate** > **New RP Candidate**.

Step 6 From the **Interface** drop-down list, choose the interface that you used for configuring the PIM feature template.

a. (Optional) In the **Access List** field, if you have configured the access list with a value, add the same value.

b. (Optional) In the **Interval** field, if you have configured the interval with a value, add the same interval value.

Step 7 In the **Priority** field, specify a higher priority on the Cisco IOS XE Catalyst SD-WAN device than on the service-side device.

Step 8 Click **Add**.

Step 9 Click **Update** to save your configuration changes.

Configure the BSR candidate

Procedure

Step 1 Repeat Step 1 through Step 4 from the Configure the RP candidate section.

Step 2 Click **BSR Candidate**.

Step 3 In the **BSR Candidate** field, chose the same interface from the drop-down list that you used for configuring the PIM feature template.

Step 4 (Optional) In the **Hash Mask Length** field, specify the hash mask length. Valid values for hash mask length range from 0 – 32.

Step 5 In the **Priority** field, specify a higher priority on the Cisco IOS XE Catalyst SD-WAN device device than on the service-side device.

Step 6 (Optional) In the **RP Candidate Access List** field, if you have configured the RP candidate access list with a value, add the same value. An RP candidate uses a standard access control list (ACL) where you can enter the name for the access list.

Step 7 Click **Update** to save your configuration changes.

Verify VRRP-Aware PIM Using the CLI

Sample VRRP-aware PIM configuration on router 1:

```
interface Vlan13
no shutdown
arp timeout 1200
vrf forwarding 1
ip address 10.0.0.1 255.255.255.0
ip pim sparse-mode
ip pim redundancy 1 vrrp dr-priority 200
ip tcp adjust-mss 1350
ip mtu 1500
ip igmp version 3
vrrp 1 address-family ipv4
vrrpv2
address 10.0.0.3
priority 200
timers advertise 100
track omp shutdown
vrrs leader 1
exit
```

Sample VRRP-aware PIM configuration on router 2:

```
interface Vlan13
no shutdown
arp timeout 1200
vrf forwarding 1
ip address 10.0.0.2 255.255.255.0
ip pim sparse-mode
ip pim redundancy 1 vrrp dr-priority 200
ip tcp adjust-mss 1350
ip mtu 1500
ip igmp version 3
vrrp 1 address-family ipv4
vrrpv2
address 10.0.0.3
priority 200
timers advertise 100
track omp shutdown
vrrs leader 1
exit
```

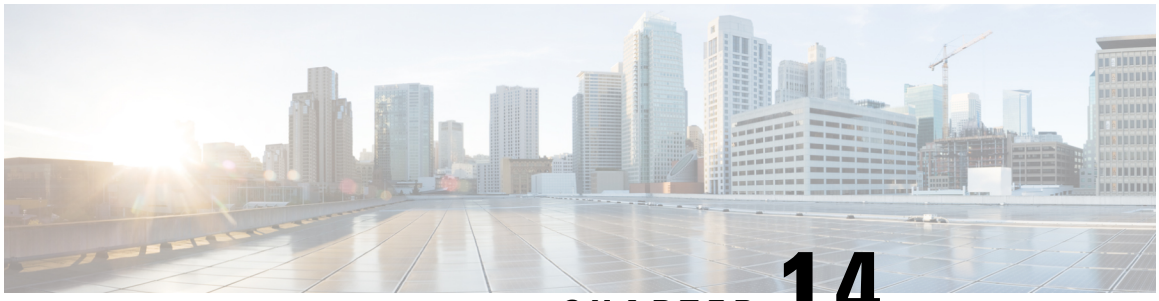

CHAPTER 14

Multicast Source Discovery Protocol

Feature history for MSDP

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

Table 43: Feature History

Feature Name	Release Information	Feature Description
Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN Domains	Cisco IOS XE Catalyst SD-WAN Release 17.11.1a Cisco Catalyst SD-WAN Control Components Release 20.11.1	This feature enables Multicast Source Discovery Protocol (MSDP) interoperability between Cisco IOS XE Catalyst SD-WAN devices in Cisco Catalyst SD-WAN and the devices in a non-SD-WAN setup. **Note** This feature does not provide support for MSDP peers formed between Cisco IOS XE Catalyst SD-WAN devices in the overlay network.

MSDP Routing Protocol Support

A Multicast Source Discovery Protocol (MSDP) is a multicast routing protocol that:

- facilitates interconnection of multiple Protocol Independent Multicast Sparse-Mode (PIM-SM) domains,
- enables a rendezvous point (RP) in a PIM-SM domain to maintain MSDP peering relationships with MSDP-enabled routers in other domains, and

- allows Cisco IOS XE Catalyst SD-WAN devices, from Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, to interoperate with other devices by converting Source Active (SA) messages received from MSDP peers into OMP routes, and vice-versa.

MSDP enables multicast sources in one PIM-SM domain to be discovered by other PIM-SM domains through the exchange of SA messages between MSDP peers.

When MSDP is enabled on Cisco IOS XE Catalyst SD-WAN devices, a rendezvous point (RP) in a PIM-SM domain maintains MSDP peering relationships with MSDP-enabled routers in other domains. When configured for interoperability, these devices convert Source Active (SA) messages received from MSDP peers into OMP routes and vice-versa. For more information about MSDP, see MSDP.

The following illustration depicts MSDP interoperability between Cisco IOS XE Catalyst SD-WAN devices in Cisco Catalyst SD-WAN and devices in a non-SD-WAN setup.

Figure 11: MSDP Interoperability

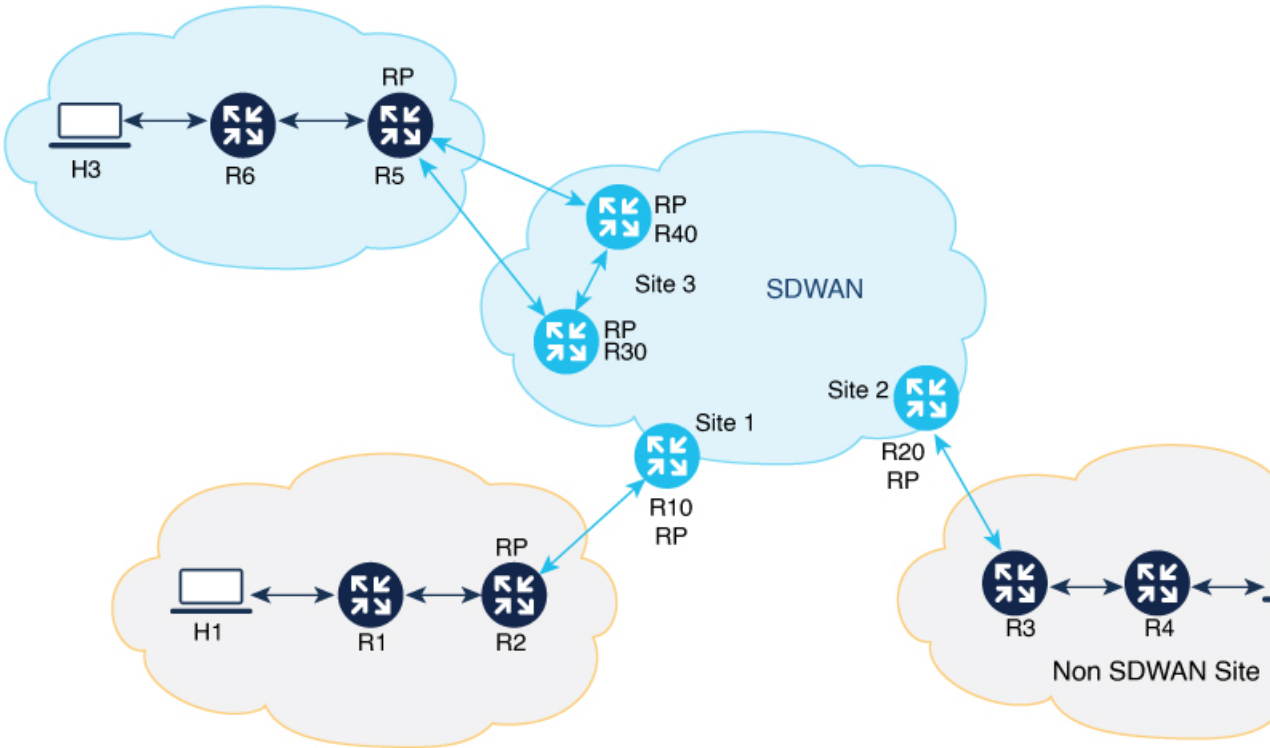

Single Homed Network

In the sample topology, the Cisco IOS XE Catalyst SD-WAN device R20 at site 2 has MSDP interoperability enabled. At the non-SD-WAN site, R3 serves as the rendezvous point (RP) for its PIM domain. MSDP peering is established between R3 and R20.

When source H2 sends multicast traffic to R4, R4 registers the data with R3. R3 then sends an MSDP Source Active (SA) message to R20. Because MSDP interoperability is enabled, R20 converts the received MSDP SA message into an OMP SA route and advertises it to all Cisco IOS XE Catalyst SD-WAN devices at other sites through the Cisco SD-WAN Controller.

When R10, a Cisco IOS XE Catalyst SD-WAN device at site 1, receives this OMP SA route, it converts the route into an MSDP SA message and advertises it to its MSDP peer R2 at the non-SD-WAN site. If R2 has receivers interested in the multicast group advertised in the MSDP SA message, it sends a (S,G) join towards the source. This process establishes an inter-domain source tree across the Cisco Catalyst SD-WAN.

As multicast packets arrive at R2 (the RP), R2 forwards them down its shared tree to group members within its domain. R20 withdraws the advertised OMP SA route only when the MSDP SA message expires.

Dual-Homed Network

A dual-homed network has two Cisco IOS XE Catalyst SD-WAN devices configured for MSDP interoperability. At the dual-homed Cisco Catalyst SD-WAN site 3, MSDP peering must be established among the Cisco IOS XE Catalyst SD-WAN devices R30, R40, and the non-SD-WAN device R5.

When the source registers its traffic with the RP R5, R5 sends an MSDP SA message to both R30 and R40. Upon receiving the MSDP SA message, R30 converts it into OMP SA routes and advertises these to all Cisco IOS XE Catalyst SD-WAN devices at other sites, as well as to R40 within site 3.

To prevent routing loops, an MSDP SA filter must be configured between R30 and R40 to drop SA messages that originate from other Cisco IOS XE Catalyst SD-WAN devices and sites via the Overlay Management Protocol (OMP).

At site 1, the Cisco IOS XE Catalyst SD-WAN device R10 receives two OMP SA routes for the same Source Group (S, G) and caches both routes. R10 then converts the OMP SA route into an MSDP SA message and advertises it to its MSDP peer R2 at the non-SD-WAN site. If R2 has receivers interested in the multicast group advertised in the MSDP SA message, it sends a (S, G) join towards the source. This sequence establishes an inter-domain source tree across the Cisco Catalyst SD-WAN.

MSDP supports the following scenarios where Cisco IOS XE Catalyst SD-WAN devices at the Cisco Catalyst SD-WAN sites are configured for MSDP interoperability with other devices located in the non-SD-WAN sites.

- Source devices located at the Cisco Catalyst SD-WAN sites, and receivers at the Cisco Catalyst SD-WAN and non-SD-WAN sites.
- Source devices located in the non-SD-WAN sites, and receivers at the Cisco Catalyst SD-WAN and non-SD-WAN sites.
- In Dual border sites, where two devices are configured for MSDP interoperability in Cisco Catalyst SD-WAN where sources and receivers are located in the Cisco Catalyst SD-WAN sites.
- In dual border sites, where two devices are configured for MSDP interoperability in non-SD-WAN, and where sources and receivers are located at the non-SD-WAN sites.
- A Replicator can be any Cisco IOS XE Catalyst SD-WAN device located in the Cisco Catalyst SD-WAN site. For more information about Replicators, see the Replicators section in PIM.

Benefits of support for MSDP to interconnect Cisco SD-WAN and non-SD-WAN

Facilitates MSDP interoperability between devices located at the Cisco SD-WAN sites and devices at the non-SD-WAN sites.

Restrictions for support for MSDP to interconnect Cisco SD-WAN and Non-SD-WAN

- Only one MSDP mesh group is supported per site in Cisco Catalyst SD-WAN.
- The MSDP peer devices must be located at the same site and cannot be spread across sites.

Configure MSDP to interconnect Cisco SD-WAN and Non-SD-WAN using a CLI Template

Before you begin

Note You cannot configure the MSDP interoperability using the feature template or the configuration groups in Cisco SD-WAN Manager.

Procedure

Step 1 Enable MSDP on Cisco IOS XE Catalyst SD-WAN device. For more information, see *Configure MSDP Using a CLI Template*.

Step 2 Configure MSDP interworking using a CLI template.

Note
By default, CLI templates execute commands in global config mode.

a) Enable MSDP on a Cisco IOS XE Catalyst SD-WAN device. For more information, see *Configure MSDP Using a CLI Template*

b) Configure a Cisco IOS XE Catalyst SD-WAN device for MSDP interoperability with other devices in the non-SD-WAN sites.

```
multicast address-family ipv4 vrf vrf-name
spt-only
msdp-interworking
```

Example:

The following is a complete configuration example to configure MSDP interoperability in Cisco Catalyst SD-WAN:

```
sdwan

 multicast address-family ipv4 vrf 1

  spt-only


  msdp-interworking
```

Verify MSDP configuration to interconnect Cisco SD-WAN and Non-SD-WAN

Procedure

The following is a sample output from the **show platform software sdwan multicast remote-nodes vrf** *1* command, which shows if MSDP interoperability is enabled or not.

Example:

```
Device# show platform software sdwan multicast remote-nodes vrf 1
            Multicast SDWAN Overlay Remote Nodes (* - Replicator, ^ - Delete Pending):
            Received                Sent
            SPT-Only MSDP             (X,G)      (S,G)      (X,G)      (S,G)
            System IP   Mode     I-Work Label  Join/Prune Join/Prune Join/Prune Join/Prune
            10.16.255.11 No       No     1003      0/0        0/0        0/0        1/0
            10.16.255.15 No       No     1003      1/0        1/0        0/0        0/0
            10.16.255.16 Yes      No     1003      1/0        1/0        0/0        0/0
            10.16.255.21 Yes      Yes    1003      0/0        0/0        0/0        0/0
```

Monitor MSDP configuration to interconnect Cisco SD-WAN and Non-SD-WAN

Procedure

Use the following show commands to monitor MSDP interoperability on Cisco IOS XE Catalyst SD-WAN devices:

```
Device# show ip msdp vrf 1 sa-cache
MSDP Source-Active Cache - 1 entries
(10.169.1.1, 12.169.1.1), RP 41.41.41.41, AS ?,6d20h/00:05:55, Peer 12.168.3.11
Device# show ip msdp vrf 1 count
SA State per Peer Counters, <Peer>: <# SA learned>
    12.168.3.11: 1
    12.168.11.15: 0
    12.168.12.12: 0
    12.168.14.14: 0
    12.168.5.24: 0
SA State per ASN Counters, <asn>: <# sources>/<# groups>
    Total entries: 1
    ?: 1/1
Device# show ip msdp vrf 1 summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
12.168.3.11      ?     Up       17w6d    0     1     ?
12.168.11.15     ?     Up       17w6d    0     0     ?
```

```
12.168.12.12     ?     Up       17w6d    0     0     ?
12.168.14.14     ?     Up       17w6d    0     0     ?
12.168.5.24      ?     Up       17w6d    1     0     ?
Device# show ip msdp vrf 1 peer 12.168.15.19 advertised-SAs
MSDP SA advertised to peer 12.168.15.19 (?) from mroute table

MSDP SA advertised to peer 12.168.15.19 (?) from SA cache

MSDP SA advertised to peer 12.168.15.19 (?) from mvpn sact table

20.169.1.1      13.169.1.1 RP 41.41.41.41 (?) 6d20h ref: 2
```

In the output above, the entry **MSDP SA advertised to peer 12.168.15.19 (?)** from **mvpn sact table** provides information about SA cache messages advertised to a peer based on the OMP SA routes received.

Example:

```
Device# show ip msdp vrf 1 peer 12.168.21.29
MSDP Peer 12.168.21.29 (?), AS ?
  Connection status:
    State: Up, Resets: 0, Connection source: GigabitEthernet5 (12.168.21.28)
    Uptime(Downtime): 16w4d, Messages sent/received: 169100/169106
    Output messages discarded: 82
    Connection and counters cleared 16w4d    ago
    Peer is member of mesh-group site3
  SA Filtering:
    Input (S,G) filter: sa-filter, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 0
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 10700/10827
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 0/10
```

Troubleshooting for MSDP

MSDP SA Cache Not Populated

Problem: MSDP SA cache is not populated on a Cisco IOS XE Catalyst SD-WAN device when a source in a site sends traffic.

Possible Cause: Check if there are any connectivity or configuration issues between the MSDP peers.

Solution:

1. Check the MSDP peering status between the Cisco IOS XE Catalyst SD-WAN device and the device in non-SD-WAN.

2. Verify that these commands **msdp-interworking** and **spt-only** are configured in the Cisco IOS XE Catalyst SD-WAN device.

OMP SA Route Not Advertised

Problem: A Cisco IOS XE Catalyst SD-WAN device does not advertise the OMP SA route when it receives a MSDP SA message from a MSDP peer.

Possible Cause: **msdp-interworking** configuration could be missing.

Solution: Configure the **msdp-interworking** command in the correct VRF.

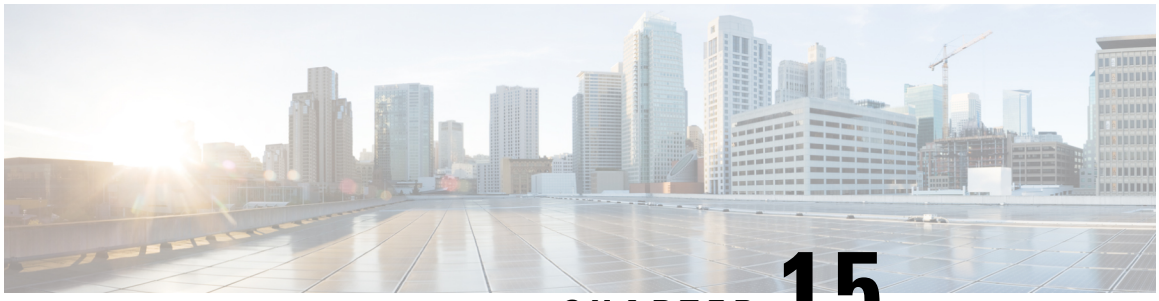

CHAPTER 15

Transport Gateway

Feature history for transport gateway

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

Table 44: Feature History Table

Feature Name	Release Information	Description
Transport Gateway	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Manager Release 20.12.1	A transport gateway operates as the hub in a hub-and-spoke routing topology. It offers the advantage of achieving this topology without requiring complex routing policy configuration. The following are some uses of a transport gateway: • Providing connectivity to routers in disjoint underlay networks • Serving as a gateway (hub) for all traffic in one discrete network to reach another discrete network, such as directing all local network traffic to a cloud gateway

Transport gateways for connecting networks in Cisco SD-WAN

A transport gateway is a network device that:

- connects routers that may or may not have direct connectivity,

- simplifies the process of providing connectivity between disjoint networks (such as between a physical LAN and a cloud-based network), and
- enables indirect connectivity without the complexity and limitations of manual control policy configuration.

A transport gateway facilitates communication between routers that are not directly connected, often bridging networks that are physically or logically separate (for example, connecting a traditional LAN to a cloud-based network).

A transport gateway connects routers that may or may not have direct connectivity. A common use case for transport gateways is to provide connectivity between routers in disjoint networks, such as between a physical LAN and a cloud-based network.

Without a transport gateway, one method of configuring indirect connectivity for these routers is to create a control policy that configures routes through an intermediate device with connectivity to both networks. This provides indirect connectivity between the disjoint routers. This approach has the following problems:

- Complexity: Configuring a control policy to advertise prefixes is complicated.
- Potential unavailable traffic endpoint: The control policy cannot detect whether a device or a configured route is unavailable. This can lead to packet loss if a route becomes unavailable.

Configuring a router to operate as a transport gateway solves the same issue, but with a simpler configuration process.

Figure 12: Transport Gateway

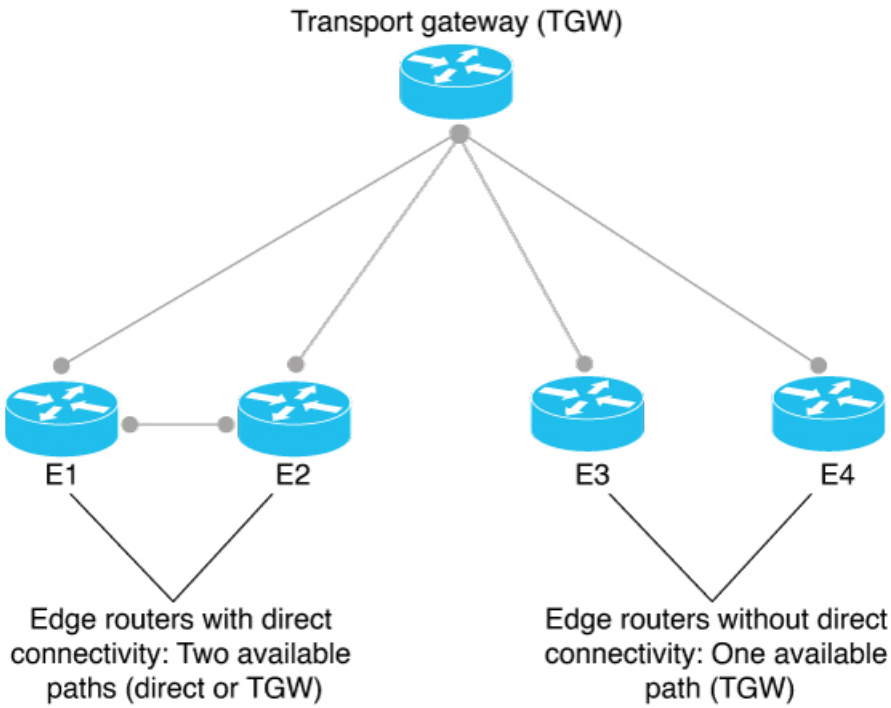

In the context of Cisco Catalyst SD-WAN, you can efficiently configure a hub-and-spoke routing topology by using transport gateways as hubs. This enables you to create the hub-and-spoke topology without requiring complex routing policy configuration. For information, see *Hub-and-Spoke*.

How a router functions as a transport gateway

Summary

Starting from Cisco Catalyst SD-WAN Manager Release 20.16.1, routes that are reoriginated from a site through the transport gateway are filtered out by the Cisco SD-WAN Controller. These reoriginated routes are not sent back to the originating site or the sites which share same site ID as the originating site. The reoriginated routes are only distributed to different sites within the Cisco Catalyst SD-WAN network.

This change in routing mechanism is also backported to Cisco Catalyst SD-WAN Manager Release 20.15.2, and Cisco Catalyst SD-WAN Manager Release 20.12.5.

When a router is configured to function as a transport gateway, it does the following for each route that it learns from the Cisco SD-WAN Controllers:

1. The transport gateway re-originates each route, substituting its own TLOCs as the next hop for the routes. This means that it substitutes its TLOCs as the next hop for each route.
2. The transport gateway advertises the re-originated routes to the Cisco SD-WAN Controllers.
3. The transport gateway attaches its own affinity attribute to routes that it re-originates. In scenarios in which routers in the network have re-originated routes available from more than one transport gateway, the routers apply affinity group preference logic to choose a route.

In the following illustration, E11 advertises prefix P1 and E22 advertises prefix P2. E11 and E22 are disjoint—they do not have direct connectivity. The transport gateway re-originates routes from E11 and E22, providing a P1 route to E22 and a P2 route to E11.

Workflow

Figure 13: Transport Gateway Re-Originating Routes

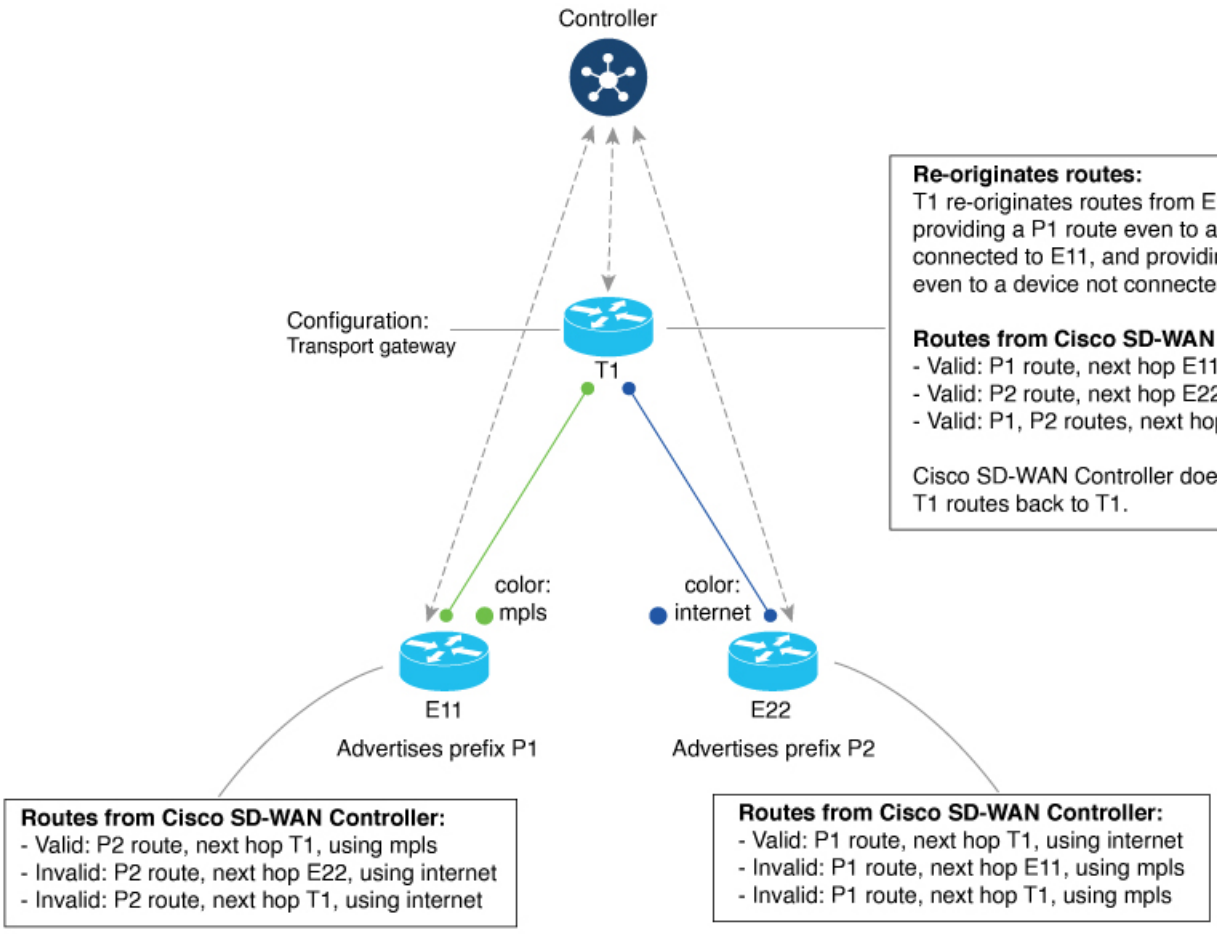

Site type

One part of configuring networks to use transport gateways is assigning a site type parameter to routers in the network. Site type helps to classify the intended function of a router, helping to define its position within the topology. Site type values include br, branch, cloud, spoke, type-1, type-2, and type-3.

After assigning site types, you can configure routers to prefer a transport gateway path only for traffic destined to a specific site type. This provides greater granularity when configuring a preference for transport gateway paths.

Site types are arbitrary, with no specific meaning, except br (border router) and spoke, which have specific uses for Multi-Region Fabric or intent-based hub-and-spoke topology, respectively.

Site Type Inheritance

Every OMP vRoute and TLOC originated from a router inherits the site type attributes of the router.

For information about configuring a site type for a router, see Configure the Site Type for a Router Using Cisco SD-WAN Manager.

OMP best path logic and transport gateway path preference

In general, when multiple paths are available between two routers, the overlay management protocol (OMP) applies best path selection logic to choose the best path. The best path selection logic is biased toward paths with fewer hops.

Summary

When you have configured a transport gateway, you can configure routers to apply a specific preference for transport-gateway-re-originated paths, if available. This alters the OMP best path calculation to include the transport gateway, according to the details of the configuration, as described below.

For information about configuring the preference for transport-gateway-re-originated paths, see Configure the Transport Gateway Path Preference.

Workflow

This table describes the best path logic.

	Router Configuration		Resulting Best Path Behavior
	Transport Gateway Path Behavior	**Specify Site Type(s)**	
1.	Not configured	Not applicable	(This is the default behavior.) Prefer a direct path.
	Prefer Transport Gateway Path	No	Prefer a transport-gateway path over a direct path.
	Prefer Transport Gateway Path	Yes	For a transport-gateway path that matches a specified site type, prefer a transport-gateway path over a direct path. For a transport-gateway path that does not match a specified site type, prefer a direct path over a transport-gateway path.
	Do ECMP Between Direct and Transport Gateway Paths	No	Treat a direct path and a transport-gateway path as equal.
	Do ECMP Between Direct and Transport Gateway Paths	Yes	For a transport-gateway path that matches a specified site type, treat a direct path and a transport-gateway path as equal. For a transport-gateway path that does not match a specified site type, prefer a direct path over a transport-gateway path.

As described earlier, a transport gateway attaches its own affinity attribute to paths that it re-originates. In scenarios in which routers in the network have re-originated paths available from more than one transport gateway, the routers apply affinity group preference logic to choose a path.

How the transport gateway configuration works

Summary

1. To configure a router to function as a transport gateway, use a System feature template or CLI add-on template. See Configure a Router as a Transport Gateway Using Cisco SD-WAN Manager.
2. To configure routers to use the transport gateway path, use an OMP feature template or CLI add-on template. See Configure the Transport Gateway Path Preference Using Cisco SD-WAN Manager. You can configure the OMP logic as follows:
   - Prefer a transport gateway path over a direct path.
   - Prefer a transport gateway path only for specific traffic, according to the site type attribute. See Configure the Site Type for a Router Using Cisco SD-WAN Manager.
   - Consider direct paths and transport gateway paths as equal.

The following figure shows how routers in a network can operate with a transport gateway, preferentially directing all traffic or specific traffic through transport gateway routes.

Workflow

Figure 14: Edge Routers and Transport Gateway Path Preference

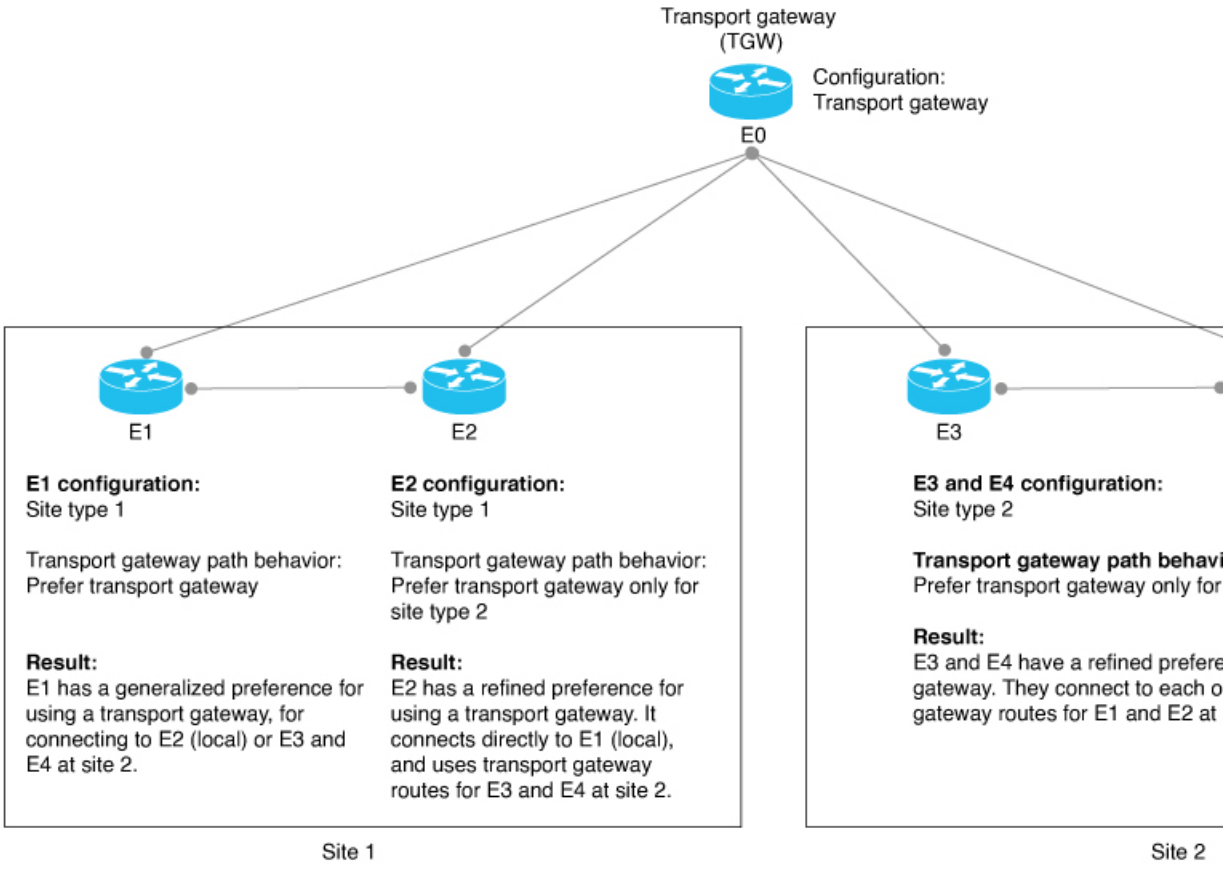

The devices in the illustration are configured as follows:

	Device	Configuration
1.	E0	**a.** Configure as a transport gateway. • By feature template: In a Cisco System template, use the **Transport Gateway** field. • By CLI add-on template: `system transport-gateway enable`

Device	Configuration

E1	**a.** Configure the site type as type-1.
	• By feature template:
	In a Cisco System template, use the **Site Type** field.
	• By CLI add-on template:
	`system site-type type-1`
	b. For best path, configure a preference for transport gateway routes.
	• By feature template:
	In an OMP template, use the **Transport Gateway Path Behavior** field. Choose the **Prefer Transport Gateway Path** option.
	• By CLI add-on template:
	`omp best-path transport-gateway prefer`

E2	**a.** Configure the site type as type-1.
	• By feature template:
	In a Cisco System template, use the **Site Type** field.
	• By CLI add-on template:
	`system site-type type-1`
	b. For best path, configure a preference for transport gateway routes for traffic to type-2 devices.
	• By feature template:
	In an OMP template, use the **Transport Gateway Path Behavior** field. Choose the **Prefer Transport Gateway Path** option. In the **Site Types** field, choose **type-2**.
	• By CLI add-on template:
	`omp best-path transport-gateway prefer transport-gateway-settings type-2`

Device	Configuration

E3 and E4	**a.** Configure the site type as type-2.
	• By feature template:
	In a Cisco System template, use the **Site Type** field.
	• By CLI add-on template:
	`system site-type type-2`
	b. For best path, configure a preference for transport gateway routes for traffic to type-1 devices.
	• By feature template:
	In an OMP template, use the **Transport Gateway Path Behavior** field. Choose the **Prefer Transport Gateway Path** option. In the **Site Types** field, choose **type-1**.
	• By CLI add-on template:
	`omp best-path transport-gateway prefer transport-gateway-settings type-1`

Use cases for transport gateways

In this use case, an organization needs to bridge a local network with a cloud services network, such as Azure or AWS. Edge routers in the local and cloud networks lack direct connectivity.

To create a transport gateway to bridge the local and cloud networks, network administrators configure the devices as follows:

Intent	Devices to Configure	Configuration
Configure the cloud gateway router with site type cloud.	Cloud gateway router	**1.** Configure the site type as cloud. • By feature template: In a Cisco System template, use the **Site Type** field. • By CLI template: **system site-type cloud**
Deploy a transport gateway to operate as a hub for cloud-destined traffic from devices in a local network. The transport gateway attracts the cloud-destined traffic and routes it to the cloud gateway for the cloud-based network.	Transport gateway router	**1.** Enable as a transport gateway. • By feature template: In a Cisco System template, use the **Transport Gateway** field. • By CLI template: **system transport-gateway enable**

Intent	Devices to Configure	Configuration
Traffic within the local network uses direct routes, not transport gateway routes. Traffic from the local network to the cloud uses a transport gateway route.	Edge routers in the local network	**1.** Use a transport gateway route for all cloud-destined traffic. • By feature template: In an OMP template, use the **Transport Gateway Path Behavior** field. Choose the **Prefer Transport Gateway Path** option. • By CLI template: **omp best-path transport-gateway prefer transport-gateway-settings cloud** **2.** Configure the site type as spoke. • By feature template: In a Cisco System template, use the **Site Type** field. • By CLI template: **system site-type spoke**

The following illustration shows the topology and configuration:

Figure 15: Transport Gateway Topology and Configuration

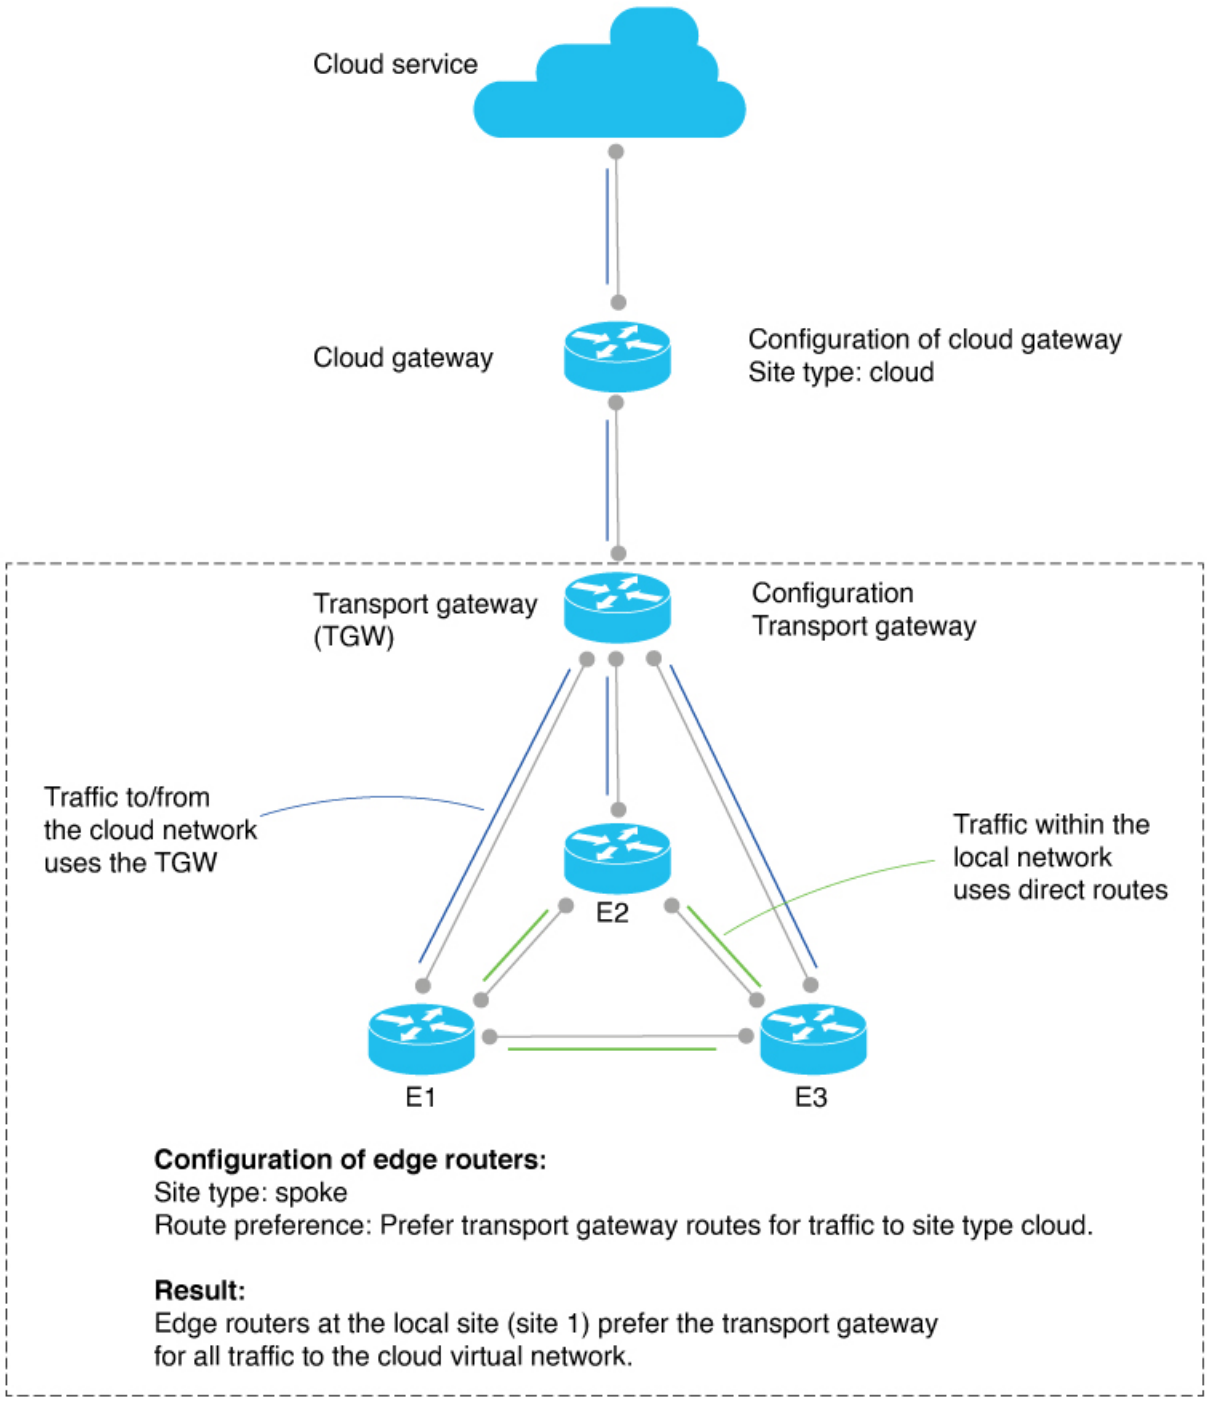

Restrictions for transport gateways

Restriction	Description
Resource demands of transport gateway functionality	Because of the resource demands of transport gateway functionality, we recommend enabling this only on a high-performance device with CPU and memory resources to handle the additional load. The specific resource requirements depend on your networking environment.
Multiple transport gateways: best path	If you enable transport gateway functionality on multiple devices, edge routers apply best path selection logic to determine the best path. This may include multiple transport gateway paths.
Multiple transport gateways: preventing routing loops	If you enable transport gateway functionality on multiple devices within network, the Cisco SD-WAN Controllers for the network do the following to avoid creating routing loops: When a Cisco SD-WAN Controller receives a route re-originated by one transport gateway, it does not advertise the route to another transport gateway. Avoiding advertising a transport gateway route to another transport gateway prevents routing loops.
On-demand tunnels	You cannot configure dynamic on-demand tunnels for a device configured as a transport gateway. However, edge routers that are not operating as transport gateways can use on-demand tunnels. For information about dynamic on-demand tunnels, see Dynamic On-Demand Tunnels in the *Cisco SD-WAN Systems and Interfaces Configuration Guide, Cisco IOS XE Release 17.x*.

Configure a router as a transport gateway

- Configure a router as a transport gateway using the CLI
- Configure a router as a transport gateway using SD-WAN Manager

Configure a router as a transport gateway using a CLI template

Before you begin

By default, CLI templates execute commands in global configuration mode.

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates.

Procedure

Step 1 Enter system configuration mode.

```
system
```

Step 2 Enable transport gateway functionality.

```
transport-gateway enable
```

Note

To disable transport gateway functionality, use the **no** form of the command.

```
system
  transport-gateway enable
```

Configure a router as a transport gateway using Cisco SD-WAN Manager

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

- To create a new System template, click **Add Template**, choose a device type, and click **Cisco System**.
- To edit an existing System template, locate a System template in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 In **Basic Configuration**, in the **Transport Gateway** field, choose **On**.

Step 5 Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure the transport gateway path preference

The following sections describe methods for configuring a router's best path decision to handle transport-gateway-re-originated paths.

- Configure the transport gateway path preference using a CLI Template
- Configure the transport gateway path reference using Cisco SD-WAN Manager

Configure the transport gateway path preference using a CLI Template

Procedure

Step 1 Enter sdwan configuration mode.

```
sdwan
```

Step 2 Enter system OMP configuration mode.

```
omp
```

Step 3 Configure the transport gateway path preference, using one of the following options:

best-path transport-gateway {prefer | ecmp-with-direct-path}

Option	Description
ecmp-with-direct path	For devices that can connect through a transport gateway and through direct paths, apply equal-cost multi-path (ECMP) to all available paths.
prefer	For devices that can connect through a transport gateway, use only the transport gateway paths, even if other paths are available.

Step 4 (Optional) Specify one or more site types to which to apply the transport gateway behavior. For information about how the Site Types parameter operates together with the Transport Gateway Path Behavior parameter, see OMP Best Path Logic and Transport Gateway Path Preference.

omp best-path transport-gateway-settings site-types *site-types*

Option	Description
site-types	Include one or more of the following site types, separated by spaces: **cloud**, **branch**, **br**, **type-1**, **type-2, type-3**

Note
To use this command, ensure that you use the **omp best-path transport-gateway prefer** command in the previous step.

The following example configures a device to prefer transport gateway routes.

```
sdwan
omp
  omp best-path transport-gateway prefer
```

The following example configures a device to prefer transport gateway routes only for traffic destined to sites with site type cloud.

```
sdwan
omp
  omp best-path transport-gateway prefer
  omp best-path transport-gateway-settings site-types cloud
```

Configure the transport gateway path preference using Cisco SD-WAN Manager

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

- To create a new OMP template, click **Add Template**, choose a device type, and click **Cisco OMP**.

- To edit an existing OMP template, locate a OMP template in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 In the **Best Path** section, in the **Transport Gateway Path Behavior** field, choose **Global** mode and choose one of the following options:

Option	Description
Do ECMP Between Direct and Transport Gateway Paths	For devices that can connect through a transport gateway and through direct paths, apply equal-cost multi-path (ECMP) to all available paths.
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.

Note
If you do not configure this field, by default, routers favor a direct path as the best path.

Step 5 (Optional) Click the **Site Types** field and choose one or more site types to which to apply the transport gateway behavior. For information about how the Site Types parameter operates together with the Transport Gateway Path Behavior parameter, see OMP Best Path Logic and Transport Gateway Path Preference.

Step 6 Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure the site type for a router

- Configure the site type for a router using a CLI template
- Configure the site type for a router using Cisco SD-WAN Manager

Configure the site type for a router using a CLI template

Before you begin

By default, CLI templates execute commands in global configuration mode.

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates.

Procedure

Step 1 Enter system configuration mode.

```
system
```

Step 2 Configure up to four site types for the router. Possible values are br, branch, cloud, spoke, type-1, type-2, and type-3.

```
site-type site-type
```

Note
To disable transport gateway functionality, use the **no** form of the command.

Example:

The following example configures a router site type as cloud:

```
system
  site-type cloud
```

The following example configure a router with site types cloud and branch:

```
system
  site-type cloud branch
```

Configure the site type for a router using Cisco SD-WAN Manager

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

- To create a new System template, click **Add Template**, choose a device type, and click **Cisco System**.
- To edit an existing System template, locate a System template in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 In **Basic Configuration**, click **Site Type** and choose a type from the drop-down list .

Step 5 Click **Save** if creating a new template, or **Update** if editing an existing template.

Verify the site type of a router using the CLI

Procedure

Use the **show sdwan omp summary** command on a device to verify the site type configuration of a router. The output includes a site-type field and the configured value.

Example:

In this example, the router is configured with a site type, spoke:

```
Device#show sdwan omp summary
...
site-type    SPOKE
...
```

Verify a transport gateway configuration using the CLI

Procedure

Use the **show sdwan running-config system** command on a device to check whether it is configured as a transport gateway. In the output, **transport-gateway enable** indicates that it is configured.

```
Device#show sdwan running-config system
system
system-ip            192.168.1.1
domain-id            1
site-id              11100
region 1
!
role                 border-router
transport-gateway enable
...
```

You can also use the **show sdwan omp summary** command on a device to check whether it is configured as a transport gateway. In the output, **transport-gateway enabled** indicates that transport gateway functionality is enabled.

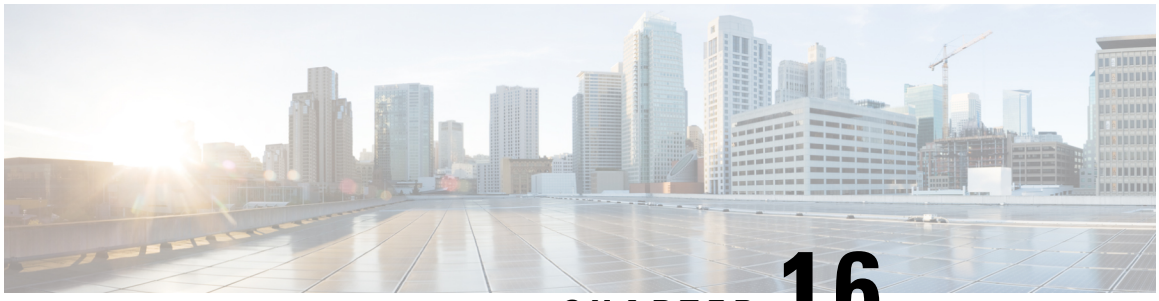

CHAPTER 16

Symmetric Routing

Feature history for symmetric routing

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

Table 45: Feature history

Feature name	Release information	Description
Symmetric Routing	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Control Components Release 20.12.1	You can use affinity groups, affinity group preferences, and RIB metric translation to ensure devices route traffic flows symmetrically across the network. Symmetric routing supports a wide range of network topologies, including Multi-Region Fabric. Transport gateways translate RIB metrics into control plane protocols such as BGP and OSPF to extend symmetric routing beyond the overlay network. This translation applies the path-preference configuration to routers outside the overlay network, including routers in a data center LAN.

Symmetric routing mechanisms for Cisco SD-WAN

A symmetric routing is a traffic flow property that

- uses the same path for traffic in both directions
- maintains a consistent return path between endpoints, and

- supports features that require bidirectional path symmetry.

Some networking functionality requires symmetric routing to operate correctly, including Cisco NBAR2, Cisco Zone-Based Firewall (ZBF), Cisco Unified Threat Defense (UTD), Cisco Application Quality of Experience (AppQoE), and network address translation (NAT).

Within a Cisco Catalyst SD-WAN network, you can use affinity groups, affinity group preference, control policy, and other mechanisms to configure the network so that the preferred route between two endpoints remains consistent for traffic in both directions. This configuration ensures symmetric routing for traffic flows between those endpoints. In some scenarios, you can also ensure symmetric routing for traffic flows that extend to a device outside the Cisco Catalyst SD-WAN overlay network.

TLOC behavior when moving away from symmetric NAT

When a TLOC starts behind symmetric NAT and then moves to any other NAT type such as full cone, port-restricted cone, or restricted cone, the TLOC does not update its public IP or port. The learned NAT type of the TLOC also does not update. As a result, some or all BFD sessions for this TLOC go down. To recover from this state, you can run **clear sdwan control connections** from the edge router.

Assumption about router operation

All of this applies only when routers stay operational during a traffic flow. If a router in the path becomes inoperable, traffic must take a new route. This change can cause temporary asymmetric routing.

Benefits of symmetric routing configuration

Before Cisco IOS XE Catalyst SD-WAN Release 17.12.1a configuring symmetric routing required complex and error-prone control policies in the overlay network. These policies set up hop-by-hop routing in both directions.

In service-side routing, it required complex route-maps to maintain path symmetry in both directions.

From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a onward, you can use affinity groups, affinity group preferences, and OMP metric redistribution to achieve symmetric routing. The following sections describe the details and supported scenarios.

Restrictions for symmetric routing

You cannot use both the **redistribute omp translate-rib-metric** command and the **redistribute omp metric** command together on the same device.

The **translate-rib-metric** option generates BGP attributes and OSPF metrics from OMP metrics, whereas the **metric** option configures the metrics explicitly. For information, see Translating OMP Metrics for Devices Outside of the Overlay Network.

Translating OMP metrics for devices outside of the overlay network

A router configured as a transport gateway and operating as a hub (TGW1 in the illustration below) may conduct traffic between devices within the Cisco Catalyst SD-WAN overlay network (WAN) and a device outside of the overlay network (LAN), such as DC1 in the following illustration. This is WAN-to-LAN traffic. Note that devices outside of the overlay network are not managed by Cisco Catalyst SD-WAN.

A transport gateway translates RIB metric information into parameters used by the BGP or OSPF protocols. It uses those parameters in its BGP or OSPF routing tables, and when the transport gateway advertises routes to its BGP or OSPF neighbors, it includes the RIB-derived parameters with the routes.

These RIB-derived parameters influence path selection by devices in the LAN, helping to ensure that the LAN chooses the same path for LAN-to-WAN traffic as the overlay network uses for WAN-to-LAN traffic.

Figure 16: Translating OMP metrics

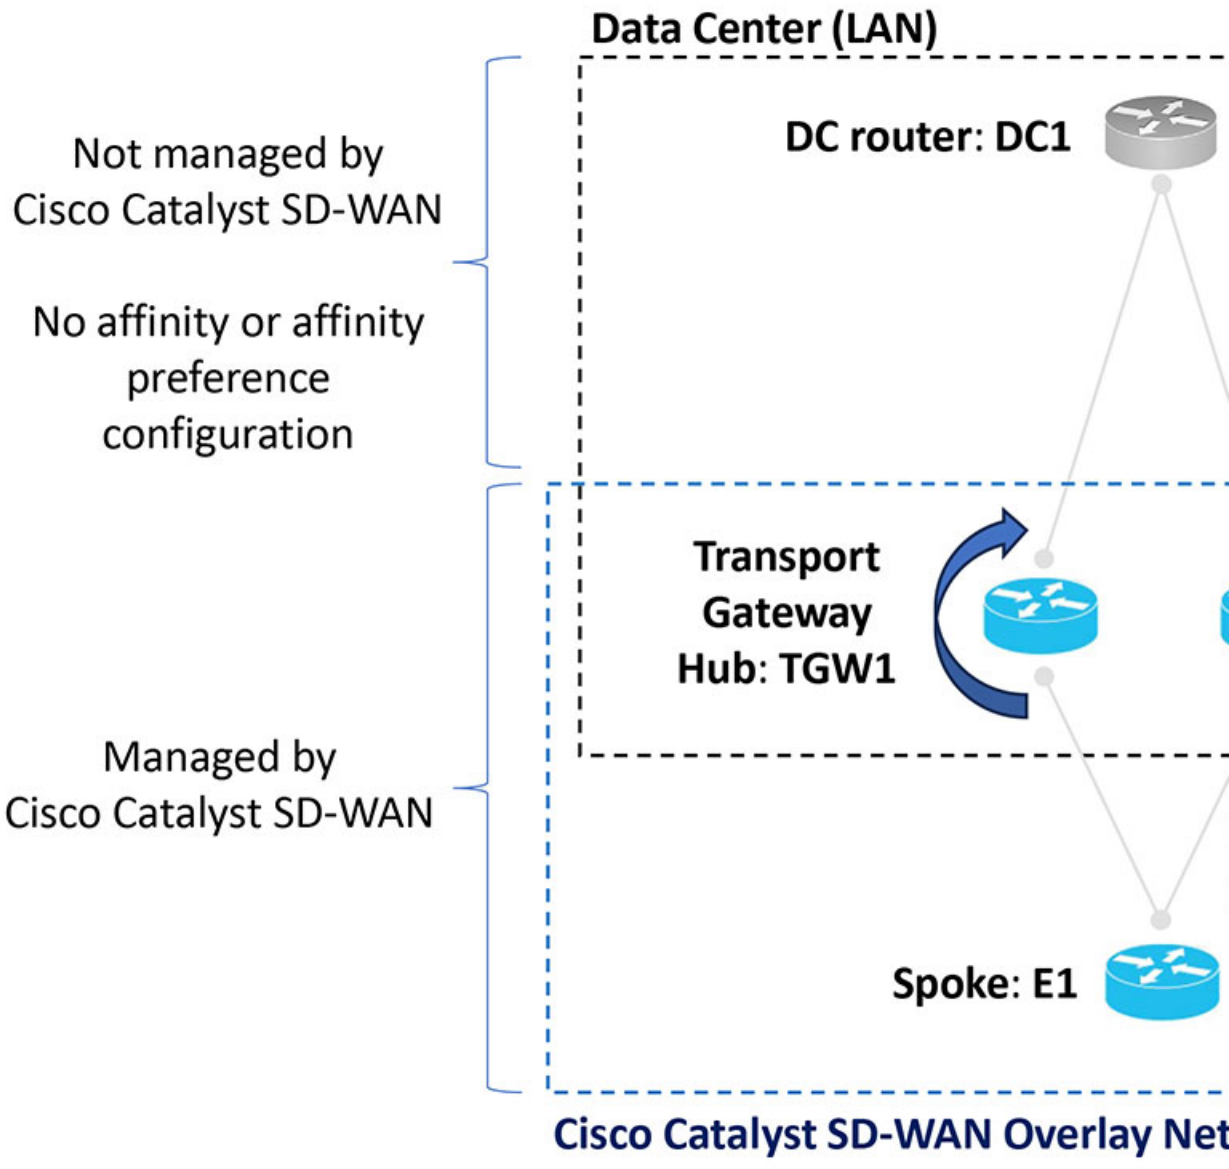

Translating OMP metrics to BGP attributes

When you enable a router to translate RIB metrics from OMP to BGP, the router uses the following OMP metric and attribute:

- OMP route metric

  Among the OMP metrics, one specific metric is named OMP.

- OMP AS-PATH

The router uses these to derive three BGP attributes:

- BGP MED
- BGP LOCAL_PREF
- BGP AS_PATH

For information about how to view OMP metrics for a route and the resulting BGP attributes, see Monitor RIB metric translation.

Translation from OMP to BGP

Table 46: Translation of OMP metrics to BGP attributes

BGP Attribute	How it is derived
BGP MED	Equal to the OMP route metric.
BGP LOCAL_PREF	255 – (OMP route metric)
BGP AS_PATH	Two possibilities: • If the **propagate-aspath** command is used: 1. If OMP AS-PATH is empty, then the router uses its own local AS value and repeats it (OMP route metric) times, with a maximum of 13 repetitions. 2. If OMP AS-PATH is not empty, then the router uses the OMP AS-PATH and prepends it with the first AS in the OMP AS-PATH (OMP route metric) times, with a maximum of 13 times. • If the **propagate-aspath** command is not used: The router uses a list of its own local AS value, repeated (OMP route metric) times, prepending the value up to a maximum of 13 repetitions.

AS-PATH propagation recommendation

In most scenarios, when you enable translation of RIB metrics using the **redistribute omp translate-rib-metric**command, also enable propagating the AS-PATH metric using the **propagate-aspath** command. If you omit this, the router treats the AS-PATH metric as empty.

The router includes these BGP attributes with the routes that it re-originates to a device in a LAN outside the overlay network that uses BGP.

BGP attributes without RIB metric translation

The table below shows combinations of OMP metrics and the BGP attributes that the router derives when RIB metric translation is not enabled.

Table 47: Translation from OMP to BGP without RIB metric translation enabled

	OMP metrics: Example combinations		Translation to BGP attributes: propagate-aspath enabled translate-rib-metric not enabled		
Example	OMP route metric	OMP AS-PATH	BGP MED	BGP LOCAL_PREF	BGP AS_PATH
1	0	100 101	1000	50	100 101
2	1	100 101	1	50	100 101
3	2	100 101	2	50	100 101
4	10	(empty)	10	50	(empty)
5	14	100 101	14	50	100 101

BGP attributes without RIB metric translation

The table below shows combinations of OMP metrics and the BGP attributes that the router derives when RIB metric translation is not enabled.

Table 48: Translation from OMP to BGP without RIB metric translation enabled

	OMP metrics: Example combinations		Translation to BGP attributes: propagate-aspath enabled translate-rib-metric not enabled		
Example	OMP route metric	OMP AS-PATH	BGP MED	BGP LOCAL_PREF	BGP AS_PATH
1	0	100 101	1000	50	100 101
2	1	100 101	1	50	100 101
3	2	100 101	2	50	100 101
4	10	(empty)	10	50	(empty)
5	14	100 101	14	50	100 101

Translating OMP metrics to an OSPF metric

If you do not configure the router to translate RIB metrics, it uses a default OSPF metric when it redistributes routes to a device outside the Cisco Catalyst SD-WAN overlay network. The default OSPF metric is 16777214 (hexadecimal FFFFFE).

When you enable the router to translate RIB metrics, it assigns the OMP route metric value as the OSPF metric. For example, if the OMP route metric is 10, the OSPF metric is also 10.

For information about viewing the OMP metrics for a route and the resulting BGP metrics, see Monitor RIB Metric Translation.

Symmetric routing scenarios

An overview of the configuration workflow helps you understand the scenarios in which Cisco Catalyst SD-WAN supports symmetric routing. The figures below shows

- a transport gateway scenario, and
- a Multi-Region Fabric scenario.

Transport gateway scenario

In the transport gateway scenario, the goal is to ensure symmetric routing between the spoke devices (E1 and E2 in the illustration) and the data center router (DC1).

Figure 17: Transport gateway scenario with a data center LAN

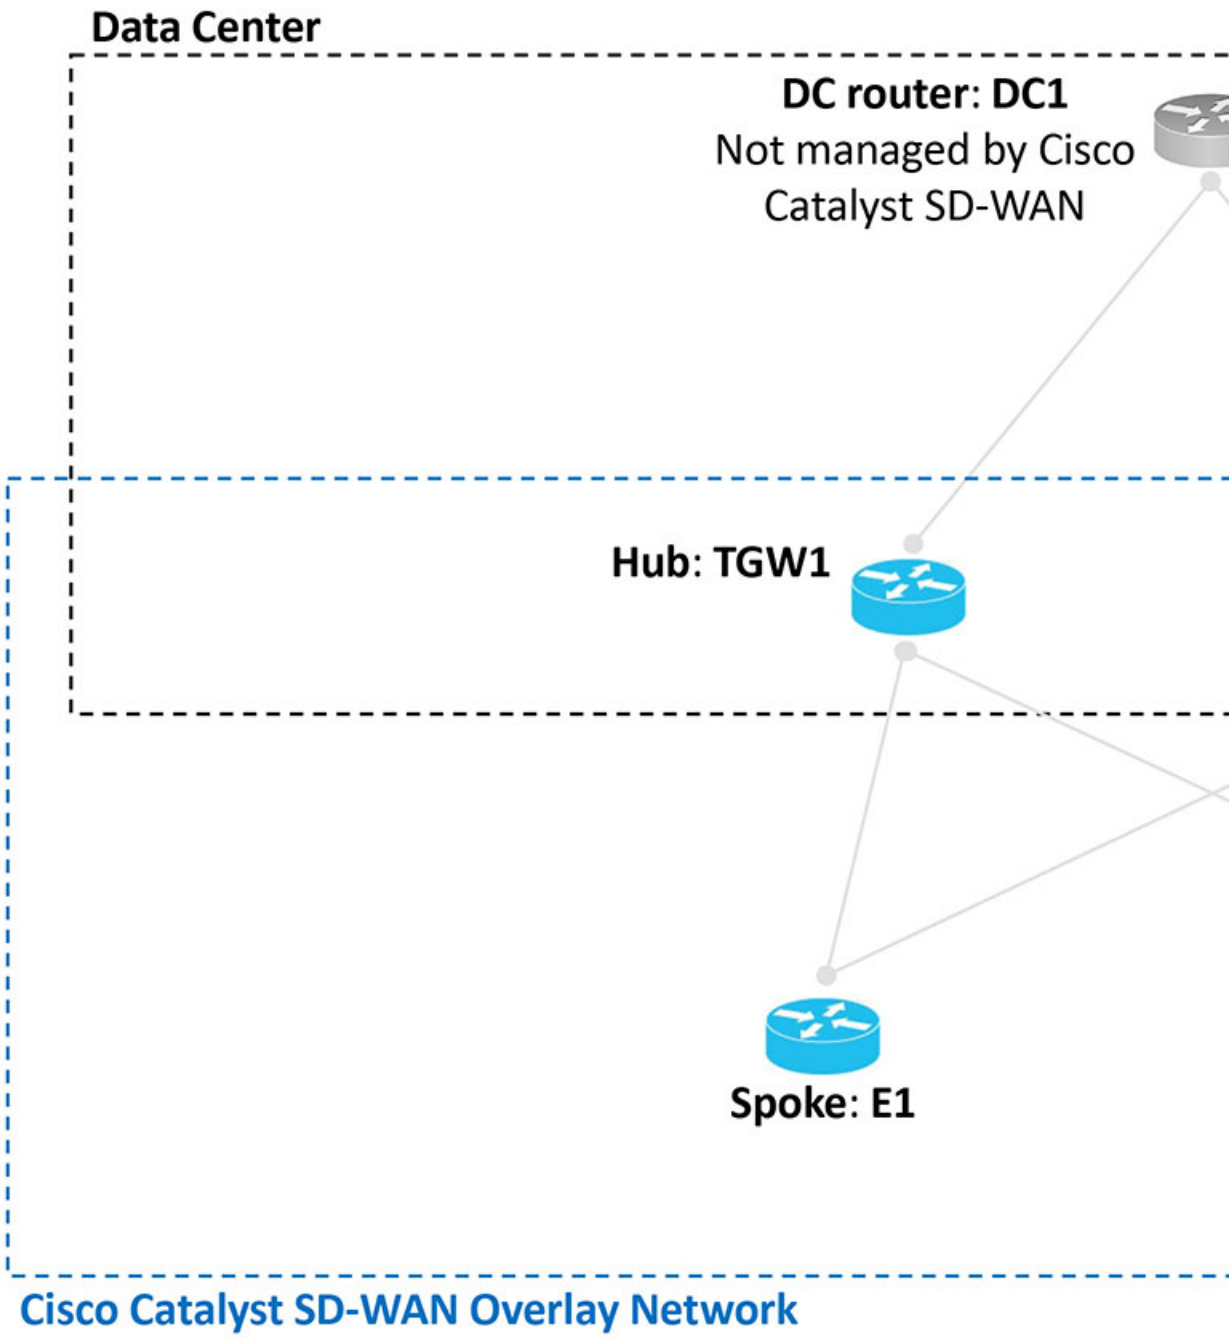

Multi-region fabric scenario

In the Multi-region fabric scenario, the goal is to ensure symmetric routing between the PC devices served by edge router ER11 in Region 1, and the PC devices served by ER21 in Region 2.

Figure 18: Multi-region fabric scenario

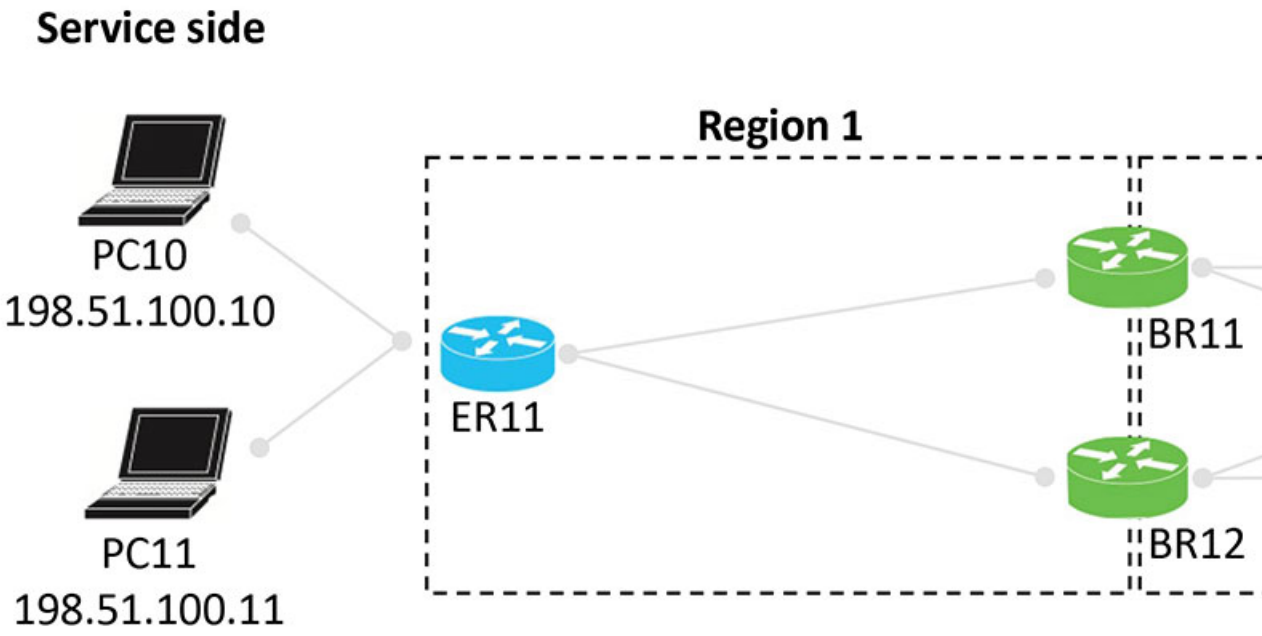

Configuration overview

The steps below provide an overview of the configuration required for symmetric routing.

Configuration step	Devices	Description
1. Configure affinity group preference	Spoke routers Edge routers in a multi-region fabric scenario	To ensure traffic symmetry within the overlay network, configure spoke routers (or edge routers in a multi-region fabric scenario) in the network with an affinity group preference. This can be a manually configured order of preference or automatic preference. With automatic affinity preference order, a spoke device or edge router prefers paths tagged with a lower affinity group number. For configuration instructions, see Configure a Router Affinity Group or Affinity Group Preference.

Configuration step	Devices	Description
2. Configure affinity groups	Transport gateways Border routers in a multi-region fabric scenario	To ensure traffic symmetry within the overlay network, configure border routers and transport gateways with (a) an affinity group number, or (b) affinity groups per VRF for some or all VRFs that the devices handle. You can configure both (a) and (b) together. For example, if a device has a VRF range of 1 to 10, you can configure a device as follows: • System-level affinity group 10 • Affinity groups per VRF: Affinity group 20 for VRF6 through VRF 10 The result is that vRoutes in the range 1 to 5 are tagged with affinity group 10 (from the system-level affinity group), and vRoutes in the range of 6 to 10 are tagged with affinity group 20. For configuration instructions, see Configure a Router Affinity Group or Affinity Group Preference.
3. Enable translation of RIB metrics	Transport gateways Border routers in a multi-region fabric scenario	To enable symmetric routing between the overlay network and a LAN, on the border routers or transport gateways that conduct traffic with a LAN, enable translation of RIB metrics for redistribution of OMP routes to LAN routing protocols. For a full explanation, see Translating OMP Metrics for Devices Outside of the Overlay Network. For configuration instructions, see Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template.

Example configurations for symmetric routing in transport gateway and multi-region fabric deployments

The illustrations below show the two scenarios described earlier, with an example configuration for each router, in accordance with the steps described here to ensure symmetric routing.

Figure 19: Transport gateway scenario with a data center LAN, showing a configuration for symmetric routing

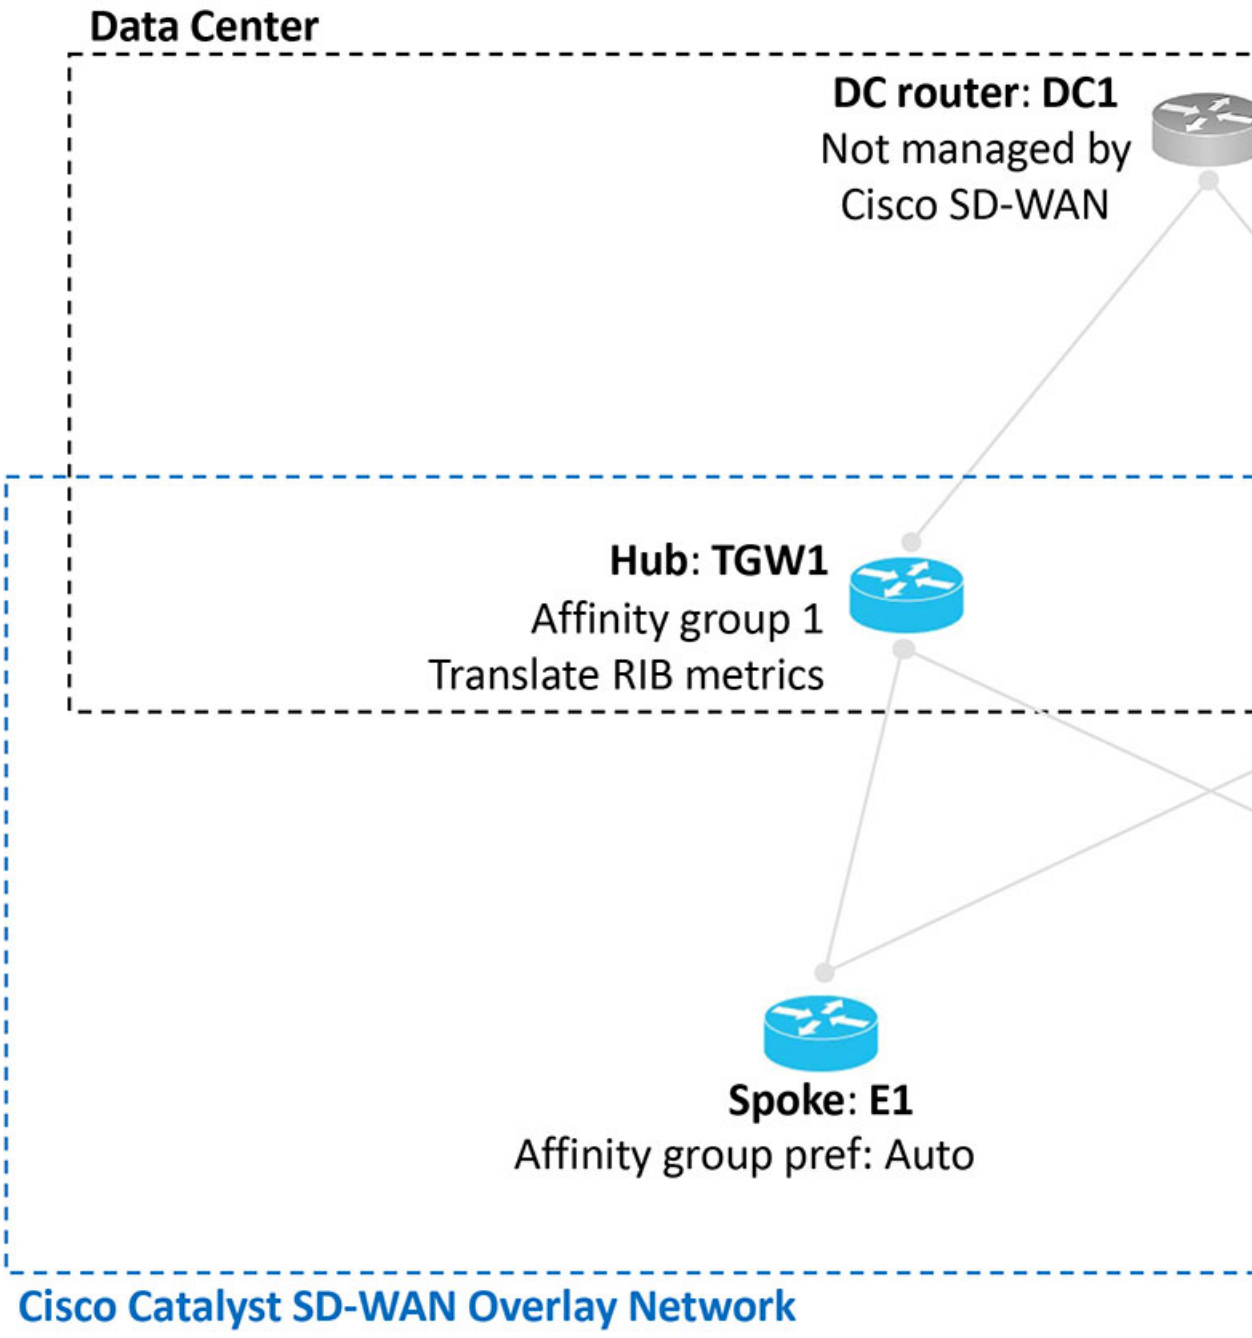

Figure 20: Multi-region fabric scenario, showing a configuration for symmetric routing

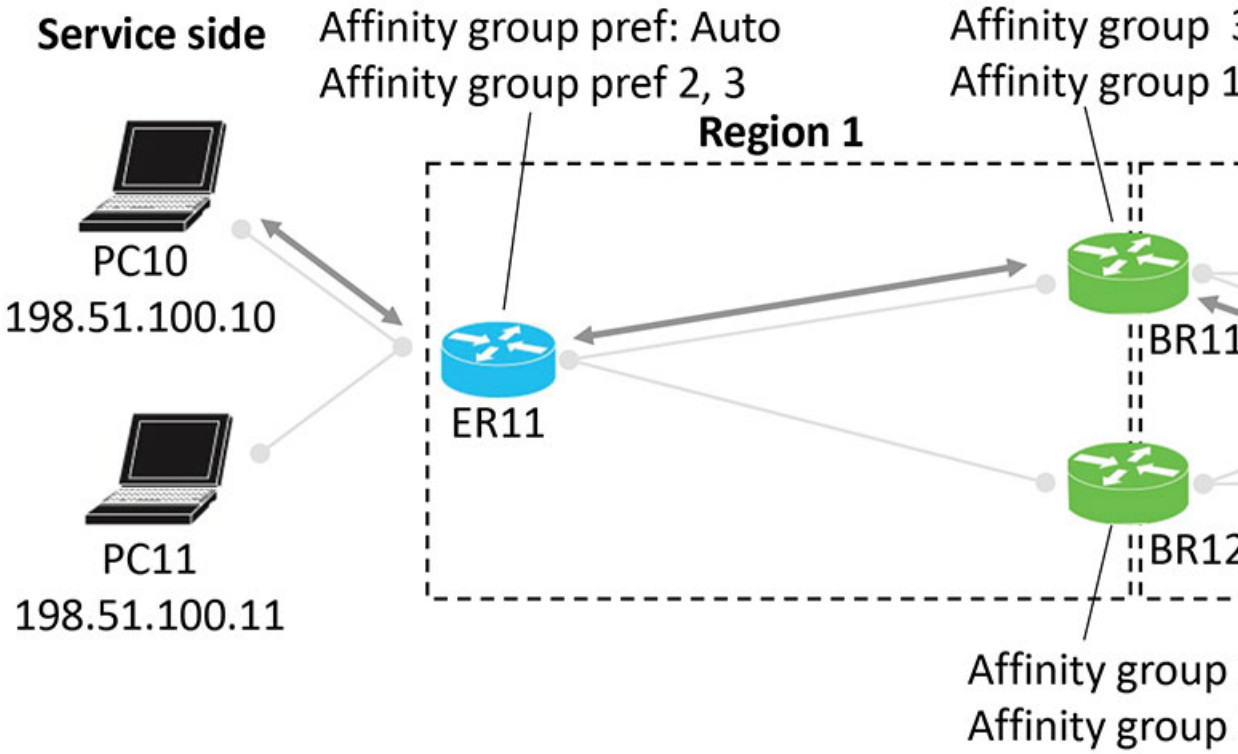

Example of configuration for symmetric routing and the mechanism

This comprehensive example describes how to configure border routers and edge routers in a Multi-Region Fabric (MRF) environment to achieve symmetric routing between PC devices behind ER11 (Region 1) and ER21 (Region 2).

The example focuses specifically on traffic between PC10 and PC20.

The step-by-step illustrations show how route re-origination and path preference ensure that traffic in both directions follows the same path across multiple hops.

Multi-region fabric scenario: configuration for symmetric routing

Figure 21: Multi-region fabric scenario, configuration for symmetric routing

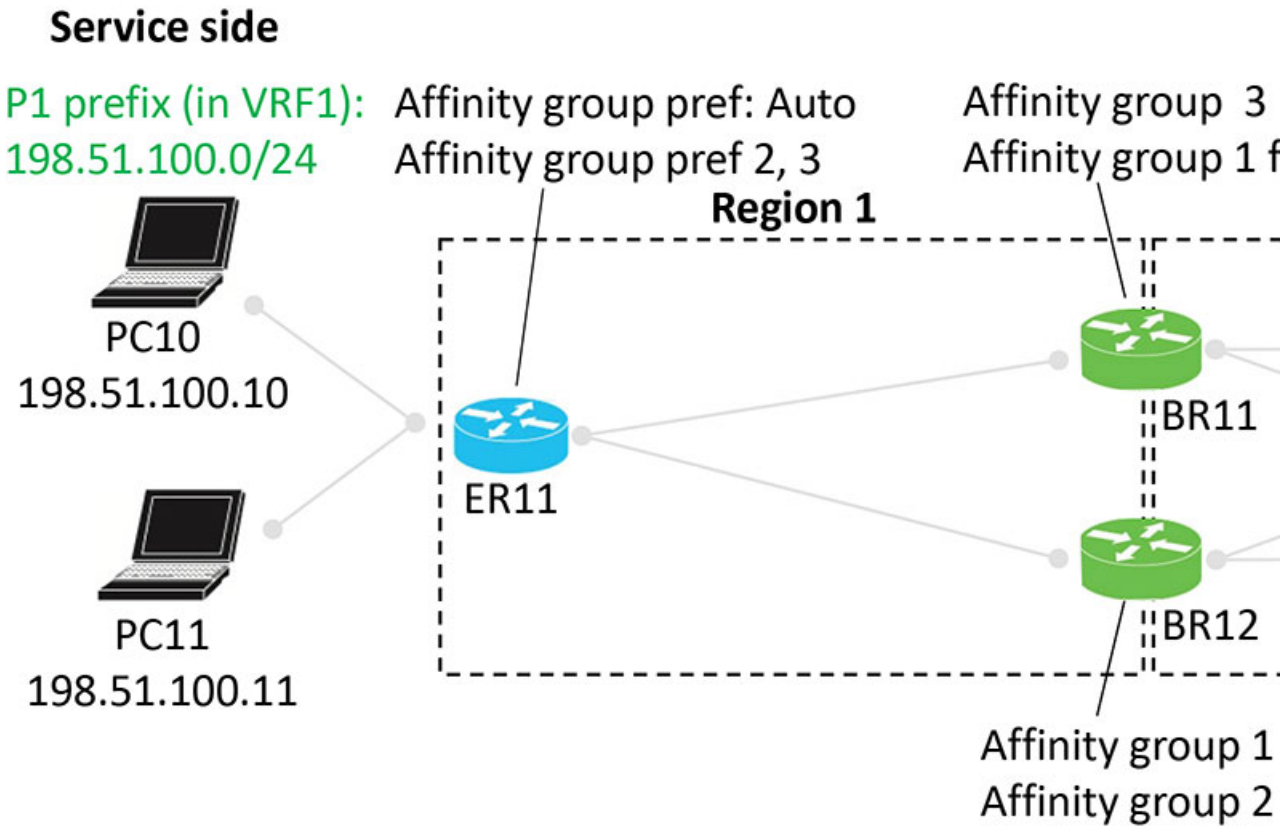

Advertising P1 routes

Edge router ER11 advertises P1 routes. These routes are re-originated as they move from Region 1 toward Region 2, passing through border routers that assign affinity groups and derived affinity groups (DAG).

Routers select preferred routes based on:

- Outside the core region: Affinity group preference
- Inside the core region: Lowest derived affinity group (dag) value

Figures:

Figure 22: Edge router ER11 advertises P1 routes

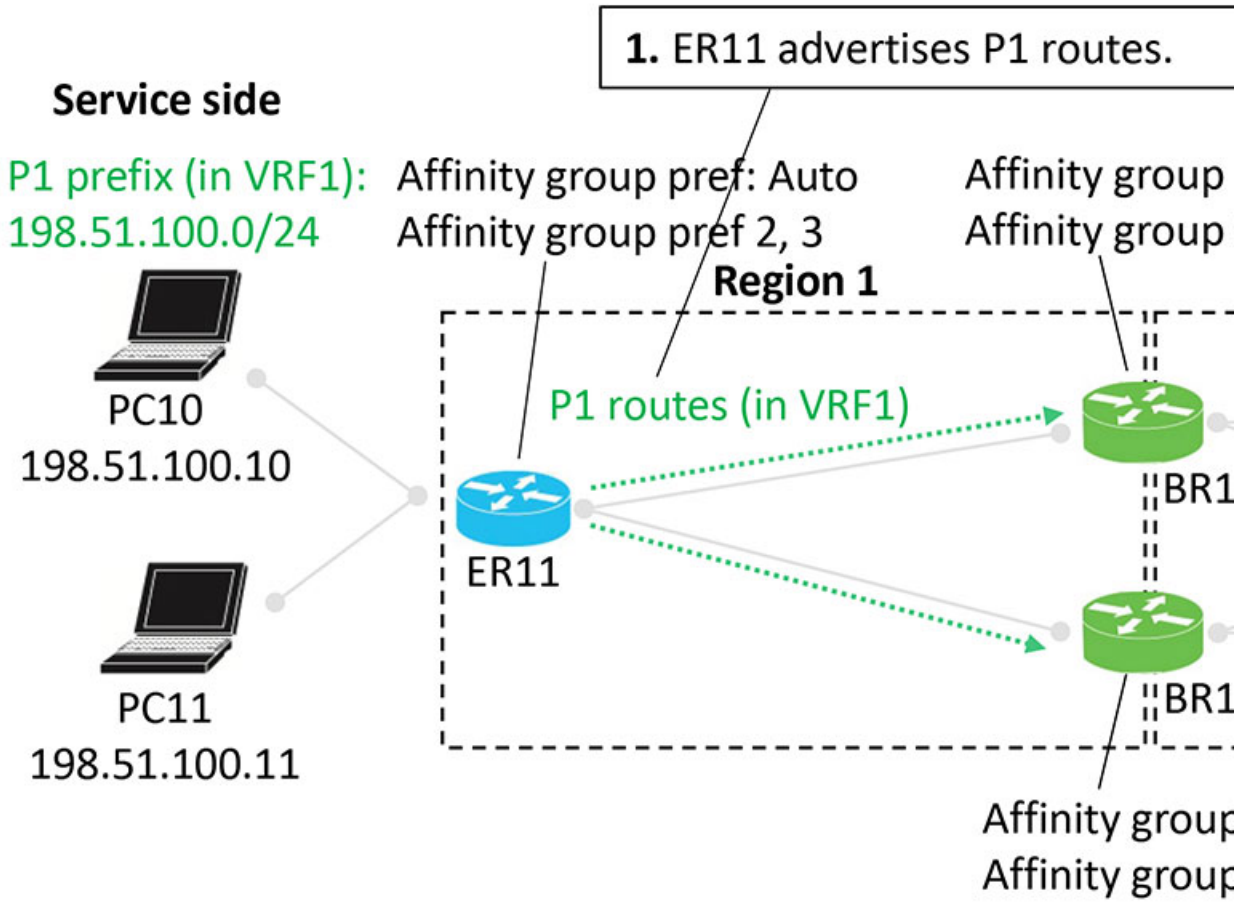

····▶	ER11 advertises P1 routes

Figure 23: Border routers BR11 and BR12 re-originate the P1 routes

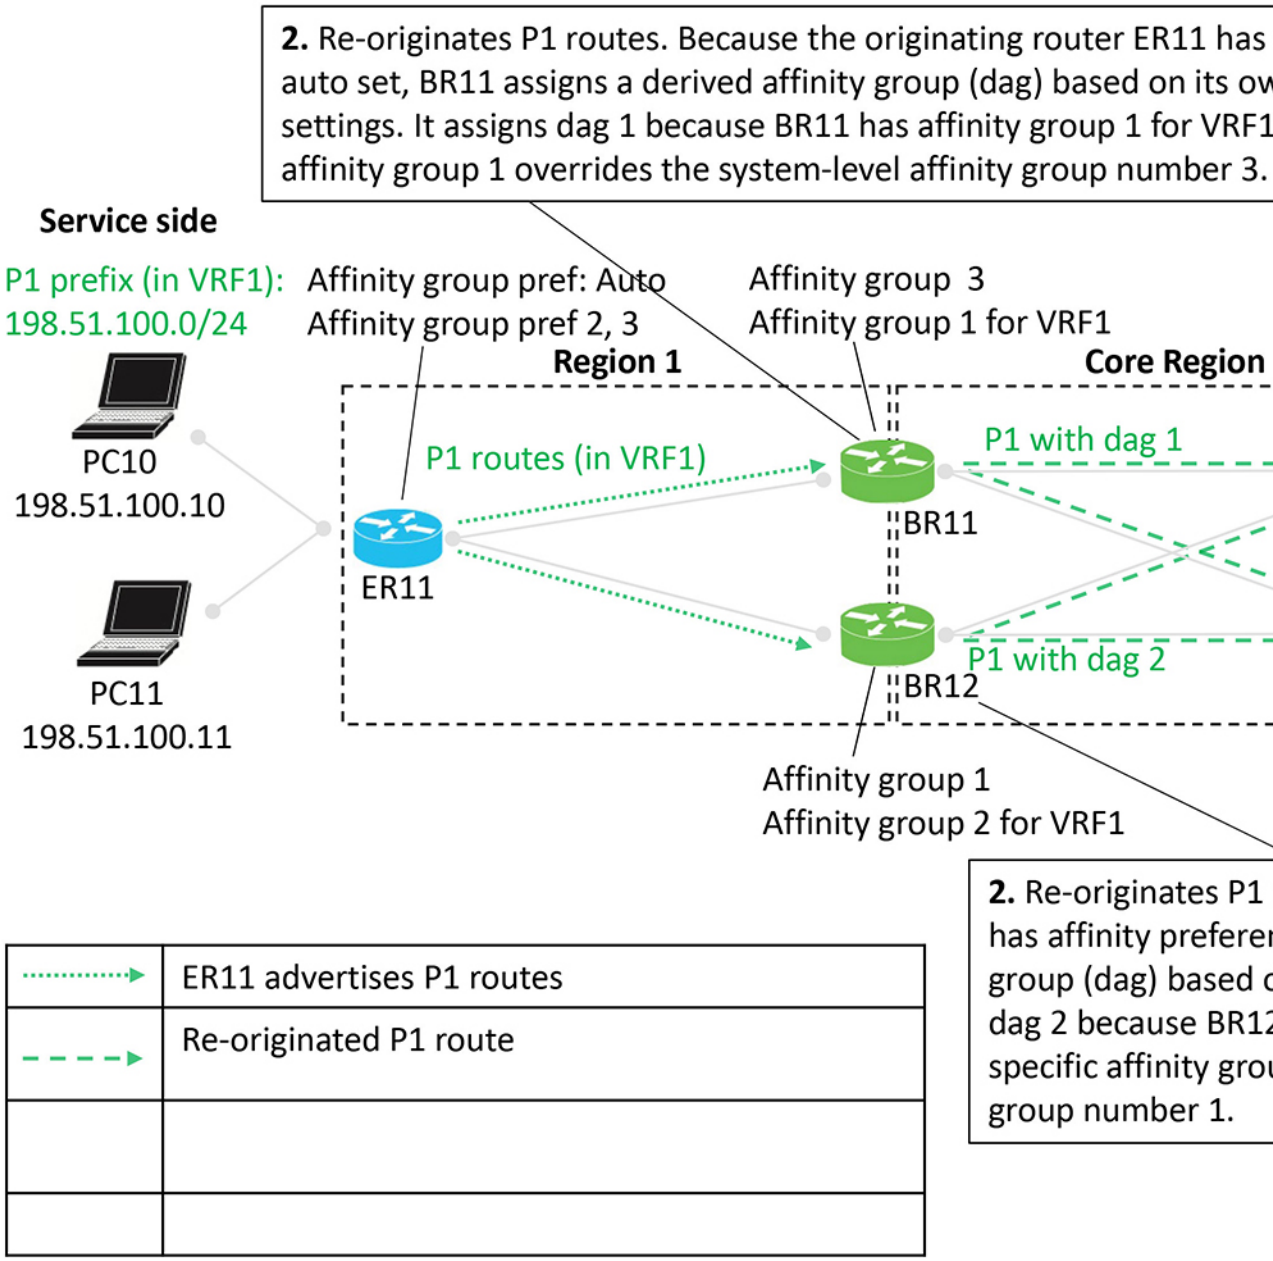

Figure 24: Border routers BR21 and BR22 re-originate the P1 routes

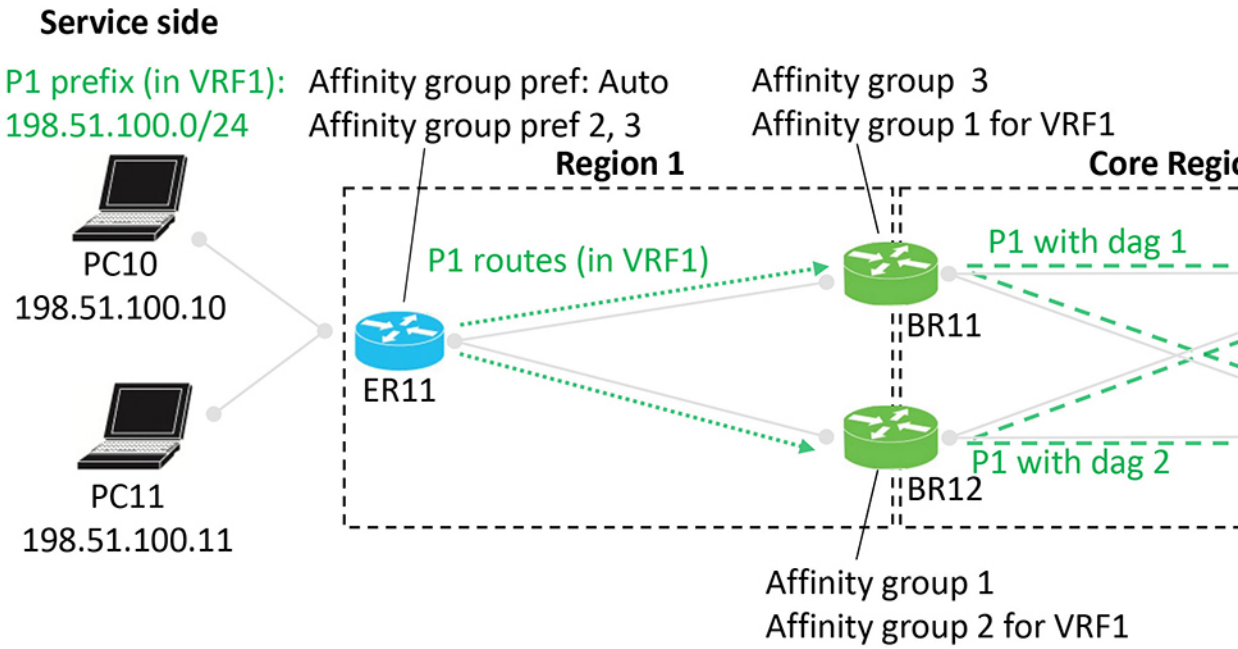

····▶	ER11 advertises P1 routes
- - - -▶	Re-originated P1 route

Figure 25: Route preference according to affinity group and derived affinity group

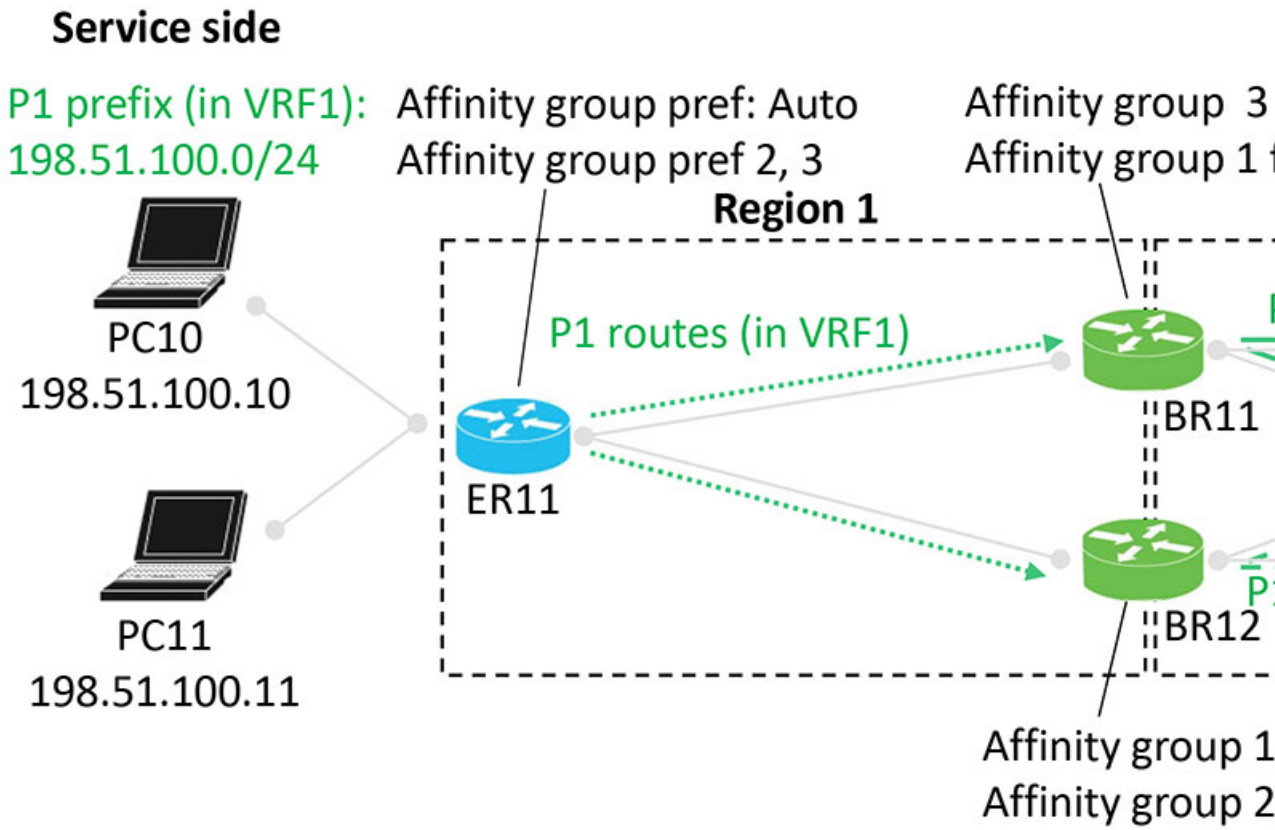

····▶	ER11 advertises P1 routes
- - - -▶	Re-originated P1 route **not preferred** by the receiving device due to higher ag or dag
——▶	Re-originated P1 route **preferred** by the receiving device receiving due to lower ag or dag

Figure 26: Resulting path of traffic to P1

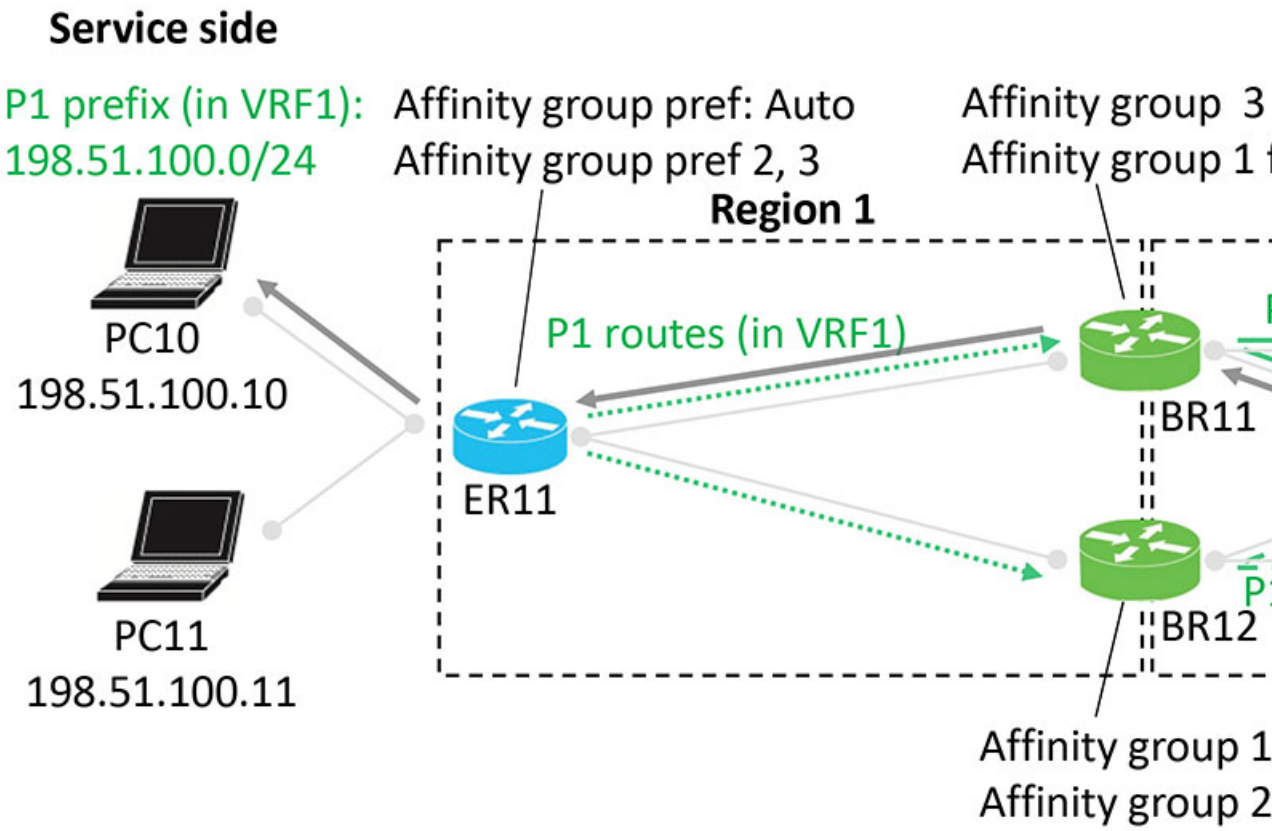

····▶	ER11 advertises P1 routes
- - - -▶	Re-originated P1 route **not preferred** by the receiving device due to higher ag or dag
——▶	Re-originated P1 route **preferred** by the receiving device receiving due to lower ag or dag
——▶	Resulting path between PC20 and PC10

Advertising P2 routes

Edge routers ER21 and ER22 advertise P2 routes. These routes are re-originated from Region 2 back toward Region 1, with border routers again assigning affinity groups and derived affinity groups during the process.

Route preference is determined using the same rules:

- Outside the core region: Affinity group preference
- Inside the core region: Lowest derived affinity group (dag) value

Figure 27: Edge router ER21 advertises P2 routes

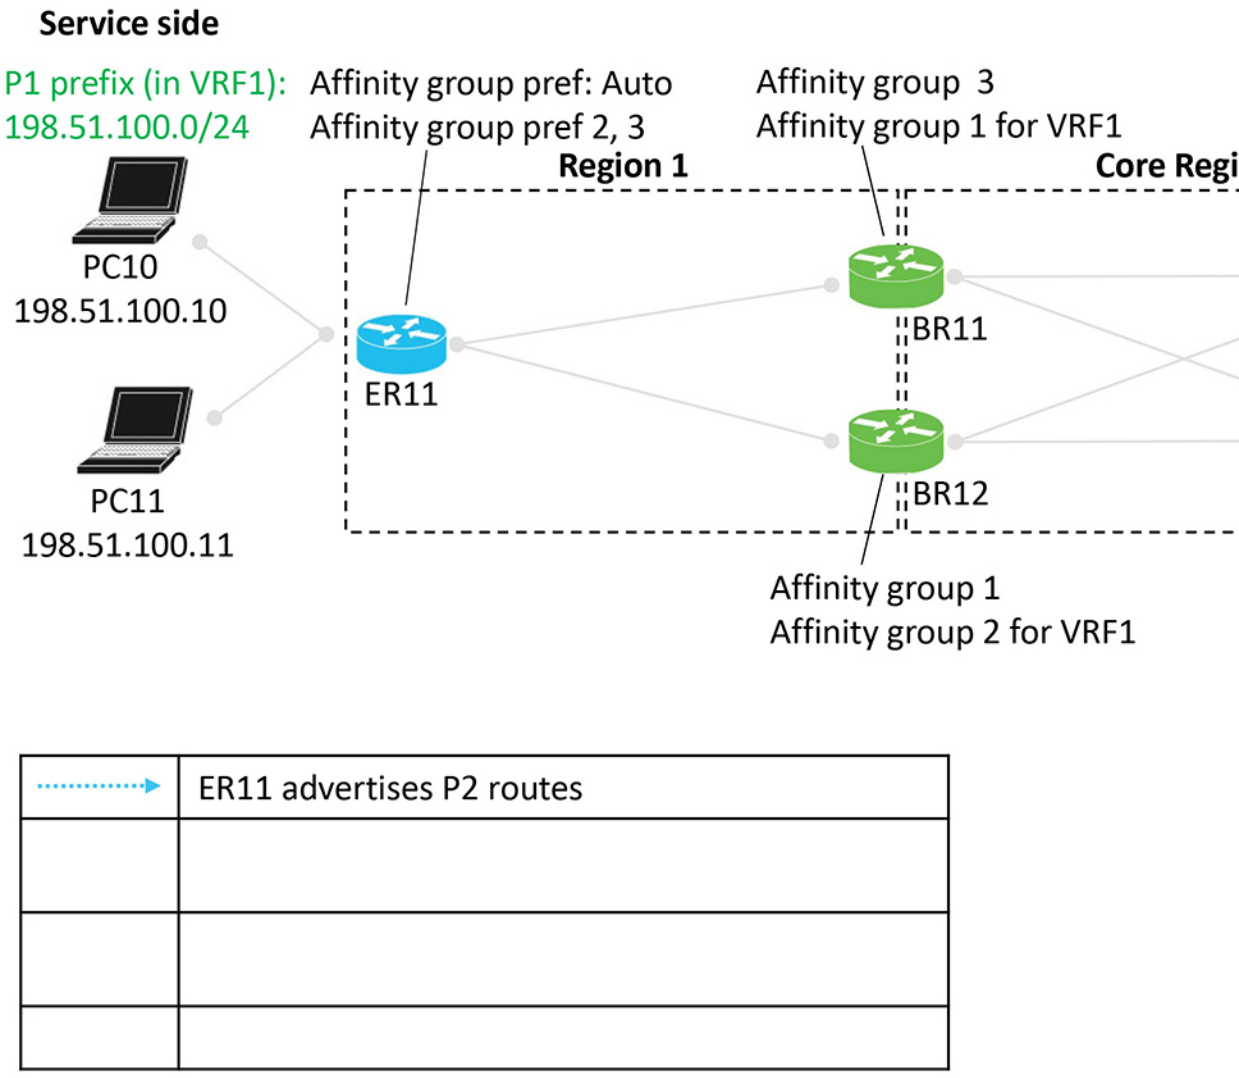

Figure 28: Border routers BR21 and BR22 re-originate the P2 routes

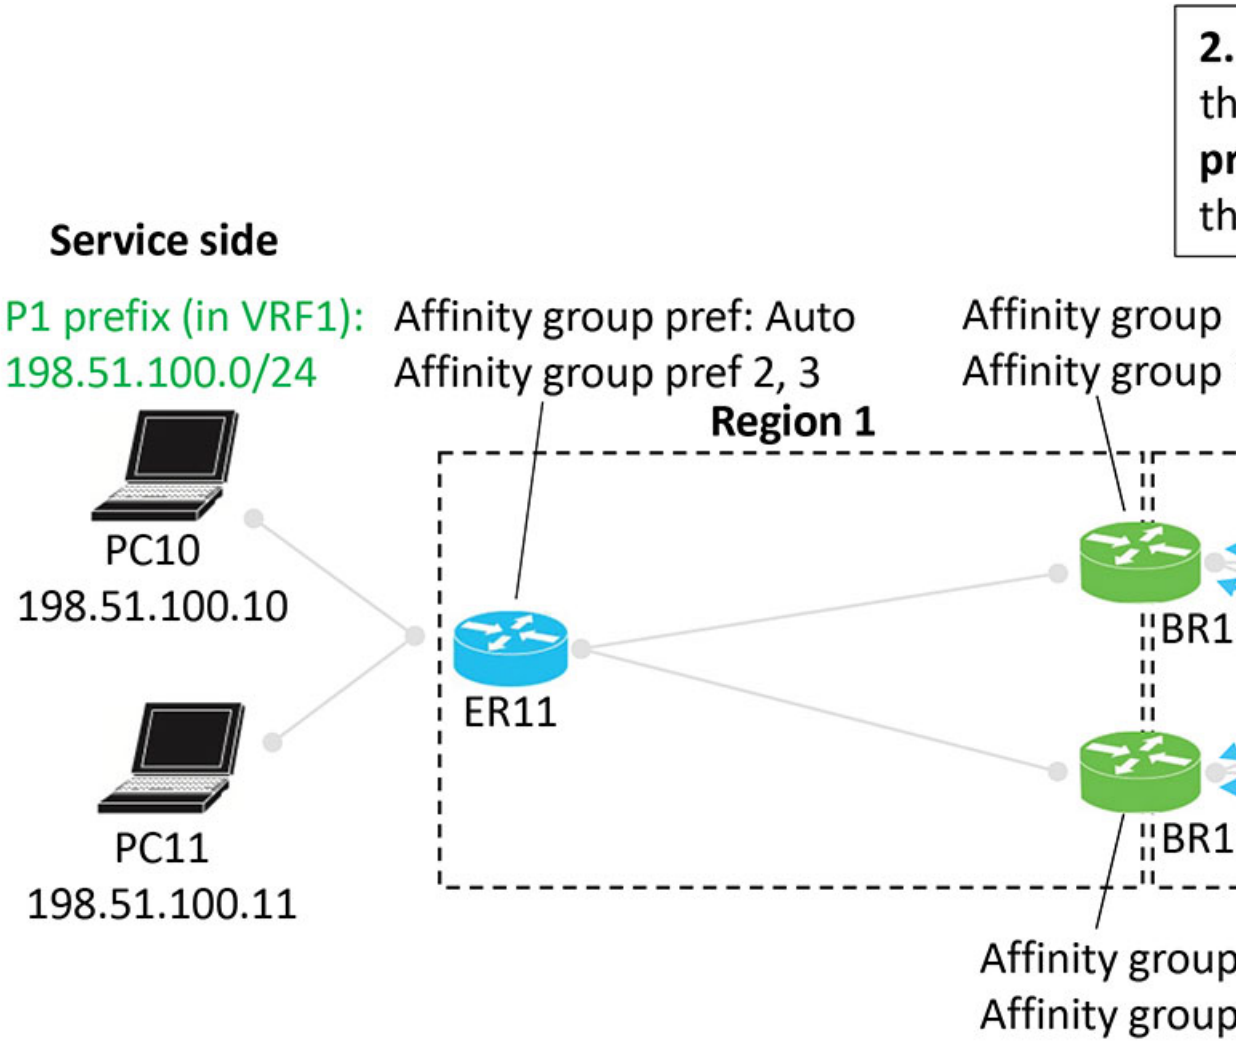

····▶	ER11 advertises P2 routes
- - - -▶	Re-originated P2 route

Figure 29: Border routers BR11 and BR12 re-originate the P2 routes

Figure 30: Route preference according to affinity group and derived affinity group

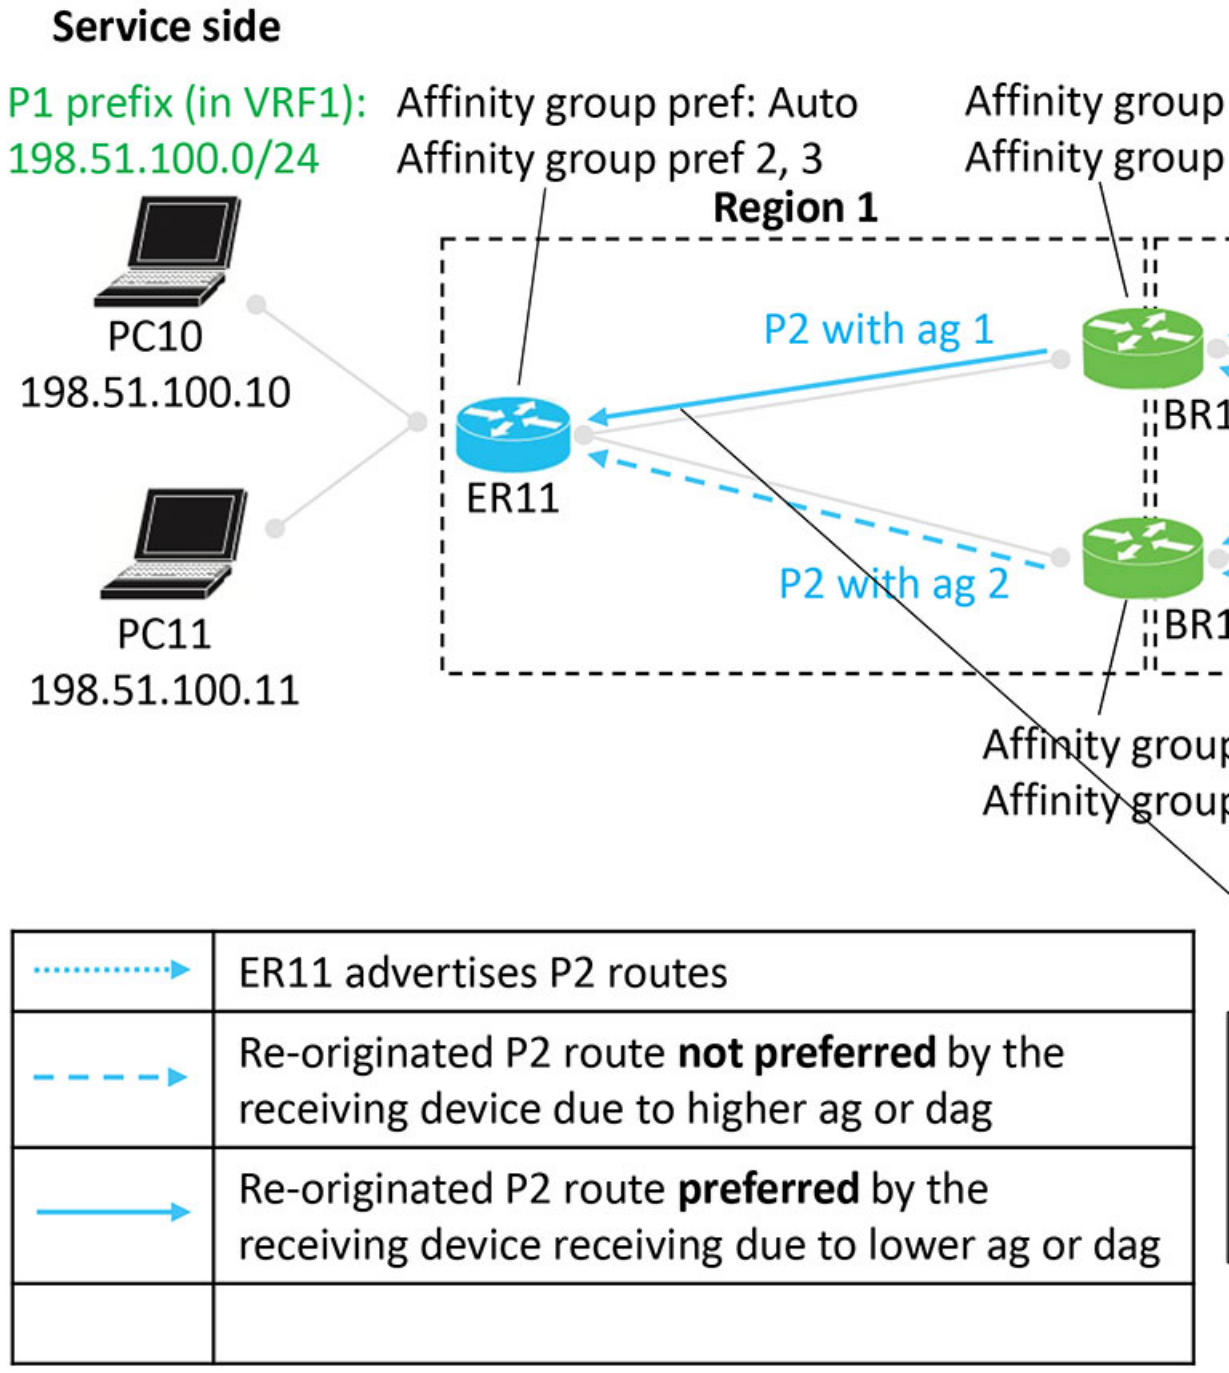

Figure 31: Resulting path of traffic to P2

Result

The following figure shows that the result of the configuration is symmetric routing for flows between, in this example, PC10 and PC20:

Figure 32: Result is symmetric routing

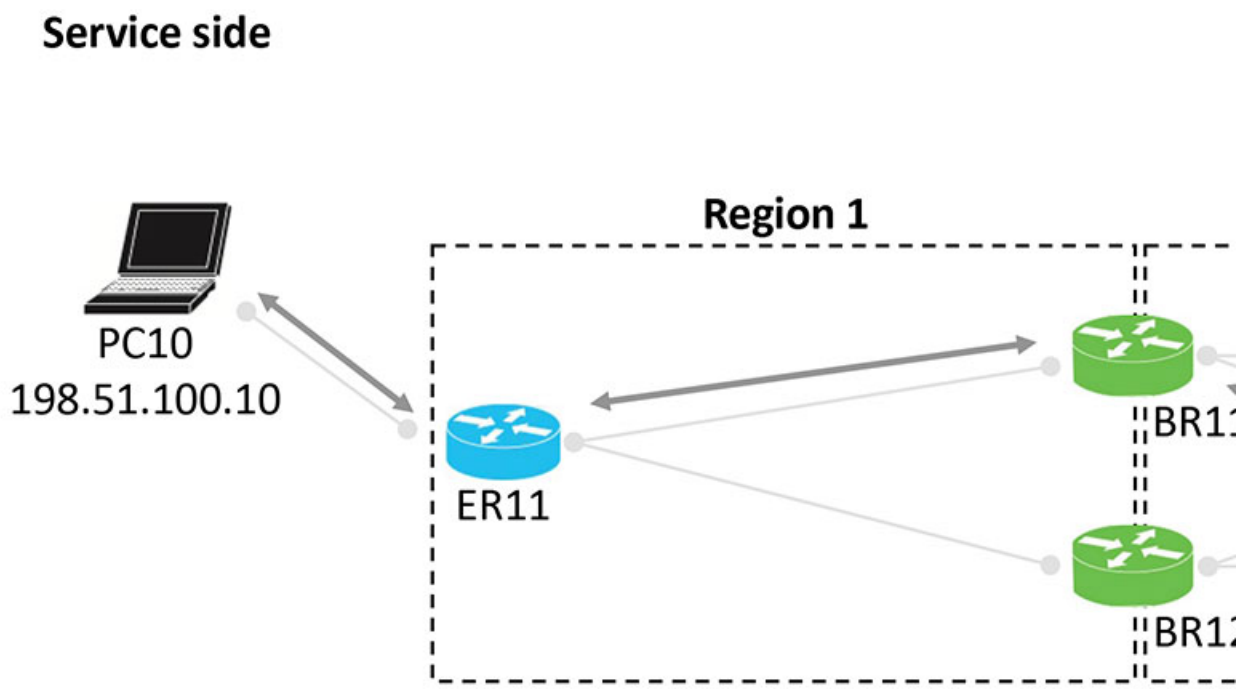

Supported scenarios

The symmetric routing configuration method described in this document applies to the following deployment scenarios:

- Hub-and-spoke topology with multiple hub routers: includes deployments where the hub router provides connectivity to a multi-homed data center.
- Multi-region fabric with multiple border routers: Covers scenarios where an MRF region contains a multi-homed data center, requiring consistent bidirectional path selection across regions.
- Multi-region fabric with transport gateways serving subregions: Applies to MRF deployments where transport gateways (TGWs) connect subregions and influence route propagation.

Scenario: Hub-and-spoke topology, multiple hubs serving a data center, active/active

In this scenario, two hubs serve a data center. The two hubs are both active, for an active/active arrangement.

The data center LAN is not part of the Cisco Catalyst SD-WAN overlay network.

For information about the redistribute omp translate-rib-metric command shown in the illustration see Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template.

Figure 33: Data center, two hubs, active/active

Scenario: Hub-and-spoke topology, multiple hubs serving a data center, active/passive

In this scenario, two hubs serve a data center. Only one hub is typically active, and the other is stand-by, in case the active hub becomes unavailable. This is an active/passive arrangement.

The data center LAN is not part of the Cisco Catalyst SD-WAN overlay network.

Figure 34: Data center, two hubs, active/passive

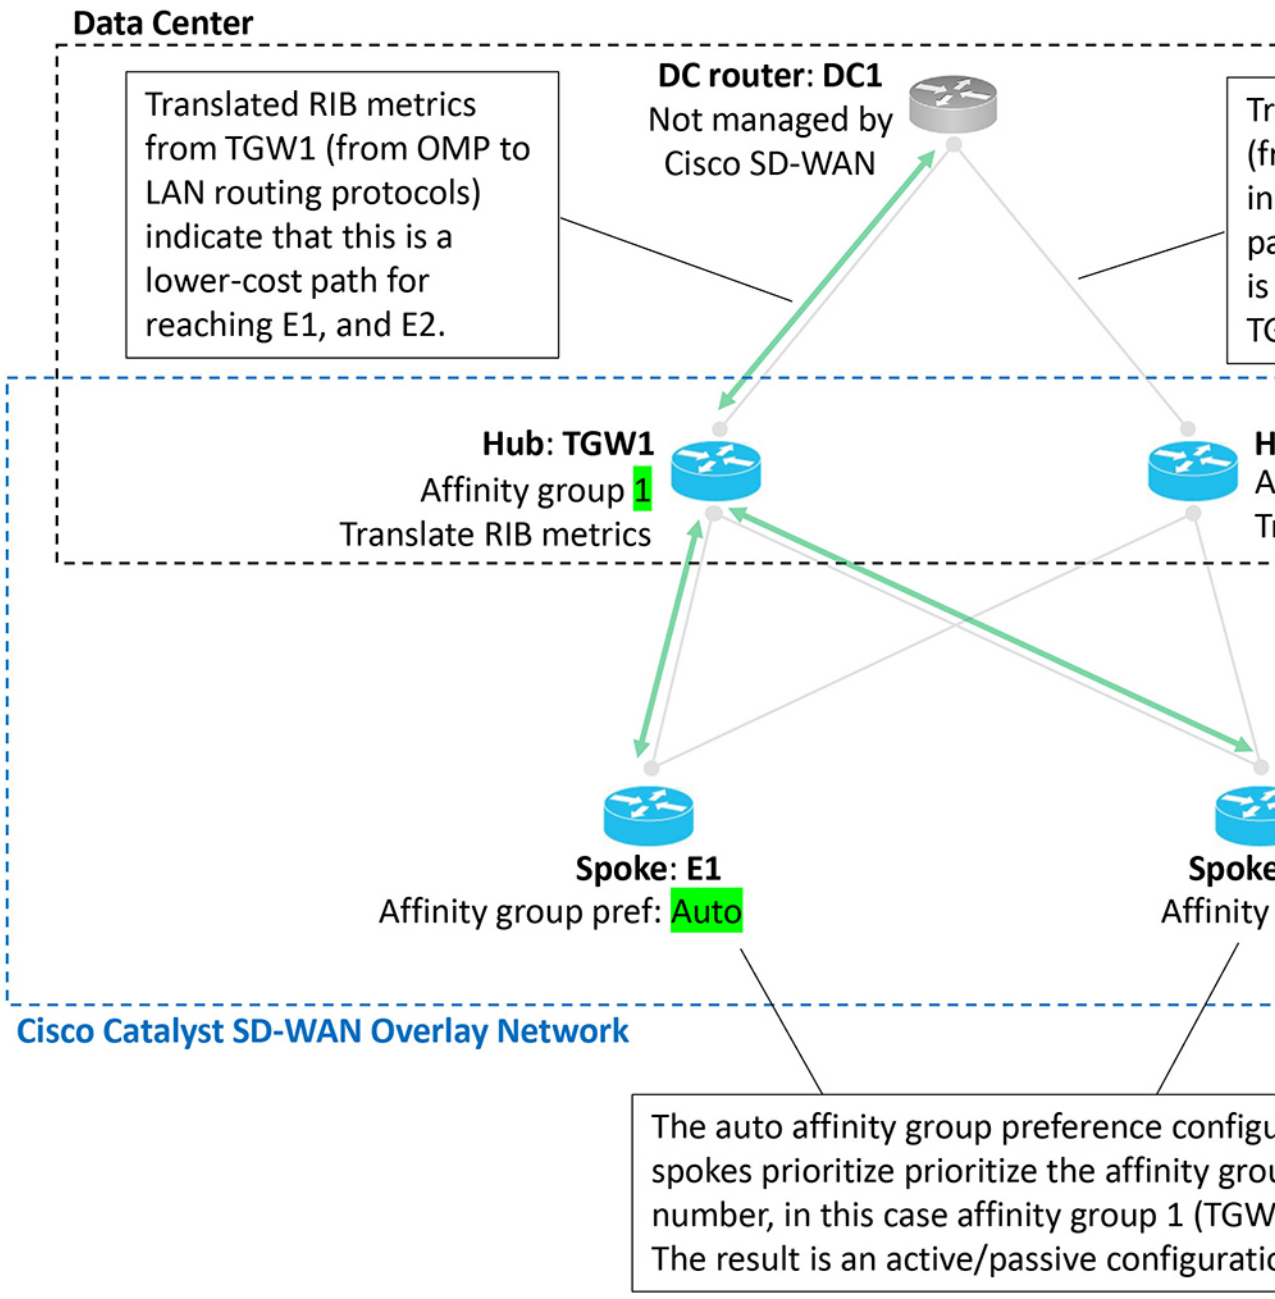

Scenario: Hub-and-spoke topology, multiple hubs serving a data center, active/active by VRF

In this scenario, two hubs serve a data center. The two hubs are both active, for traffic in one of the two VRFs. This is an active/active arrangement, segregated by VRF. The hub TGW1 is active for VRF1 and the hub TGW2 is active for VRF2. Both hubs can operate as stand-by for the other VRF.

The data center LAN is not part of the Cisco Catalyst SD-WAN overlay network.

Figure 35: Data center, two hubs, active/active, segregated by VRF

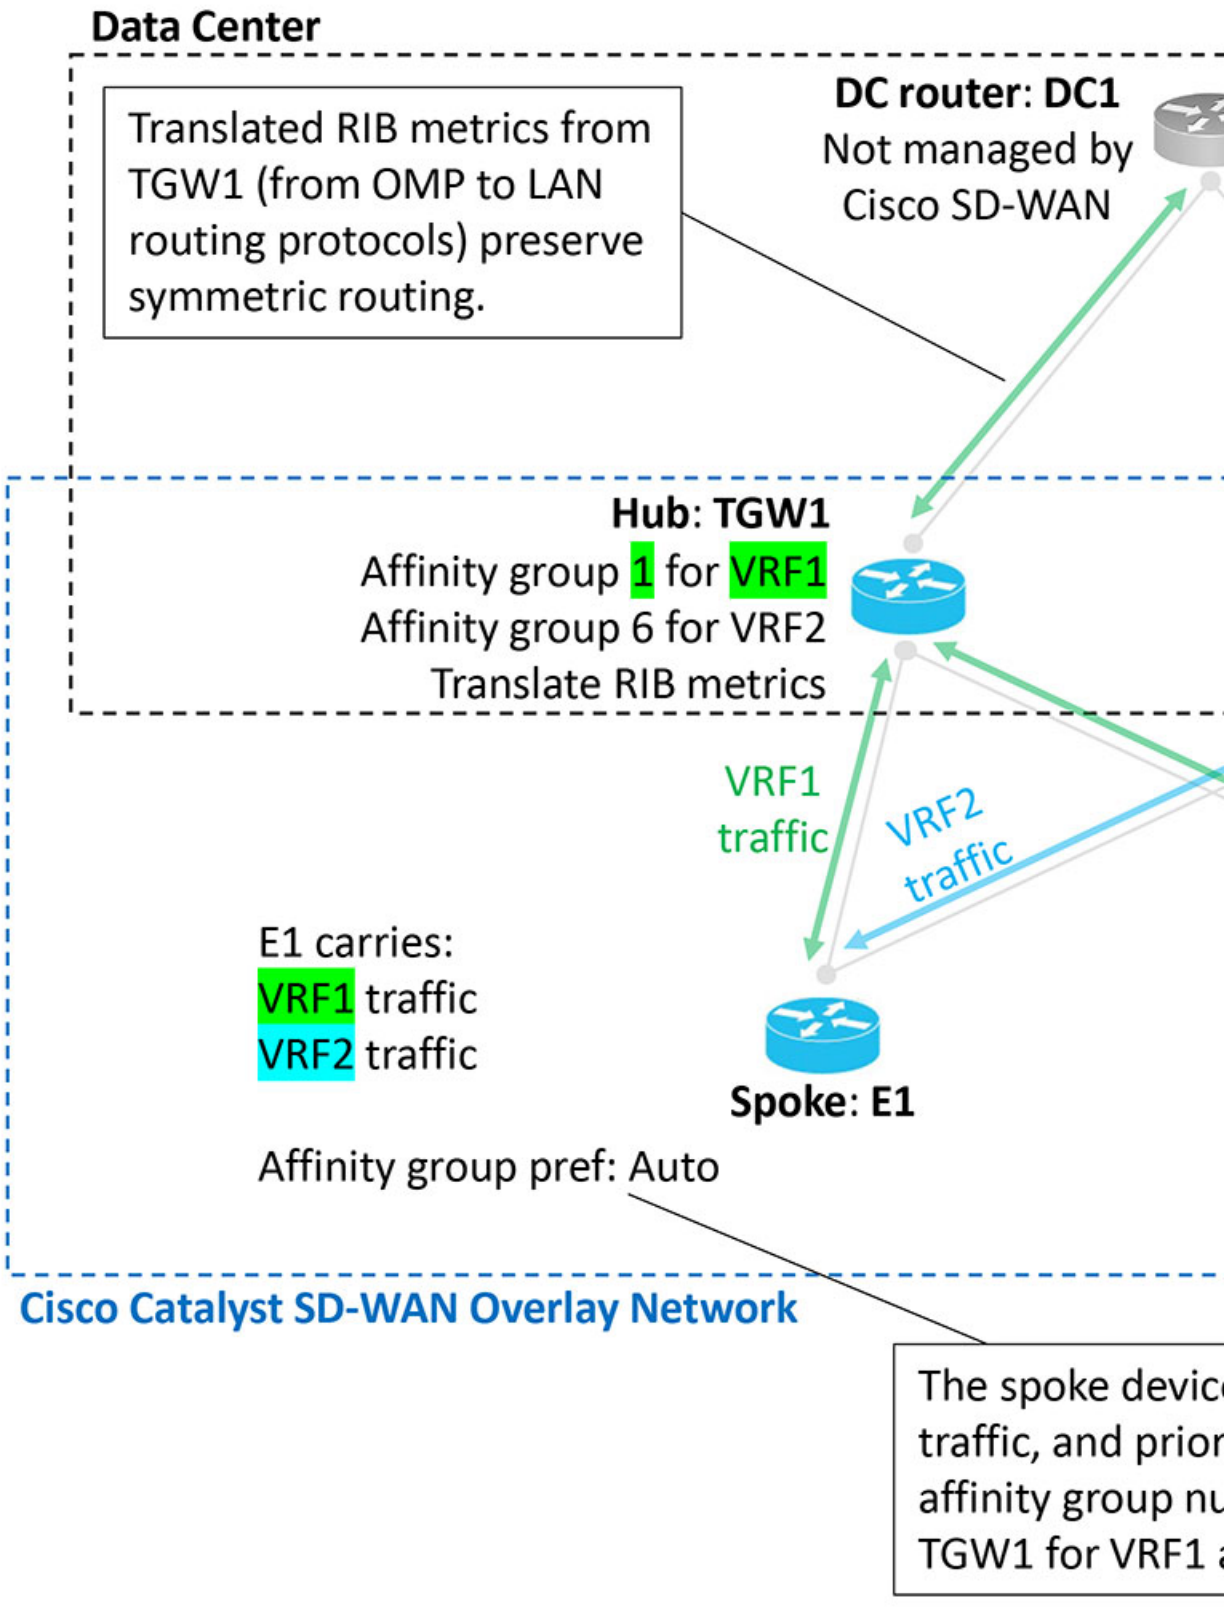
Data Center
Translated RIB metrics from TGW1 (from OMP to LAN routing protocols) preserve symmetric routing.
DC router: DC1
Not managed by Cisco SD-WAN
Hub: TGW1
Affinity group 1 for VRF1
Affinity group 6 for VRF2
Translate RIB metrics
VRF1 traffic
VRF2 traffic
E1 carries:
VRF1 traffic
VRF2 traffic
Spoke: E1
Affinity group pref: Auto
Cisco Catalyst SD-WAN Overlay Network
The spoke devic
traffic, and prior
affinity group nu
TGW1 for VRF1 a

Scenario: Multi-region fabric, transport gateways serving subregions

A Multi-Region Fabric scenario in which transport gateways serve two subregions closely resembles the comprehensive example described in Example of Configuration for Symmetric Routing and the Mechanism.

Similarly to the border routers in the comprehensive example, transport gateways assign a derived affinity group (dag) to routes that they re-originate to other transport gateways. As described in the illustration:

- When transport gateways re-originate routes, they assign derived affinity group (dag) values to the routes.
- Routers choose a preferred route as follows:
  - Between edge routers and transport gateways: According to affinity group preference
  - Between transport gateways in different subregions: According to the lowest derived affinity group value

Figure 36: Multi-region fabric with transport gateways serving subregions

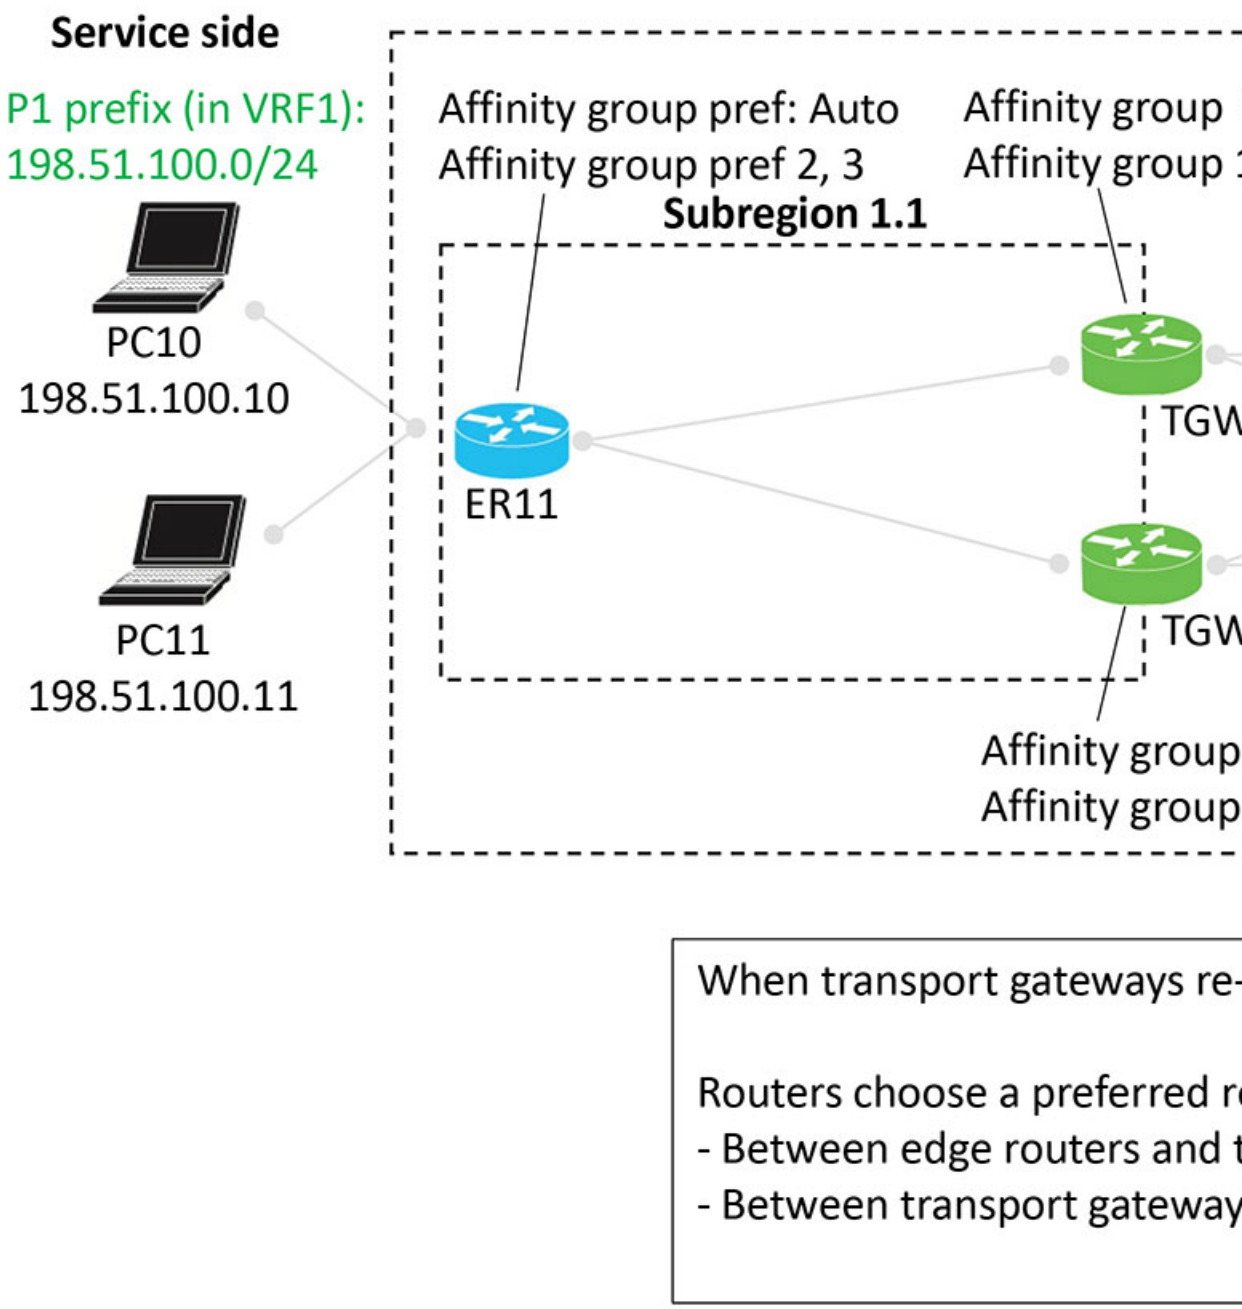

Scenario: Multi-region fabric with route leaking

A Multi-Region Fabric scenario in which transport gateways serve two subregions, with route leaking, closely resembles the comprehensive example described in Example of Configuration for Symmetric Routing and the Mechanism.

Similarly to the border routers in the comprehensive example, transport gateways assign a derived affinity group (dag) to routes that they re-originate to other transport gateways. This scenario is similar to the one described in Scenario: Multi-Region Fabric, Transport Gateways Serving Subregions, but with route leaking. As described in the illustration:

When transport gateways re-originate routes, they assign derived affinity group (DAG) values to those routes.

Routers select a preferred route in the following ways:

- Between edge routers and transport gateways: They use affinity group preference.
- Between transport gateways in different subregions: They choose the route with the lowest derived affinity group value.

In this scenario, a control policy on the Cisco SD-WAN Controllers leaks routes from VRF1 to VRF2 and from VRF2 to VRF1. This route leaking enables endpoints in different VRFs to communicate.

This route-leaking scenario clearly shows how transport gateways (or border routers) assign a derived affinity group (DAG) when they re-originate routes. The logic works subtly, but this example highlights it well.

Default behavior

In this example, the edge routers and transport gateway routers operate as follows:

- ER11 subscribes only to VRF1 and advertises prefix P1 in VRF1.
- ER21 subscribes only to VRF2 and advertises prefix P2 in VRF2.

All transport gateway routers handle traffic for both VRF1 and VRF2, so they re-originate both P1 (in VRF1) and P2 (in VRF2).

By default, the network enforces VRF isolation. When a device advertises routes in different VRFs, the Cisco SD-WAN Controllers filter those routes before sending them to other devices. A controller advertises a VRF x route only to devices that subscribe to VRF x.

Therefore, in this example:

- ER11, which subscribes only to VRF1, does not receive P2 routes from VRF2.
- ER21, which subscribes only to VRF2, does not receive P1 routes from VRF1.

As a result, VRF isolation blocks traffic between ER11 and ER21 because each router subscribes exclusively to a different VRF.

Route leaking

Route leaking allows devices to advertise routes across VRFs by exporting (“leaking”) a route from one VRF into another.

- Source VRF: the route’s original VRF
- Current VRF: the VRF into which the route was exported

When routers advertise exported routes, they track both the source VRF and the current VRF, preserving the background of each route. This tracking becomes important in the logic described later.

In this example, the following route-leaking policies apply:

- An inbound control policy for ER11 instructs it to receive VRF1 routes and export them into VRF2.

Result: ER11 advertises prefix P1 in both VRF1 and VRF2 to its transport gateways, TGW11 and TGW12.

- An inbound control policy for ER21 instructs it to receive VRF2 routes and export them into VRF1.

  Result: ER21 advertises prefix P2 in both VRF2 and VRF1 to its transport gateways, TGW21 and TGW22.

As mentioned earlier, after leaking routes, devices continue to track each route's source VRF and current VRF.

Calculating the derived affinity group (DAG)

A transport gateway device or a border router in a similar scenario assigns a derived affinity group (DAG) to any route it re-originates using the following logic:

1. If the originating router **is** configured with affinity group preference auto (see ER11 in the example), then the re-originating device (for example, TGW11) determines the dag according to its own (TGW11's) affinity group configuration, as follows:
   a. For the leaked route, consider its source VRF and current VRF. Choose the numerically lower of the two values. Call this *x*.
   b. Do one of the following:
      - If the re-originating device only has a system-level affinity group, not VRF-specific affinity groups, then:

        Use the system-level affinity group number for the dag. Assign a dag of that number when re-originating the route.
      - If the re-originating device has a VRF-specific affinity group configured for VRF *x* described in step **a**, then:

        Use this VRF-specific affinity group number for the dag. Assign a dag of this number when re-originating the route.

2. If the originating router **is not** configured with affinity group preference auto (see ER21 in the example), then the re-originating device (for example, TGW21) must consider the affinity preference order configured on the originating device when determining the dag for re-originated routes, as follows:
   a. For the leaked route, consider its source VRF and current VRF. Choose the numerically lower of the two values. Call this *x*.
   b. Do one of the following
      - If the re-originating device only has a system-level affinity group, not VRF-specific affinity groups, then:

        Check the affinity group preference order of the originating device (see ER21). Determine the item number of where the system-level affinity group number occurs in the preference order (item 1, 2, 3, and so on, in the preference order list). Assign a dag of this item number when re-originating the route.

        In the example of TGW21 and ER21, determine where affinity group 2 occurs in the preference order of ER21, which is (1, 2). It is item 2 in the list. So assign a dag of 2 when re-originating the route.
      - If the re-originating device has a VRF-specific affinity group configured for VRF *x* described in step **a**, then:

Using this VRF-specific affinity group, check the affinity group preference order of the originating device. Determine the item number of where the VRF-specific affinity group number occurs in the preference order (item 1, 2, 3, and so on, in the preference order list). Assign a dag of this item number when re-originating the route.

Hypothetically, in the example, if TGW21, in addition to having a system-level affinity group of 2, also had a VRF-specific affinity group of 1 for VRF1, then when TGW21 received from ER21 a P2 route leaked to VRF1, it would consider the preference order of the originating device (ER21). In this hypothetical example with a VRF-specific affinity group of 1, for a route received from ER21, it would check where affinity group 1 occurs in the preference order of ER21, which is (1, 2). It is item 1 in the list. So TGW2 would assign a dag of 1 when re-originating the route.

Example

In the scenario shown in the illustration, a route leaked from VRF2 to VRF1 has a source VRF value of 2 and a current VRF value of 1. When a transport gateway re-originates this route, it assigns a DAG based on the number 1, which is the lower of the two VRF numbers. For example, if TGW12 re-originates a route with a source VRF value of 1 and a current VRF value of 2, it chooses 1 because it is the lower of the two VRF numbers. It therefore calculates the DAG according to VRF1. TGW12 has a system-level affinity group of 1 and a VRF-specific affinity group of 2 for VRF1. Since it calculates the DAG according to VRF1, it assigns the re-originated route a DAG value of 2, taken from the VRF-specific affinity group.

In a hypothetical scenario, if TGW12 had a system-level affinity group of 5 and a VRF1-specific affinity group of 7, then for a route with source VRF 1 and current VRF 2, TGW12 would assign a DAG of 7, taken from the VRF-specific affinity group of 7 for VRF1.

Figure 37: Multi-region fabric with subregions, route leaking

Configure symmetric routing

Use these procedures to configure symmetric routing.

- Configure automatic affinity group preference
- Configure router affinity groups for specific VRFs
- Configure a router to translate OMP metrics to BGP or OSPF
- Configure an affinity group on a router
- Configure Affinity Group Preference on a Router

Configure a router to use automatic affinity group preference

Use any one of these methods to configure a router to use automatic affinity group preference.

- Feature template
- CLI commands

Configure a router to use automatic affinity group preference using templates

If you configure both a manual affinity preference order and an auto preference order on a router, the router gives priority to the auto preference order when selecting the next hop.

However, the manually configured preference list is still useful for path filtering using the **filter route outbound affinity-group preference** command. For information about filtering out paths for routers that are not on the device's affinity list, see Information About Router Affinity Groups and see the filter route outbound affinity-group preference command reference in the Cisco IOS XE SD-WAN Qualified Command Reference.

Procedure

Step 1 From the Cisco SD-WAN Manager menu, choose **ConfigurationTemplates**.

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

- To create a System template for a device, click **Add Template**, choose a device type, and click **Cisco System**.
- To edit an existing System template, locate a System template in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 In the **Affinity Group Preference Auto** field, choose **On**.

Step 5 Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure a router to use automatic affinity group preference, using CLI commands

If you configure a router with both **affinity-group preference-auto** and **affinity-group preference list**, the **affinity-group preference-auto** command has priority for selecting a next hop.

However, the **affinity-group preference list** command is still useful for path filtering using the **filter route outbound affinity-group preference** command.

For information about filtering out paths for routers that are not on the device's affinity list, see Information About Router Affinity Groups and see the filter route outbound affinity-group preference command reference in the Cisco IOS XE SD-WAN Qualified Command Reference.

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates. By default, CLI templates execute commands in global configuration mode.

Procedure

Step 1 Enter system configuration mode.

```
system
```

Step 2 Configure automatic affinity group preference.

```
affinity-group preference-auto
```

Configure router affinity groups for specific VRFs

Use any one of these methods configure router affinity groups for specific VRFs

- Feature template
- CLI commands

Configure router affinity groups for specific VRFs using templates

For more information about using CLI templates, see CLI Add-On Feature Templates and CLI Templates. By default, CLI templates execute commands in global configuration mode.

Procedure

Step 1 Enter system configuration mode.

```
system
```

Step 2 Configure an affinity group to apply to a specific VRF or range of VRFs.

```
affinity-per-vrf affinity-group vrf-range vrf-range
```

The following example configures affinity group 1 for VRF1:

```
system
  affinity-per-vrf 1 vrf-range 1
```

The following example configures affinity group 4 for the VRF range 3 to 6:

```
system
  affinity-per-vrf 4 vrf-range 3-6
```

Configure router affinity groups for specific VRFs using CLI commands

Use these steps to configure router affinity groups for specific VRFs.

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

- To create a System template for a device, click **Add Template**, choose a device type, and click **Cisco System**.
- To edit an existing System template, locate a System template in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 For Affinity Group Number for VRFs, there are two fields. In the left field, enter an affinity group number. In the right field, enter a VRF number or a range of numbers–for example, 2-4. To configure addition group numbers for specific VRFs, click the plus button.

In Cisco SD-WAN Manager, you can configure up to four ranges. If you need to configure more, you can use a CLI template or CLI add-on template. See Configure Router Affinity Groups for Specific VRFs Using a CLI Template.

Step 5 Click **Save** if creating a new template, or **Update** if editing an existing template.

Verify symmetric routing

Use these commands to verify the configurations required for symmetric routing. For more information see, Cisco IOS XE Catalyst SD-WAN Qualified Command Reference

Verify the next hops for a specific prefix on a router

Use **show sdwan omp routes** prefix on a router to show the next hops for a specific prefix.

```
Device# show sdwan omp routes 10.1.1.0/24
```

Verify the path to a destination router

Use **traceroute vrf** *vrf-number destination-ip-address* **numeric** on any device in the network to show the path from the device to a specified destination device, for a specified VRF.

The output shows a list of each hop in the path to the destination device. The last item in the list is the destination device.

```
Device# traceroute vrf 1 10.1.1.1 numeric
Type escape sequence to abort.
Tracing the route to 10.1.1.1
VRF info: (vrf in name/id, vrf out name/id)
  1 209.165.200.225 3 msec 1 msec 1 msec
  2 209.165.200.226 2 msec 1 msec 1 msec
  3 10.1.1.1 4 msec * 4 msec
```

Verify the VRF-specific affinity group configuration on a router

Use **show platform software sdwan rp active internal "omp daemon"** on a transport gateway, or a border router in a Multi-Region Fabric scenario, to show the VRF-specific affinity group configuration on a router. The output shows the affinity group for each configured VRF range.

See the procedures below for configuring VRF-specific affinity groups:

- Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN ManagerConfigure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager
- Configure Router Affinity Groups for Specific VRFs Using a CLI Template Configure Router Affinity Groups for Specific VRFs Using a CLI Template

```
Device# show platform software sdwan rp active internal "omp daemon" | include Affinity
…
Affinity per VRF:

Affinity Group Number: 1 for VRF Range: 1-1
Affinity Group Number: 5 for VRF Range: 2-8
```

Verify a control policy for route leaking

Use **show running-config policy control-policy** on a Cisco SD-WAN Controller to show a control policy that configures route leaking from one VRF to another, if such a policy exists. Exporting routes from one VRF to another is called leaking routes.

For information about configuring a control policy that matches routes of a VRF list and exports the routes to a specific VRF, see Configure Centralized Policies Using the CLI.

Verify control policy application

Use **show running-config apply-policy** on a Cisco SD-WAN Controller to show the sites to which a control policy is applied.

This example shows a control policy that matches VRF1 routes and exports them to VRF2, and matches VRF2 routes and exports them to VRF1.

```
sdwanController# show running-config policy control-policy
policy
 control-policy LEAK_1_TO_2
  sequence 1
   match route
    vpn-list VRF1
   !
   action accept
    export-to
     vpn 2
    !
```

```
   !
  !
  default-action accept
 !
 control-policy LEAK_2_TO_1
  sequence 1
   match route
    vpn-list VRF2
   !
   action accept
    export-to
     vpn 1
    !
   !
  !
  default-action accept
 !
!
Example 2
The following example shows the sites to which the two policies configured in the previous
 example are applied.

sdwanController#show running-config apply-policy
apply-policy
 site-list SL1100
  control-policy LEAK_1_TO_2 in
 !
 site-list SL1300
  control-policy LEAK_2_TO_1 in
 !
!
```

Verify the derived affinity group of a route

Use **show sdwan omp routes***prefix***detail** on a transport gateway, or a border router in a Multi-Region Fabric scenario, to show the derived affinity group assigned to a prefix. The derived-affinity-group parameter in the output shows the value.

In this example the derived affinity group is 2.

```
Device# show sdwan omp routes 192.168.1.0/24 detail
…
  preference       not set
  affinity group  None
  derived-affinity-group  2
  affinity-preference-order  None
  region-id          0
  br-preference      not set
```

Monitor RIB metric translation

For complete information about how a transport gateway translates RIB metrics, see Translating OMP Metrics for Devices Outside of the Overlay Network.

- OMP Metrics
- BGP Metrics
- OSPF Metrics

View OMP metrics

To view the OMP RIB metrics for a route, use the **show ip route** command on a transport gateway that is translating OMP RIB metrics.

The example below shows the OMP RIB metrics for the 10.1.1.1 route. The following metrics are shown in bold in the output:

- OMP Route Metric: 3
- OMP AS-PATH: 100 101

```
Device# show ip route vrf 1 10.1.1.1 protocol-internal
Routing Table: 1
Routing entry for 10.1.1.1/32
  Known via "omp", distance 251, metric 3, type omp
  Redistributing via bgp 1
  Advertised by bgp 1
  Last update from 10.100.1.2 00:04:35 ago
  Routing Descriptor Blocks:
  * 10.100.1.2 (default), from 10.100.1.2, 00:04:35 ago
      opaque_ptr 0x7FC8D1470748
        pdb 0x111111111110, ndb 0x111111111120, rdb 0x111111111130
        OMP attribute 0x7FC8D1470748, ref 2
        aspath 0x7FC8D1474870, ref 2, length 10, value 100 101
        Total OMP attr count 1, aspath 1, community 0
      Route metric is 3, traffic share count is 1
```

OMP route metric for IPv4 routes

To show the OMP route metric for each IPv4 route prefix that a transport gateway is redistributing, use the **show ip route** command on the transport gateway. The OMP route metric, which is 66, is shown in bold in the output, and the administrative distance is 251.

```
device# show ip route vrf 1 omp

Routing Table: 1

      10.0.0.0/32 is subnetted, 1 subnets
m        10.10.10.10 [251/66] via 172.16.0.1, 00:09:15
```

OMP route metric for IPv6 routes

To show the OMP route metric for each IPv6 route prefix that a transport gateway is redistributing, use the **show ipv6 route** command on the transport gateway. The OMP route metric, which is 66, is shown in bold in the output, and the administrative distance is 251.

```
device# show ipv6 route vrf 1 omp
m   2001:DB8::/128 [251/66]
     via 172.16.0.1%default
…
```

View BGP metrics

To view the derived BGP metrics for a route, use the **show ip bgp**command on a transport gateway that is translating OMP RIB metrics.

This example shows the derived BGP metrics for the 10.1.1.1 route. Though the example shows an IPv4 route, IPv6 routes are also supported. The following metrics are shown in bold in the output:

- BGP MED: 3
- BGP LOCAL_PREF: 252
- BGP AS_PATH: 100 100 100 100 101 (This is 100 100 100 (3 copies), plus the original 100 101 of the OMP AS-PATH value.)

```
device# show ip bgp vpnv4 all 10.1.1.1
BGP routing table entry for 1:1:10.1.1.1/32, version 2
Paths: (1 available, best #1, table 1)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  100 100 100 100 101
    10.100.1.2 (via default) from 0.0.0.0 (10.100.1.1)
      Origin incomplete, metric 3, localpref 252, valid, sourced, best
      Extended Community: SoO:0:0
      mpls labels in/out 16/nolabel
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 12 2023 19:08:17 EST
```

View OSPF metrics

To show that the redistribute omp translate-rib-metric command is active on a router, use the **show ip ospf** command. The result shown in bold in the output shows that the router is configured to translate RIB metrics.

- OSPF Metric for IPv4 Routes
- OSPF Metric for IPv6 Routes

OSPF Metric for IPv4 Routes

To show the OSPF metric that the transport gateway uses when distributing IPv4 routes to OSPF, use the **show ip ospf** command on the transport gateway. The OSPF metric, which is determined by the OMP route metric, is 66 in this example, and is shown in bold in the output.

```
Router#show ip ospf 1 rib redistribution
            OSPF Router with ID (192.168.0.1) (Process ID 1)

  Base Topology (MTID 0)

OSPF Redistribution
10.10.10.10/32, type 2, metric 66, tag 0, from OMP_AGENT Router
   via 172.16.0.1, unknown interface
…
```

OSPF Metric for IPv6 Routes

TTo show the OSPF metric that the transport gateway uses when distributing IPv6 routes to OSPF, use the **show ospfv3** command on the transport gateway. The OSPF metric, which is determined by the OMP route metric, is 66 in this example, and is shown in bold in the output.

```
Router#show ospfv3 vrf 1  ipv6 rib redistribution
          OSPFv3 10 address-family ipv6 vrf 1 (router-id 192.168.0.1)

2001:DB8::/128, type 2, metric 66, tag 0, from omp
  via 172.16.0.1
…
```

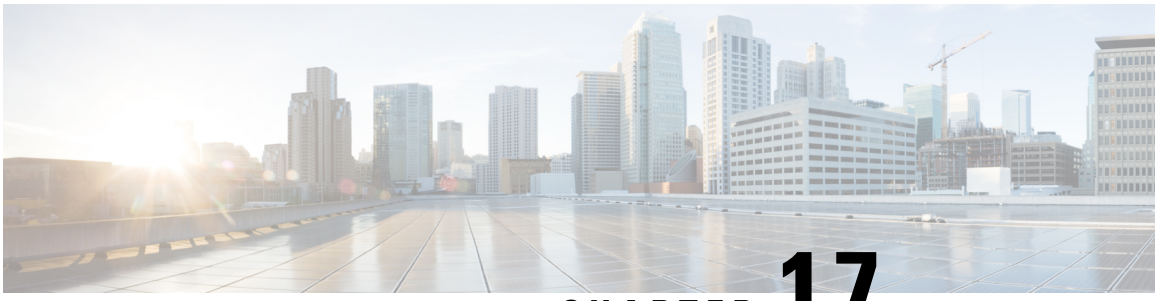

CHAPTER 17

Hub-and-Spoke

Hub-and-spoke topology configuration refers to a simplified method for setting up a hub-and-spoke network, particularly within a Cisco Catalyst SD-WAN. This configuration approach streamlines the process by eliminating the need for complex centralized control policies, which were traditionally required and often lengthy.

Feature history for Hub-and-Spoke

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

Table 49: Feature history

Feature Name	Release Information	Description
Hub-and-Spoke configuration method	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a Cisco Catalyst SD-WAN Manager Release 20.12.1	Hub-and-spoke configuration simplifies the process of configuring a hub-and-spoke topology. This approach makes complex centralized control policy unnecessary. The configuration requires only a few simple configurations: a single command each on: • The Cisco Catalyst SD-WAN Controllers serving a network • A router that serves as a hub • The routers that operate as spokes.

Hub-and-Spoke configuration

Hub-and-spoke configuration is a method that simplifies the process of establishing a hub-and-spoke topology. Historically, this topology required complex expertise and lengthy centralized control policies in a Cisco Catalyst SD-WAN environment. This new configuration method, available from Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, eliminates the need for complex control policy, making configuration faster.

This method involves configuring the Cisco Catalyst SD-WAN Controller that serve the network to enable hub-and-spoke and configuring transport gateway functionality on a router that will serve as a hub.

Note The resulting hub-and-spoke topology applies to all VRFs.

Configuration overview

Hub-and-spoke configuration for Cisco Catalyst SD-WAN has three parts, as described in the following table:

Intent	Devices or Controllers to Configure	Configuration
1. Enable a hub-and-spoke topology in the network.	Cisco SD-WAN Controllers that serve the network	Enable hub-and-spoke configuration in the network. See the following: • Configure a Cisco Catalyst SD-WAN Controller to Enable Hub-and-Spoke Using Cisco SD-WAN Manager. • Configure a Cisco SD-WAN Controller to Enable Hub-and-Spoke Using a CLI Template. The CLI template method uses the **topology hub-and-spoke enable** command.
2. Configure a router as a transport gateway to function as a hub.	Router designated as hub	Enable transport gateway functionality on the router. See Configure a router as a transport gateway using a CLI template, on page 228 The CLI template method uses the **transport-gateway enable** command.
3. Configure routers to function as spokes.	Routers designated as spokes	Configure the device site type as spoke. See Configure the site type for a router using a CLI template, on page 231. The CLI template method uses the **site-type** command.

Result of configuration

This configuration results in the following network behavior:

- Cisco Catalyst SD-WAN Controllers in the network filter the TLOC and route information that they advertise to each router in the network.
  - Routers operating as hubs (transport gateways) receive all TLOC and route information.
  - Routers operating as spokes receive TLOC and route information for the hubs (transport gateways) in the network. They do not receive TLOCs or routes for other spokes. Consequently, there are no bidirectional forwarding detection (BFD) sessions between spoke devices.
- All spoke-to-spoke traffic flows through the transport gateway, which re-originates routes for each spoke.

Taken together, the result is a hub-and-spoke topology. Routers operating as spokes receive TLOC and route information for the hubs (transport gateways) in the network. They do not receive TLOCs or routes for other spokes. Consequently, there are no bidirectional forwarding detection (BFD) sessions between spoke devices.
`

If there are non-spoke sites in the network, spoke sites continue to receive TLOCs or routes from such sites and BFD is established from spoke sites to the non-spoke sites. In this case, it is not a true hub-and-spoke topology.

Benefits of Hub-and-Spoke

A hub-and-spoke topology offers several applications and benefits, including the following:

- Operating each spoke network with a degree of isolation allows for applying different policies, transport mechanisms, and other configurations to each discrete spoke.
- Decreasing the number of peers for the edge routers serving each spoke reduces the resource demands on those edge routers.
- Routing all inter-spoke traffic through a hub enables the application of network services, such as firewall policy, to all inter-spoke traffic.
- The described configuration process simplifies the setup of a hub-and-spoke topology, avoiding the need for complex centralized control policy.

Restrictions for Hub-and-Spoke

When implementing a hub-and-spoke topology, adhere to the following restrictions to ensure proper functionality and avoid misconfigurations.

Key restrictions for hub-and-spoke configurations include:

- Transport gateway site type: When using a transport gateway as a hub, do not configure its site type as spoke.
- On-demand tunnels: In a hub-and-spoke topology, on-demand tunnels are not supported. This is because spoke-to-spoke direct tunnels are not supported in the hub-and-spoke topology.
- Migration: There is no automatic procedure for migrating from a hub-and-spoke topology defined by control policy to the hub-and-spoke configuration method described here.

Hub-and-Spoke connectivity example

This section provides a detailed example demonstrating how network connectivity changes when a full-mesh network is converted to a hub-and-spoke topology.

The following table details the devices, their intended roles, IP addresses, interfaces, and prefixes used in this example, along with their corresponding color coding for illustrations.

Table 50: Devices, IP Addresses, Roles, Interfaces, and Prefixes

Device	Intended Role	Interfaces	Prefixes
Device0 172.16.255.15 Color in illustration: Purple	Hub	10.0.20.15 (3g) 10.1.15.15 (LTE)	None
Device1 172.16.255.35 Color in illustration: Green	Spoke1	10.5.1.35 (LTE)	10.20.35.0/24 Color in illustration: Green highlight
Device2 172.16.255.45 Color in illustration: Blue	Spoke2	10.0.6.45 (LTE)	10.20.45.0/24 Color in illustration: Blue highlight
SDWAN-Controller09 172.16.255.19 Color in illustration: Dark red	Cisco SD-WAN Controller	Not applicable	Not applicable
SDWAN-Controller10 172.16.255.20 Color in illustration: Red	Cisco SD-WAN Controller	Not applicable	Not applicable

The following figure shows the initial state of the network, with full-mesh connectivity before configuring hub-and-spoke.

Figure 38: Network Connectivity Before Hub-and-Spoke Configuration

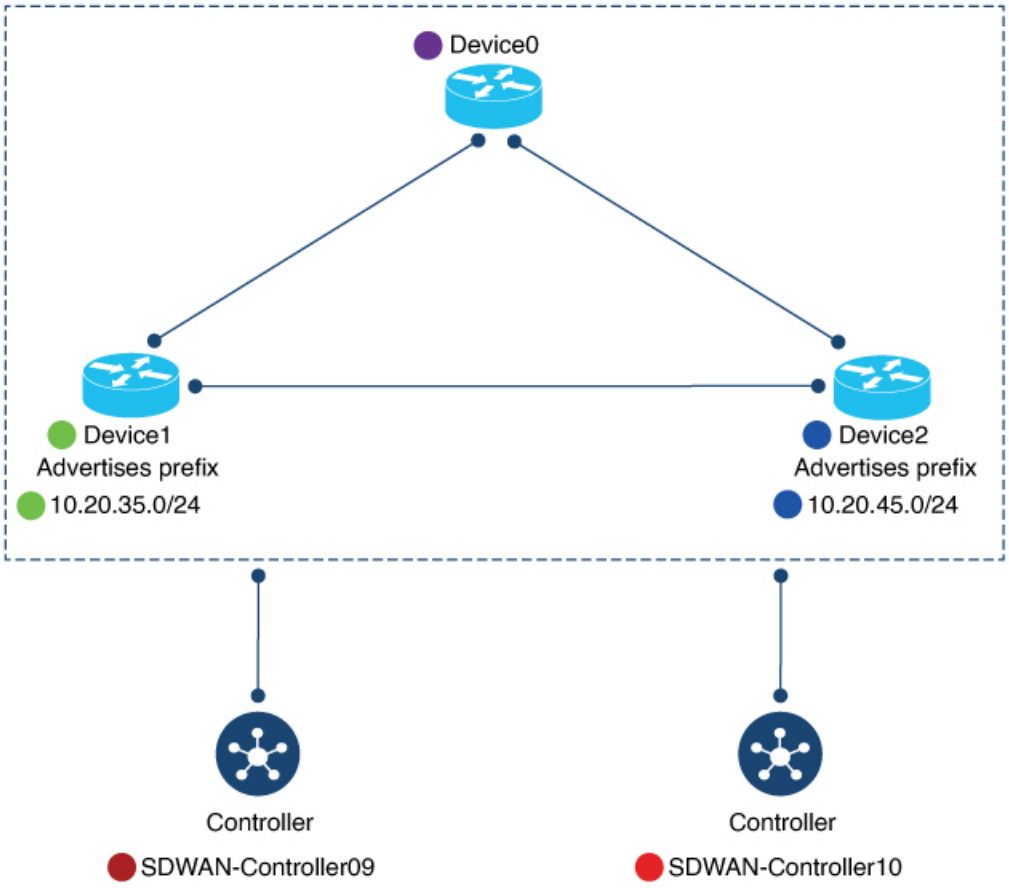

The following figure shows the network connectivity after configuring hub-and-spoke.

Figure 39: Network Connectivity After Hub-and-Spoke Configuration

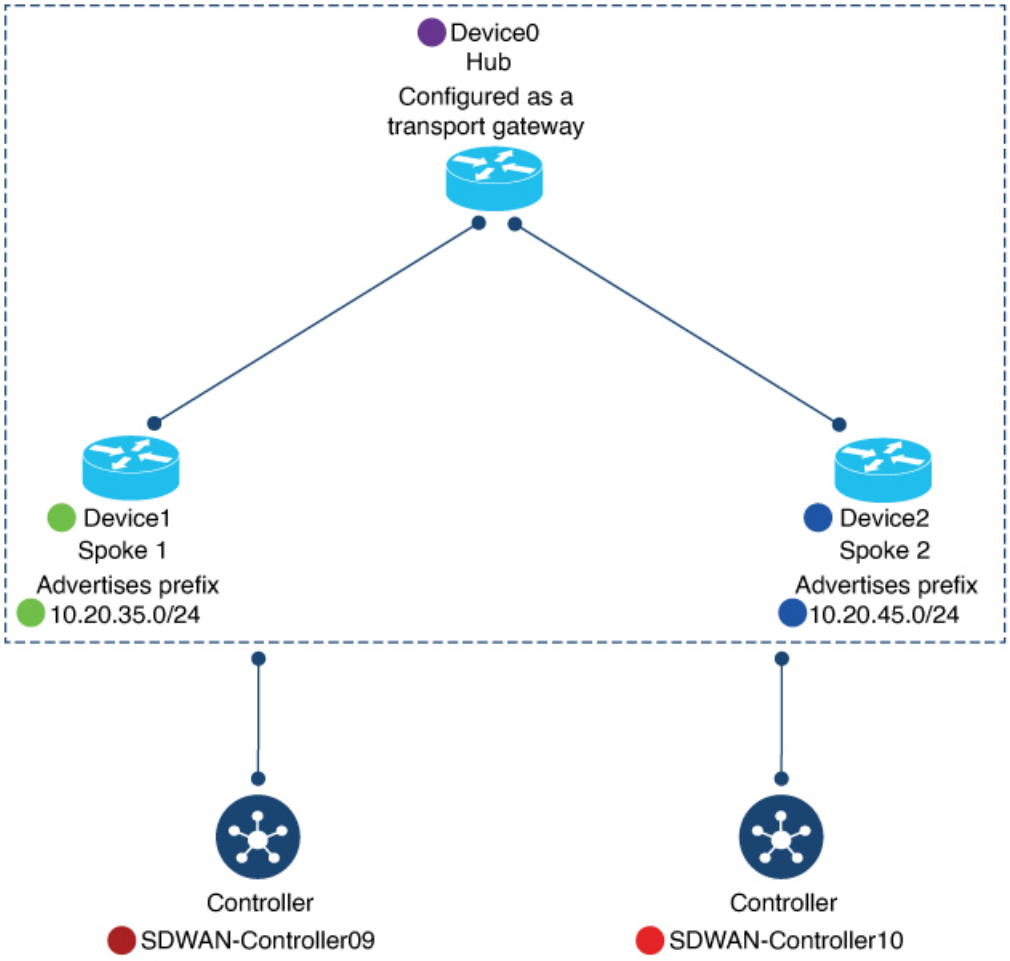

Device0 (Hub) connectivity before and after

This section details the observed connectivity for Device0, which functions as the hub, both before and after the hub-and-spoke configuration. It includes information regarding BFD sessions, OMP routes, and IP routes.

BFD Sessions on Device0 (Hub)

The following describes the state of BFD sessions on Device0.

- Before Configuration: The **show sdwan bfd sessions** command shows that it has BFD sessions with both Device1 (Spoke1) and Device2 (Spoke1).
- After configuration: Device0 retains the same BFD sessions with both Device1 (Spoke1) and Device2 (Spoke2).

Figure 40: Hub: BFD Sessions Before and After

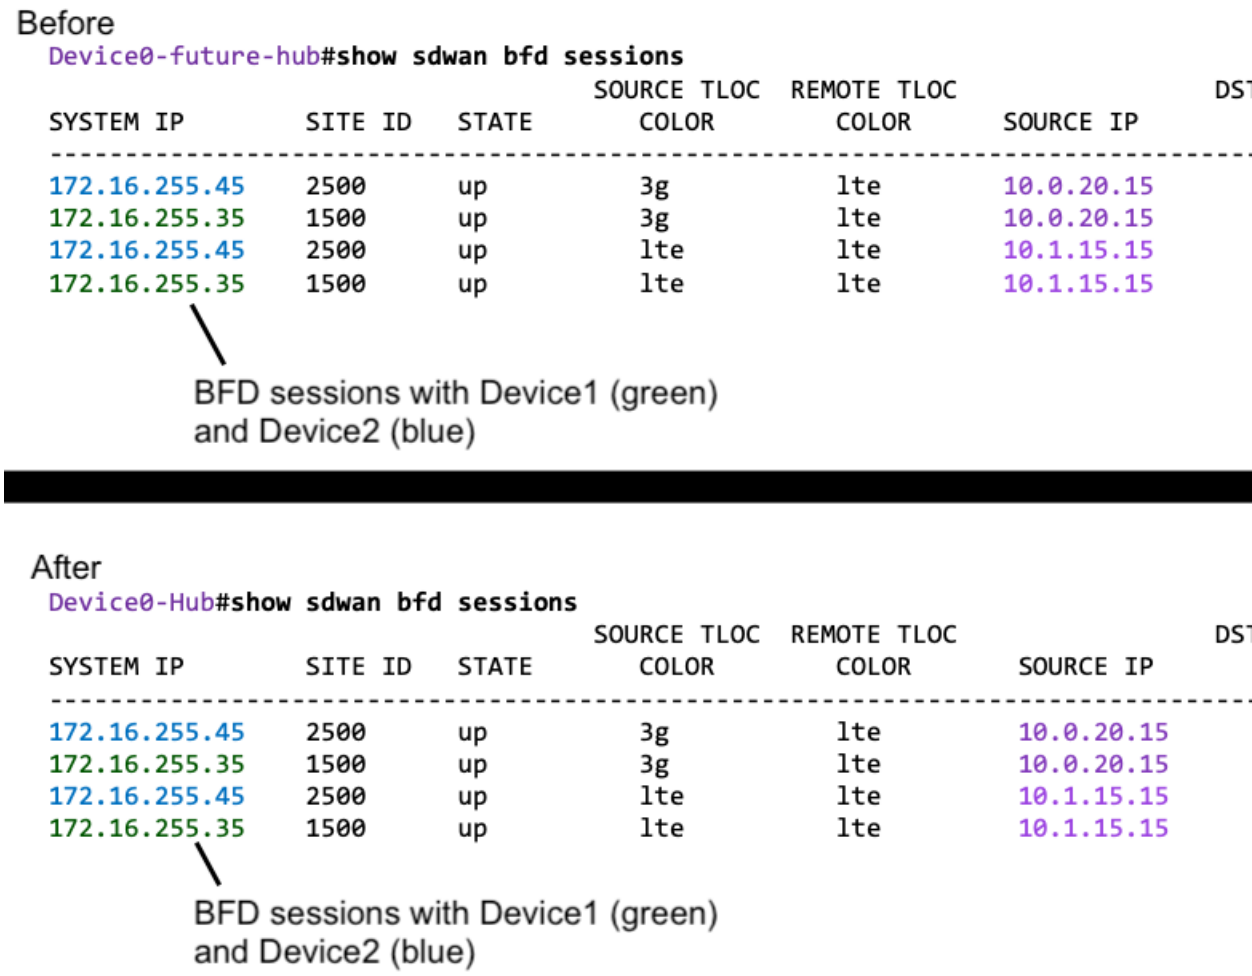

OMP Routes on Device0 (Hub)

The following describes the state of OMP routes on Device0.

- Before Configuration: The **show sdwan omp route vpn 1** command shows that the prefixes advertised by Device1 (Spoke1) and Device2 (Spoke2) are reachable only through Device1 (Spoke1) and Device2 (Spoke2), respectively.
- After configuration: The Device1 (Spoke1) prefix and the Device2 (Spoke2) prefix are reachable through the hub itself (indicated by 0.0.0.0 in the FROM PEER column).

Figure 41: Hub: OMP Routes Before and After

IP Routes on Device0 (Hub)

The following describes the state of IP routes on Device0.

- Before Configuration: The **show ip route vrf 1** command shows that the prefixes advertised by Device1 (Spoke1) and Device2 (Spoke2) are reachable through Device1 (Spoke1) and Device2 (Spoke2), respectively.

- After configuration: This connectivity remains unchanged for Device0.

Figure 42: Hub: IP Routes Before and After

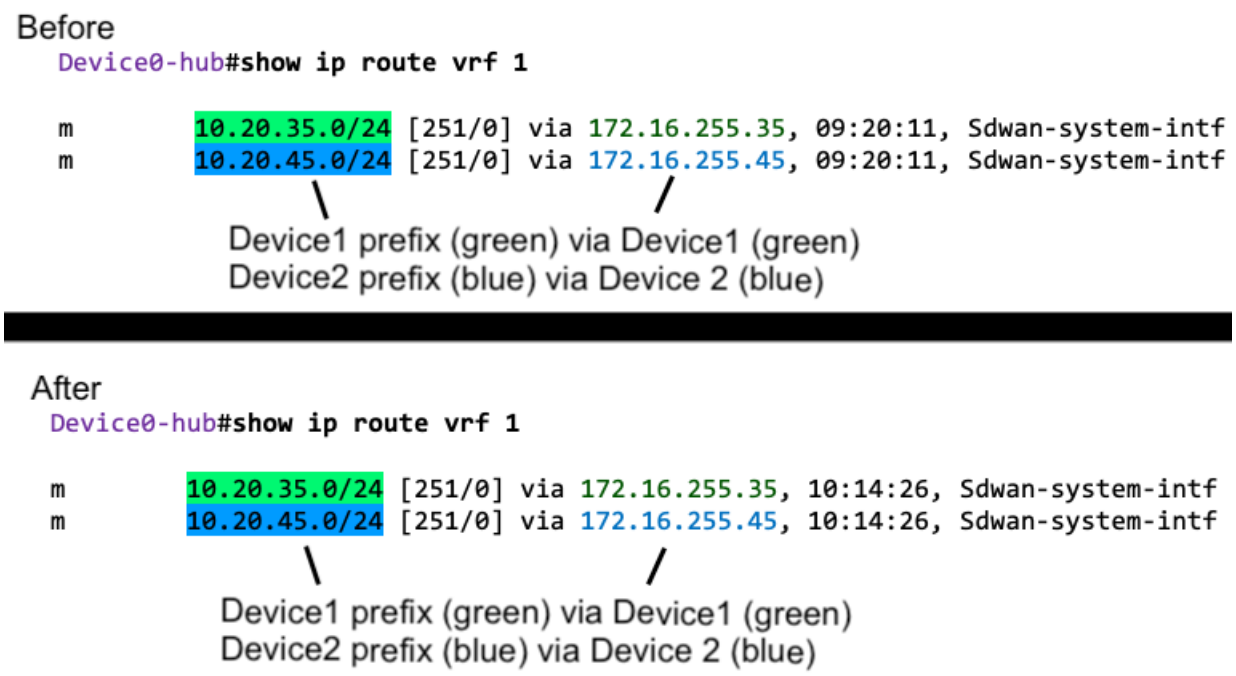

Device1 (Spoke1) connectivity before and after

This section details the observed connectivity for Device1, which functions as Spoke1, both before and after the hub-and-spoke configuration. It includes information regarding BFD sessions, OMP routes, and IP routes.

BFD Sessions on Device1 (Spoke1)

The following describes the state of BFD sessions on Device1.

- Before Configuration: The **show sdwan bfd sessions** command shows BFD sessions with both Device0 (future hub) and Device2 (future Spoke2).
- After configuration: Device1 only has BFD sessions with the hub; there are no BFD sessions with other spokes (for example, Spoke2).

Figure 43: Spoke1: BFD Sessions Before and After

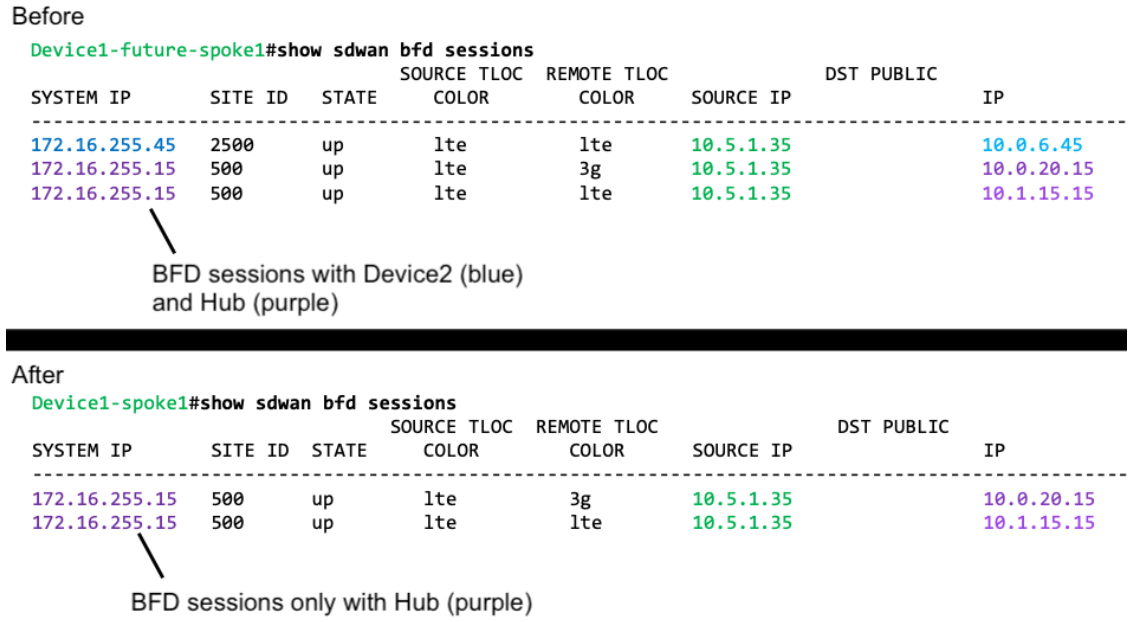

OMP Routes on Device1 (Spoke1)

The following describes the state of OMP routes on Device1.

- Before Configuration: The **show sdwan omp route vpn 1** command shows that it can reach the Device2 (Spoke2) prefix directly through Device2. This is evident because the TLOC IP column shows the system IP of Device2.
- After configuration: Device1 can reach the Device2 (Spoke2) prefix only through the hub.

Figure 44: Spoke1: OMP Routes Before and After

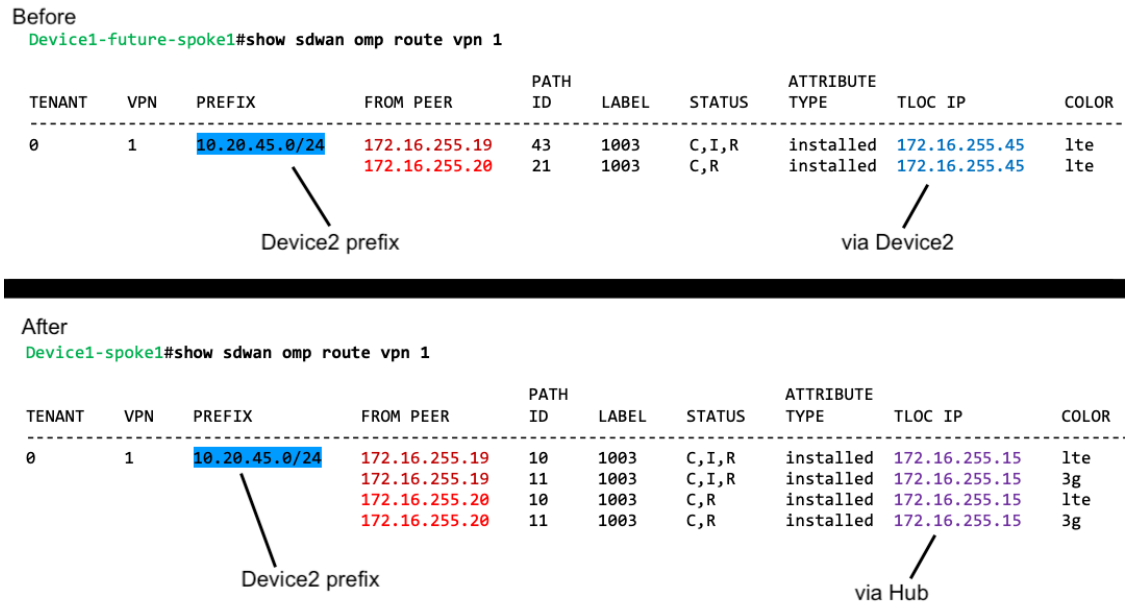

IP Routes on Device1 (Spoke1)

The following describes the state of IP routes on Device1.

- Before Configuration: The **show ip route vrf 1** command shows Device1 could reach the Device2 prefix directly through Device2.
- After configuration: Device1 (Spoke1) can reach the Device2 (Spoke2) prefix only through the hub..

Figure 45: Spoke1: IP Routes Before and After

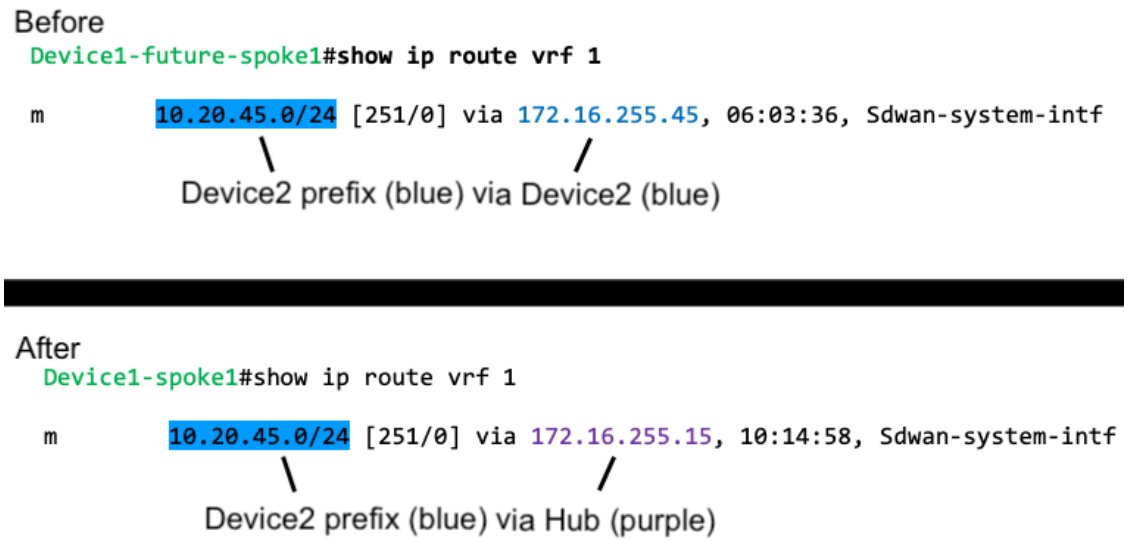

Device2 (Spoke2) connectivity before and after

This section details the observed connectivity for Device2, which functions as Spoke2, both before and after the hub-and-spoke configuration, mirroring the changes observed for Device1. It includes information regarding BFD sessions, OMP routes, and IP routes.

BFD Sessions on Device2 (Spoke2)

The following describes the state of BFD sessions on Device2.

- Before Configuration: The **show sdwan bfd sessions** command shows Device2 had BFD sessions with both Device0 (future hub) and Device1 (future Spoke1).
- After configuration: Device2 only has BFD sessions with the hub; there are no BFD sessions with other spokes (for example, Spoke1).

Figure 46: Spoke2: BFD Sessions Before and After

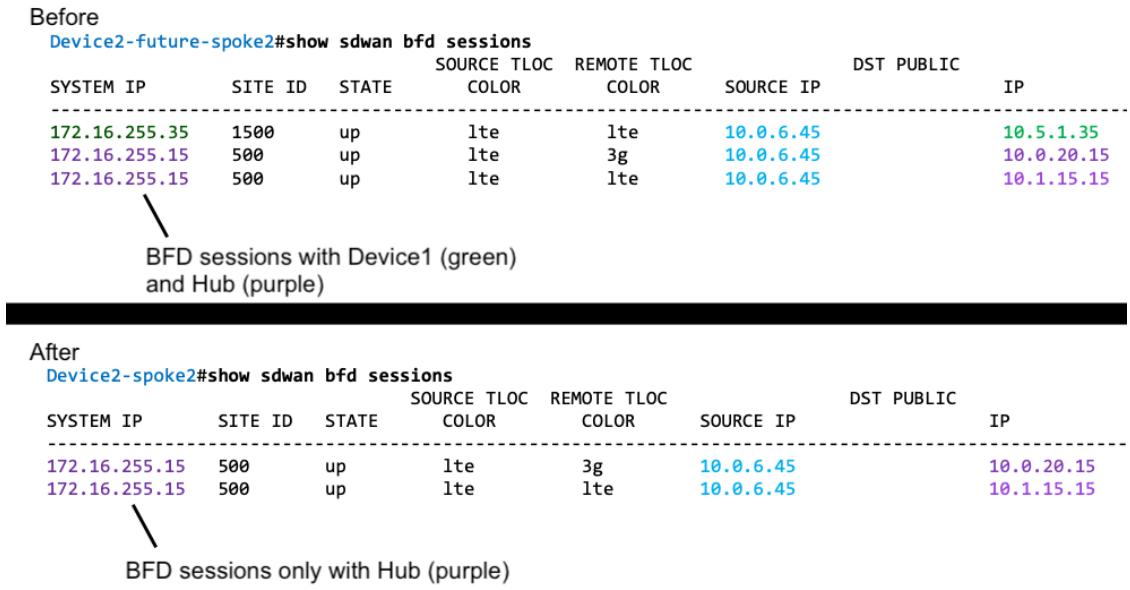

OMP Routes on Device2 (Spoke2)

The following describes the state of OMP routes on Device2.

- Before Configuration: The **show sdwan omp route vpn 1** command shows that Device2 could reach the Device1 (Spoke1) prefix directly through Device1 (TLOC IP column shows the system IP of Device1).
- After configuration: Device2 can reach the Device1 (Spoke1) prefix only through the hub.

Figure 47: Spoke2: OMP Routes Before and After

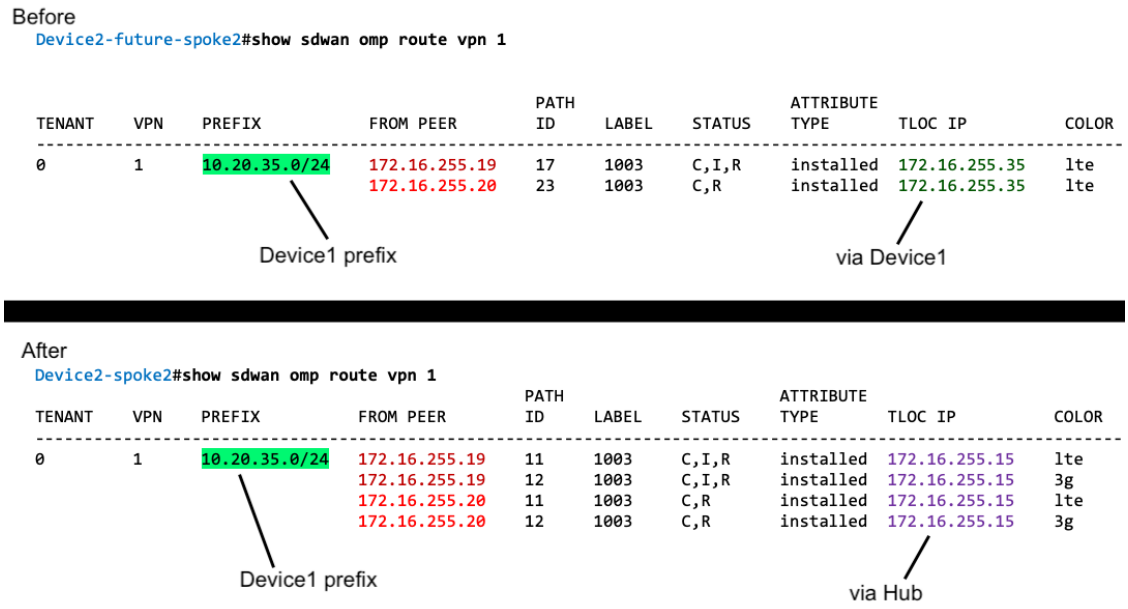

IP Routes on Device2 (Spoke2)

The following describes the state of IP routes on Device2.

- Before Configuration: The **show ip route vrf 1** command shows that Device2 could reach the Device1 prefix directly through Device1.
- After configuration: Device2 can reach the Device1 (Spoke1) prefix only through the hub.

Figure 48: Spoke2: IP Routes Before and After

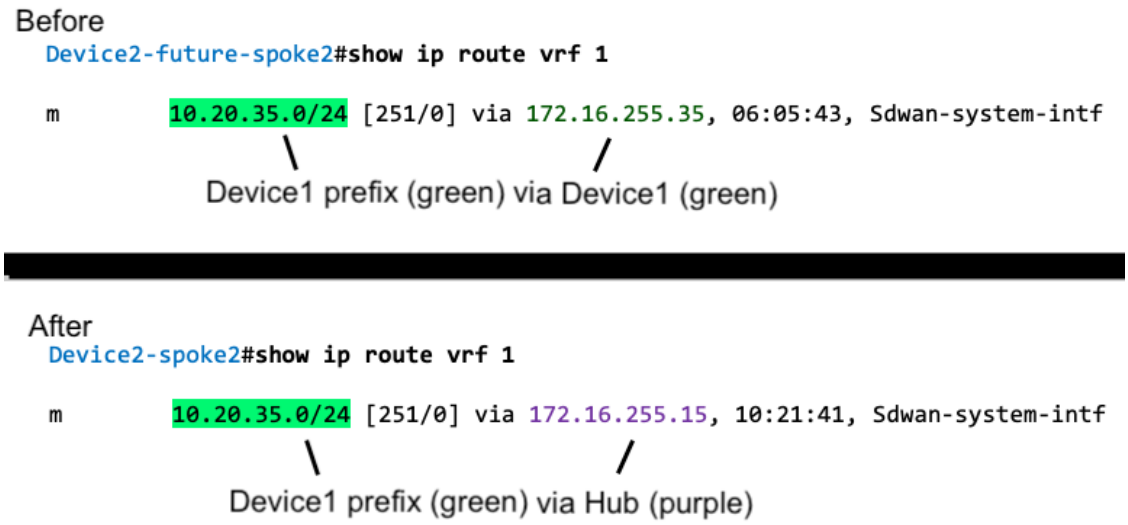

Hub-and-Spoke use cases

Hub-and-spoke use cases describe practical scenarios where this network topology is effectively applied to meet specific organizational needs and leverage its benefits.

Example (Centralized Network Services)

Consider an organization's network with the following characteristics:

- Headquarters Site: Features a single device designated as a hub, running numerous network services such as an enterprise firewall.
- Branch Sites: Consists of three branch locations, each equipped with an edge router.

Network administrators choose to implement a hub-and-spoke topology to route all traffic between branch sites through the headquarters hub. This strategic decision allows for the consistent application of centralized network services to all inter-branch traffic, enhancing security and policy enforcement.

The following illustration depicts the configured hub-and-spoke topology for this use case.

Figure 49: Hub-and-Spoke Topology

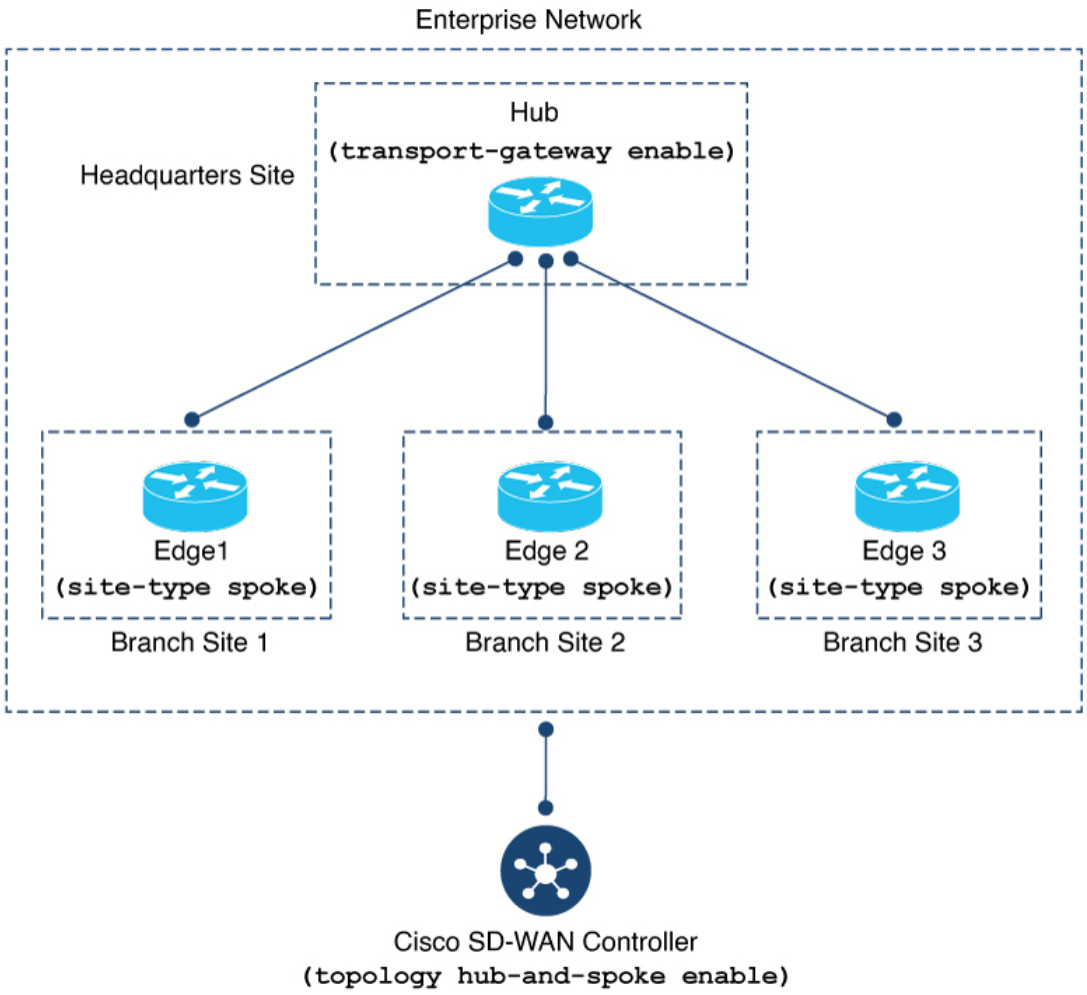

Configure a Hub-and-Spoke topology

To provide a high-level overview and guide users through the complete process of setting up a hub-and-spoke topology using simplified configuration methods.

Before you begin

Follow these steps to configure a hub-and-spoke topology:

Procedure

Step 1 Configure a Cisco SD-WAN Controller to enable Hub-and-Spoke using Cisco SD-WAN Manager. For more information, see Configure a Cisco Catalyst SD-WAN Controller to enable Hub-and-Spoke using Cisco SD-WAN Manager, on page 294.

Step 2 Configure a Cisco SD-WAN Controller to enable Hub-and-Spoke using a CLI Template. For more information, see Configure a Cisco SD-WAN Controller to enable Hub-and-Spoke using a CLI template, on page 295.

Step 3 Configure a Router as a Transport Gateway, for Hub-and-Spoke. Hub-and-spoke configuration makes use of site types and transport gateways. See the following procedures in the transport gateway documentation:

Configure a router as a transport gateway using Cisco SD-WAN Manager, on page 229.

Configure a router as a transport gateway using a CLI template, on page 228.

Step 4 Configure the Site Type for a Router, for Hub-and-Spoke. Hub-and-spoke configuration makes use of site types and transport gateways. See the following procedures in the transport gateway documentation:

Configure the site type for a router using Cisco SD-WAN Manager, on page 232.

Configure the site type for a router using a CLI template, on page 231.

Configure a Cisco Catalyst SD-WAN Controller to enable Hub-and-Spoke using Cisco SD-WAN Manager

Use this procedure to enable the simplified hub-and-spoke topology configuration method on your Cisco SD-WAN Controller via the Cisco SD-WAN Manager graphical user interface.

This configuration is a foundational step for implementing the simplified hub-and-spoke topology in your Cisco Catalyst SD-WAN environment. It prepares the controllers to filter and advertise TLOC and route information appropriately for hub and spoke devices.

Before you begin

Follow these steps to enable hub-and-spoke configuration on your Cisco SD-WAN Controller:

Procedure

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

Step 2 Click **Feature Templates**.

Step 3 Do one of the following:

a) To create a new System template for Cisco SD-WAN Controllers, click **Add Template**, choose **Controller**, and click **System**.

b) To edit an existing Cisco SD-WAN Controller System template, locate a template of type **Controller System** in the table of existing feature templates, click **…** adjacent to the template, and choose **Edit**.

Step 4 In the **Topology** field, choose **Hub and Spoke**.

Step 5 Click **Save** if creating a new template, or **Update** if editing an existing template.

The Cisco SD-WAN Controller is now configured to enable the hub-and-spoke topology, allowing for the simplified configuration of hub and spoke routers in the network.

Configure a Cisco SD-WAN Controller to enable Hub-and-Spoke using a CLI template

Use this procedure to enable the simplified hub-and-spoke topology configuration method on your Cisco SD-WAN Controllers by applying a CLI template.

This method provides an alternative to using Cisco SD-WAN Manager for enabling hub-and-spoke functionality on controllers. It is suitable for automated deployments or when direct CLI configuration is preferred. For more information about using CLI templates, see *CLI Add-On Feature Templates* and *CLI Templates*. By default, CLI templates execute commands in global configuration mode.

Before you begin

Follow these steps to enable hub-and-spoke configuration on your Cisco SD-WAN Controller using a CLI template:

Procedure

Step 1 Enter system configuration mode.

```
system
```

Step 2 Enable a hub-and-spoke topology.

```
topology hub-and-spoke enable
```

Note
To disable hub-and-spoke functionality, use the **no** form of the command.

Example

The following example shows how to enable the hub-and-spoke topology:

```
system
topology hub-and-spoke enable
```

The Cisco SD-WAN Controller is now configured to enable the hub-and-spoke topology via the CLI template, allowing for the simplified configuration of hub and spoke routers in the network.

Hub-and-Spoke configuration verification

Hub-and-spoke configuration verification involves confirming the correct setup and operational status of a hub-and-spoke topology within a Cisco Catalyst SD-WAN environment. This includes examining configurations on Cisco SD-WAN Controllers, individual hub and spoke routers, and observing expected network behavior.

Methods for Verifying Cisco SD-WAN Controller Configuration

Hub-and-spoke configuration makes use of transport gateways and the site type parameter, which are described in the *transport gateway documentation*. See Transport gateways for connecting networks in Cisco SD-WAN, on page 217.

- For information about verifying a transport gateway configuration, see Verify a transport gateway configuration using the CLI, on page 233.
- For information about verifying the site type, see Verify the site type of a router using the CLI, on page 232.
- For information about verifying BFD sessions, OMP routes, and IP routes on the devices in the network after configuring hub-and-spoke, see the example in the introduction to this feature, here: Hub-and-Spoke connectivity example, on page 282.

To verify that a Cisco SD-WAN Controller configuration includes the topology hub-and-spoke enable command, use the show running-config command.

In the following example, the Cisco SD-WAN Controller is configured to enable a hub-and-spoke topology.

```
sdwanController# show running-config
...
system
 topology hub-and-spoke
  enable

To verify that the topology hub-and-spoke enable command has taken effect, use the show omp
 summary command. The output indicates the topology. In the following example, the topology
 is hub-and-spoke.

sdwanController# show omp summary
per-state UP
admin-state UP
...
topology hub-and-spoke
```