



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

First Published: 2020-04-30

Last Modified: 2023-12-19

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/c/en/us/about/legal/trademarks.html>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2019–2023 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	Read Me First	1
------------------	----------------------	----------

CHAPTER 2	What's New in Cisco IOS XE (SD-WAN)	3
------------------	--	----------

CHAPTER 3	Unicast Overlay Routing	5
	Supported Protocols	5
	OMP Routing Protocol	5
	OMP Route Advertisements	6
	OMP Route Advertisements for Cisco Catalyst SD-WAN Controllers	11
	OMP Route Redistribution	11
	Administrative Distance	14
	OMP Best-Path Algorithm	15
	OMP Graceful Restart	18
	BGP and OSPF Routing Protocols	19
	OSPFv3	20
	EIGRP	21
	Routing Information Protocol (RIP)	22
	Information About Routing Information Protocol Support	22
	Prerequisites for Using Routing Information Protocol	24
	Restrictions for Using Routing Information Protocol	25
	Configure Unicast Overlay Routing	25
	Configure BGP	25
	Configure BGP Using CLI	34
	Configure OSPF	39
	Configure OSPF Using CLI	45
	Configure OMP	46

Configure OMP Using CLI	50
Configure OSPFv3	55
Configure OSPFv3 Using CLI	60
Configure EIGRP	62
Configure EIGRP Using CLI	65
Verify EIGRP Configuration Using CLI	66
Configure Routing Information Protocol (RIPv2) Using the CLI	67
Verify RIPv2 Configurations Using the CLI	70
Configure RIPv2 Using the CLI	71
Configuration Example for RIPv2	73
Verify RIPv2 Configurations Using the CLI	74

CHAPTER 4**Multicast Overlay Routing 75**

Multicast Overlay Routing	75
Supported Protocols	77
PIM	77
IGMP	78
MSDP	78
Traffic Flow in Multicast Overlay Routing	79
Configure Multicast Overlay Routing	80
Configure Multicast	81
Configure Multicast Using Configuration Groups	81
Configure Multicast Using the CLI	86
Configure an ACL for Multicast Using a CLI Add-On Template	86
Configure PIM	87
Rendezvous Point Selection Process by a PIM BSR	90
Sample Topology for RP Selection by a PIM BSR	91
Configure a PIM BSR	92
CLI Configurations for PIM BSR Selection	95
Verify VRRP-Aware PIM Using the CLI	97
Configure IGMP	98
Configure PIM and IGMP Using the CLI	100
Configure MSDP Using a CLI Template	101
Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN	102

Information About Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN	102
Benefits of Support for MSDP to Interconnect Cisco SD-WAN and non-SD-WAN	104
Prerequisites for Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN	104
Restrictions for Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN	104
Configure MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN	105
Configure MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN Using a CLI Template	105
Verify MSDP Configuration to Interconnect Cisco SD-WAN and Non-SD-WAN	106
Monitor MSDP Configuration to Interconnect Cisco SD-WAN and Non-SD-WAN	106
Troubleshooting	107
MSDP SA Cache Not Populated	107
OMP SA Route Not Advertised	107

CHAPTER 5**Radio Aware Routing 109**

Supported Devices for RAR	110
Prerequisites for RAR	110
Benefits of RAR	110
Restrictions for RAR	110
Information about RAR	111
Overview of RAR	111
System Components of RAR	113
Configure RAR	114
Configure the RAR Feature in Bypass Mode	115
Configure the RAR Feature in Aggregate Mode	117

CHAPTER 6**Route Leaking Between VPNs 119**

Supported Protocols	120
Restrictions for Route Leaking and Redistribution	121
Information About Route Leaking	122
Use Cases for Route Leaking	123
How Route Preference is Determined	124
Workflow to Configure Route Leaking Using Cisco SD-WAN Manager	124
Configure Localized Route Policy	124
Configure and Enable Route Leaking between Global and Service VPNs	126
Configure Route Leaking Between Service VPNs	128

Attach the Service Side VPN Feature Template to the Device Template	129
Configure and Verify Route Leaking Using the CLI	129
Configure Route Redistribution Between Global VRF and Service VPNs Using the CLI	135
Verify Route Redistribution	138
Configure Route Leaking Between Service VPNs Using a CLI Template	139
Verify Route-Leaking Configurations Between Service VPNs Using the CLI	140
Configure VRRP Tracker for Tracking Leaked Service VPNs Using the CLI	141
Verify VRRP Tracking	143
Configuration Example for Route Leaking	144

CHAPTER 7**BFD for Routing Protocols in Cisco Catalyst SD-WAN 147**

Information About BFD for Routing Protocols	148
Overview of BFD	148
How BFD Works in Cisco Catalyst SD-WAN	148
Supported Protocols and Interfaces	150
Limitations and Restrictions	151
Configure BFD for Routing Protocols	151
Enable BFD for Routing Protocols	152
Configure BFD for Service-Side BGP	152
Configure BFD for Transport-Side BGP	153
Configure BFD for Service-Side EIGRP	154
Configure BFD for Service-Side OSPF and OSPFv3	155
Attach Feature Template to Device Template	156
Configure BFD for Routing Protocols Using CLI	157
Monitor and Verify BFD Configuration	159
Troubleshoot Common BFD Errors	160

CHAPTER 8**Cisco Catalyst SD-WAN BFD 163**

Information About Cisco Catalyst SD-WAN BFD	164
Information About Automatically Suspending BFD Sessions	164
Benefits of Automatically Suspending BFD Sessions	165
How Automatically Suspending BFD Sessions Works	165
Restrictions for Automatically Suspending BFD Sessions	166
Configure Automatic Suspension of BFD Sessions Using a CLI Template	167

Verify Automatic Suspension of BFD Sessions 168

CHAPTER 9

Cisco Catalyst SD-WAN Controller Route Filtering by TLOC Color 171

Information About Cisco SD-WAN Controller Route Filtering by TLOC Color 172

Supported Devices for Cisco SD-WAN Controller Route Filtering by TLOC Color 174

Prerequisites for Cisco SD-WAN Controller Route Filtering by TLOC Color 174

Restrictions for Cisco SD-WAN Controller Route Filtering by TLOC Color 175

Configure Cisco SD-WAN Controller Route Filtering by TLOC Color Using a CLI Template 175

 Enable Route Filtering Using a CLI Template 175

 Configure the Update Interval for Route Filtering by TLOC Color Using a CLI Template 176

 Override Default TLOC Color Compatibility for Cisco SD-WAN Controller Route Filtering by TLOC Color Using a CLI Template 176

Monitor Cisco SD-WAN Controller Route Filtering by TLOC Color 177

 View TLOC Colors for a Device 178

 Check TLOC Color Compatibility 178

CHAPTER 10

Transport Gateway 179

Transport Gateway 180

Information About Transport Gateways 180

 Site Type 182

 OMP Best Path Logic and Transport Gateway Path Preference 182

 Configuration Overview 183

Restrictions for Transport Gateways 186

Use Cases for Transport Gateways 187

Configure a Router as a Transport Gateway Using Cisco SD-WAN Manager 189

Configure a Router as a Transport Gateway Using a CLI Template 190

Configure the Transport Gateway Path Preference 190

 Configure the Transport Gateway Path Preference Using Cisco SD-WAN Manager 190

 Configure the Transport Gateway Path Preference Using a CLI Template 191

Configure the Site Type for a Router Using Cisco SD-WAN Manager 192

Configure the Site Type for a Router Using a CLI Template 192

Verify the Site Type of a Router Using the CLI 193

Verify a Transport Gateway Configuration Using the CLI 193

CHAPTER 11	Hub-and-Spoke	195
	Hub-and-Spoke	195
	Information About Hub-and-Spoke	196
	Example: Hub-and-Spoke Connectivity	197
	Device0 (Hub) Before and After	199
	Device1 (Spoke1) Before and After	202
	Device2 (Spoke2) Before and After	203
	Benefits of Hub-and-Spoke	205
	Restrictions for Hub-and-Spoke	206
	Use Cases for Hub-and-Spoke	206
	Configure a Hub-and-Spoke Topology	207
	Configure a Cisco Catalyst SD-WAN Controller to Enable Hub-and-Spoke Using Cisco SD-WAN Manager	207
	Configure a Cisco SD-WAN Controller to Enable Hub-and-Spoke Using a CLI Template	208
	Configure a Router as a Transport Gateway, for Hub-and-Spoke	208
	Configure the Site Type for a Router, for Hub-and-Spoke	208
	Verify a Hub-and-Spoke Configuration	208
	Verify that a Cisco Catalyst SD-WAN Controller Has Enabled Hub-and-Spoke Configuration	209

CHAPTER 12	Symmetric Routing	211
	Symmetric Routing	212
	Information About Symmetric Routing	212
	Benefits of Symmetric Routing Configuration	212
	Mechanisms for Ensuring Symmetric Routing	213
	Translating OMP Metrics for Devices Outside of the Overlay Network	216
	Translating OMP Metrics to BGP Attributes	217
	Translating OMP Metrics to an OSPF Metric	219
	Configuration Overview	220
	Example of Configuration for Symmetric Routing and the Mechanism	223
	Supported Scenarios	229
	Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active	230
	Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Passive	230
	Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active by VRF	231

Scenario: Multi-Region Fabric Environment	232
Scenario: Multi-Region Fabric, Transport Gateways Serving Subregions	232
Scenario: Multi-Region Fabric with Route Leaking	233
Prerequisites for Symmetric Routing	236
Restrictions for Symmetric Routing	237
Configure Symmetric Routing	237
Configure a Router to Use Automatic Affinity Group Preference Using Cisco SD-WAN Manager	237
Configure a Router Affinity Group or Affinity Group Preference	238
Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager	238
Configure Router Affinity Groups for Specific VRFs Using a CLI Template	238
Configure a Router to Use Automatic Affinity Group Preference Using a CLI Template	239
Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template	239
Verify Symmetric Routing	241
Verify the Next Hops for a Specific Prefix on a Router	241
Verify the Path to a Destination Router	241
Verify the VRF-Specific Affinity Group Configuration on a Router	242
Verify a Control Policy for Route Leaking	242
Verify the Derived Affinity Group of a Route	243
Monitor RIB Metric Translation	244
OMP Metrics	244
BGP Metrics	245
OSPF Metrics	245
CHAPTER 13	Troubleshoot Cisco Catalyst SD-WAN Routing
	247
Overview	247
Support Articles	247
Feedback Request	248
Disclaimer and Caution	248



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 (SD-WAN)



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.



Note Cisco is constantly enhancing the Cisco Catalyst SD-WAN solution with every release and we try and keep the content in line with the latest enhancements. The following table lists new and modified features we documented in the Configuration, Command Reference, and Hardware Installation guides. For information on additional features and fixes that were committed to the Cisco Catalyst SD-WAN solution, see the *Resolved and Open Bugs* section in the Release Notes.

[What's New in Cisco IOS XE Catalyst SD-WAN Release 17.x](#)



CHAPTER 3

Unicast Overlay Routing



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

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

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

- [Supported Protocols, on page 5](#)
- [Configure Unicast Overlay Routing, on page 25](#)

Supported Protocols

This section explains the protocols supported for unicast routing.

OMP Routing Protocol

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

- Orchestration of overlay network communication, including connectivity among network sites, service chaining, and VPN or VRF topologies
- Distribution of service-level routing information and related location mappings

- Distribution of data plane security parameters
- Central control and distribution of routing policy

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

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

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

OMP Route Advertisements

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

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

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

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

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

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



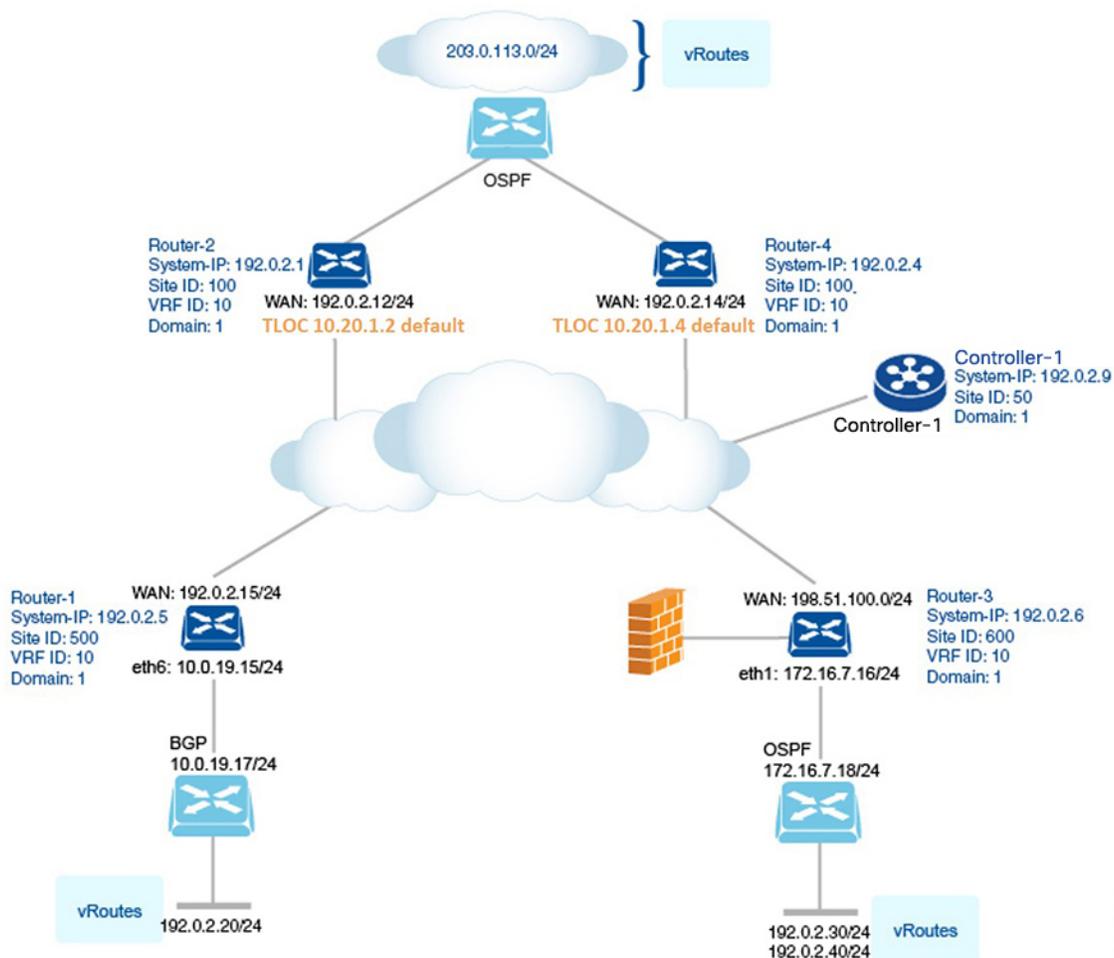
Note Route advertisements to OMP are done by either applying the configuration at the global level or at the specific VPN level. To configure route advertisements to OMP at the global level, use the OMP feature template. On the other hand, to configure route advertisements to OMP at the specific VPN level, use the VPN feature template. For more information about configuring route advertisements to OMP, see [Configure OMP](#), on page 46.

OMP advertises the following types of routes:

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

The following figure illustrates the two types of OMP routes.

Figure 1: Different Types of OMP Routes



520013

OMP Routes

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

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

- Connected (also known as direct)
- Static
- BGP
- EIGRP
- LISP
- OSPF (inter-area, intra-area, and external)
- OSPFv3 (inter-area, intra-area, and external)

- IS-IS

OMP routes advertise the following attributes:

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

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

TLOC Routes

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

TLOC routes advertise the following attributes:

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

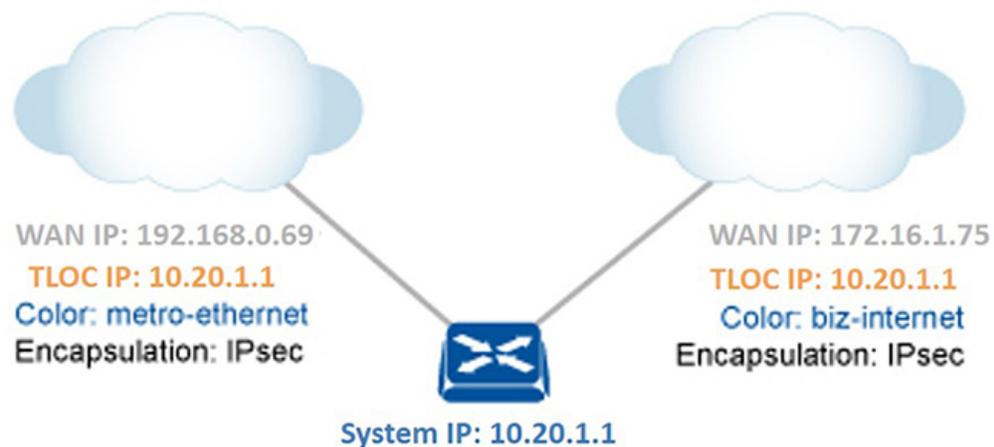
- **Tag**—Optional, transitive path attribute that an OMP speaker can use to control the flow of routing information toward a TLOC. When an OMP route is advertised along with its TLOC, both or either can be distributed with a community TAG, to be used to decide how to send traffic to or receive traffic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

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

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

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

Figure 2: Router Attributes



368487

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

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

OMP Route Advertisements for Cisco Catalyst SD-WAN Controllers

Table 1: Feature History

Feature Name	Release Information	Description
Increased OMP Path Limit for Cisco Catalyst SD-WAN Controllers	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a	This feature extends the limit on the number of OMP routes that can be exchanged between Cisco Catalyst SD-WAN Controllers to 128. Prior to this release, the limit was 16.

Overview

The transport location (TLOC) information is advertised to the OMP peers including Cisco Catalyst SD-WAN Controllers and its local-site branches. Starting from Cisco IOS XE Catalyst SD-WAN Release 17.5.1a, the limit on the number of OMP paths that can be exchanged between Cisco Catalyst SD-WAN Controllers per VPN per prefix is extended to a maximum of 128.

Limitations

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

Configure Path Limit

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

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

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



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

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

For more information about configuring path limits, see [controller-send-path-limit](#) command page.

OMP Route Redistribution

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

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

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

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

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

The following is an example for advertise network configuration:

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

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

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

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

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

```
omp
  no shutdown
```

```

ecmp-limit      6
graceful-restart
no as-dot-notation
timers
  holdtime      15
  graceful-restart-timer 120
exit
address-family ipv4
  advertise ospf external <-- This configuration implies OSPF Inter-Area/Intra-Area routes
& External routes are redistributed into OMP
  advertise connected
  advertise static
!
```

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

```

omp
no shutdown
ecmp-limit      6
graceful-restart
no as-dot-notation
timers
  holdtime      15
  graceful-restart-timer 120
exit
address-family ipv4 vrf 1
  advertise ospf external
  advertise ospf route-map RLB
!
```

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

The following example shows the redistribution of OSPFv3 external routes:

```

omp
no shutdown
ecmp-limit      6
graceful-restart
no as-dot-notation
timers
  holdtime      15
  graceful-restart-timer 120
exit
address-family ipv6
  advertise ospfv3
  advertise ospf external
!
```



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

Table 2:

OMP Route Origin Type	OMP Route Origin Subtype
BGP	External Internal

OMP Route Origin Type	OMP Route Origin Subtype
Connected	—
OSPF	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
OSPFv3	Intra-area, Inter-area, External-1, External-2, NSSA-External-1 and NSSA-External-2
Static	—
EIGRP	<ul style="list-style-type: none"> • EIGRP Summary • EIGRP Internal • EIGRP External
LISP	—
IS-IS	Level 1 and level 2

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

Administrative Distance

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

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

Table 3:

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

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

OMP Best-Path Algorithm

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

Table 4: Best Path Algorithm

Step	Applies to	Description
1	Edge devices Cisco Catalyst SD-WAN Controller	Path validity Check whether the OMP path is valid. If not, ignore it.
2	Edge devices Cisco Catalyst SD-WAN Controller	Active vs. stale paths Prefer an active path over a stale path. An active path is a one from a peer with which an OMP session is up. A stale path is one from a peer with which an OMP session is in Graceful Restart mode. Note A stale path will only be advertised if the stale version is similar to the Route Information Base (RIB) version. Otherwise, the stale path is dropped.
3	Edge devices	Administrative distance Select the OMP path with the lower administrative distance. Example: A path that the device learns locally via BGP would be preferred over a path that it learns from a Cisco SD-WAN Controller via OMP. For information about administrative distance, see Administrative Distance, on page 14 .
4	Edge devices Cisco Catalyst SD-WAN Controller	OMP path preference Select the OMP path with the higher OMP path preference value.
5	Cisco Catalyst SD-WAN Controller	Access region Cisco SD-WAN Controller drops advertisement from border router (BR) to BR in the same region.

Step	Applies to	Description
6	Edge devices	<p>Core region</p> <p>Cisco SD-WAN Controller allows advertisement between BRs in the same access region, but receiving BR drops advertisement.</p>
7	Multi-Region Fabric scenario only Edge devices	<p>Region path length</p> <p>Compare region-path-length. Prefer lower. If region-path-length-ignore is configured, then skip this step. (This addresses secondary regions in Multi-Region Fabric.)</p>
8	Multi-Region Fabric scenario only Border routers	<p>Access region vs. core region</p> <p>Prefer access region paths over core region paths.</p>
9	Edge devices	<p>Direct vs. transport gateway path</p> <p>Prefer a direct path over a transport gateway path.</p> <p>This step can be modified by the transport gateway path preference options, which can (a) cause the transport gateway path to be preferred, or (b) result in the paths to be considered equal. See Configure the Transport Gateway Path Preference in the <i>Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide</i>.</p>
10	Multi-Region Fabric scenario only Edge devices	<p>Multi-Region Fabric subregion comparison</p> <ul style="list-style-type: none"> • Prefer paths from the router's own subregion. • When comparing two paths that are not from the router's subregion, prefer a path that is not part of any subregion.
11	Multi-Region Fabric scenario only Edge devices	<p>Border router preference</p> <p>Prefer a path with a higher border router preference value.</p>
12	Edge devices	<p>Derived affinity</p> <p>Prefer a path with a lower derived affinity value.</p>
13	Edge devices with an affinity preference configured	<p>Affinity preference</p> <p>Depending on the affinity preference configured on the device, prefer a path whose affinity is earlier in the preference list (higher priority). If the device uses affinity-preference-auto, then prefer a path with a numerically lower affinity group.</p> <p>Note When comparing two paths with similar reorigination types, one with an affinity value and one without, prefer the path with an affinity value.</p>

Step	Applies to	Description
14	Edge devices	<p>TLOC preference</p> <p>Select an OMP path with a higher TLOC preference value.</p>
15	Edge devices Cisco Catalyst SD-WAN Controller	<p>Origin type and subtype</p> <p>Compare the origin type and subtype, and select the first match in the following list:</p> <ul style="list-style-type: none"> • Connected • Static • EIGRP Summary • BGP External • EIGRP Internal • OSPF/OSPFv3 Intra-area • OSPF/OSPFv3 Inter-area • IS-IS Level 1 • EIGRP External • OSPF/OSPFv3 External (External OSPF Type1 is preferred over External OSPF Type2) • IS-IS Level 2 • BGP Internal • Unknown
16	Edge devices Cisco Catalyst SD-WAN Controller	<p>Origin metric</p> <p>Select an OMP path that has a lower origin metric.</p>
17	Cisco Catalyst SD-WAN Controller	<p>Path source</p> <p>Prefer a path sourced from an edge router over the same path coming from a Cisco Catalyst SD-WAN Controller.</p>
18	Edge devices Cisco Catalyst SD-WAN Controller	<p>Private IP address</p> <p>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.</p>



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

Here are some examples of choosing the best path:

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

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

OMP Graceful Restart

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

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



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

BGP and OSPF Routing Protocols

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

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

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

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

BGP Community Propagation

Table 5: Feature History

Feature Name	Release Information	Description
Ability to Match and Set Communities during BGP to OMP Redistribution	Cisco IOS XE Catalyst SD-WAN Release 17.4.1a Cisco vManage Release 20.4.1	This feature enhances the implementation of match and set clauses for redistribution of routes from BGP to OMP and vice versa on Cisco IOS XE Catalyst SD-WAN devices. You can redistribute the routes from a BGP into an OMP routing to allow route traffic to help increase the accessibility within the network. The <code>route-maps</code> 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: <pre>route-map advertise bgp route-map bgp-to-omp redistribute omp route-map omp-to-bgp</pre>

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

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

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

OSPFv3

Table 6: Feature History

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

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

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

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

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

EIGRP

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

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

See [Introduction to EIGRP](#) for more information in EIGRP.

Benefits of EIGRP

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

Limitations and Restrictions

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

Routing Information Protocol (RIP)

Table 7: Feature History

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

Information About Routing Information Protocol Support

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

RIPv2 (RIP for IPv4)

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

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

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

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



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

See [Configure Routing Information Protocol Using the CLI](#) for more details on configurations using the CLI.

RIPng (RIP for IPv6)

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

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

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

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

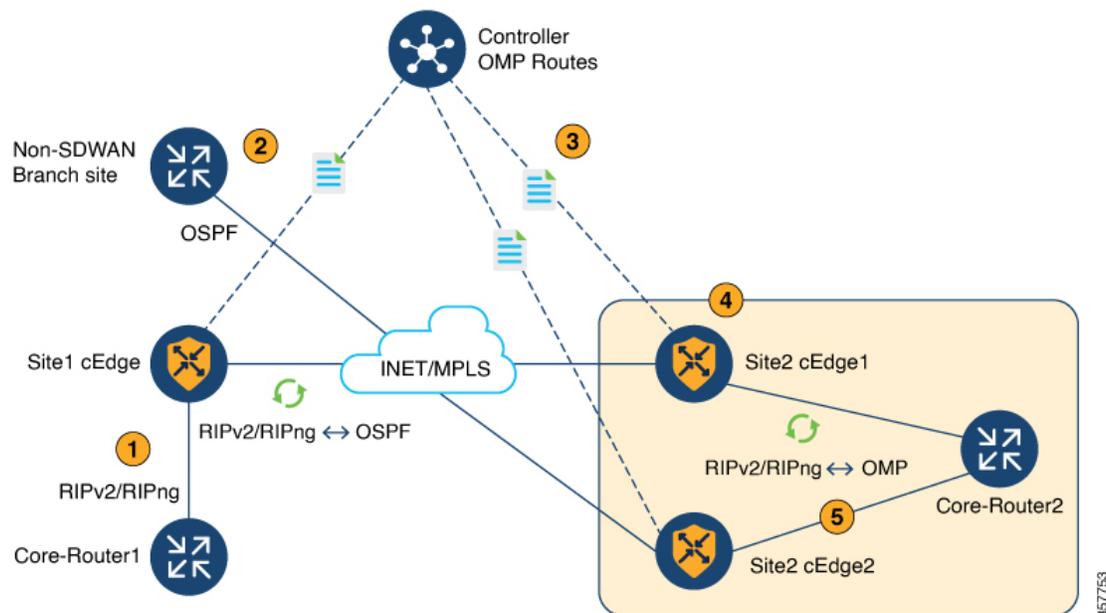
- The IPv6 prefix of the destination.
- Metric: Total cost of the metric advertised for the address.
- Route Tag: A route attribute that must be advertised and redistributed with the route.
- Next-hop IPv6 address of the destination.
- Various timers associated with the routes.

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

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

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

Figure 3: RIPv2 and RIPv2 Topology



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

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

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

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

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

See [Configure RIPv2 Using the CLI](#) for more details on RIPv2 configurations using the CLI.

Prerequisites for Using Routing Information Protocol

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

Restrictions for Using Routing Information Protocol

RIPv2 (IPv4)

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

RIPng (IPv6)

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

Configure Unicast Overlay Routing

This topic describes how to provision unicast overlay routing.

Transport-Side Routing

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

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

Configure BGP

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



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

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

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

Create a BGP Template

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



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

3. Click **Create Template**
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
6. To create a template for **VPN 0** or **VPN 512**:
 - a. Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
 - b. Under **Additional VPN 0 Templates**, click **BGP**.
 - c. From the **BGP** drop-down list, click **Create Template**. The BGP template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining BGP parameters.
7. To create a template for VPNs **1** through **511**, and **513** through **65530**:
 - a. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
 - b. Click the **Service VPN** drop-down list.
 - c. Under **Additional VPN Templates**, click **BGP**.
 - d. From the **BGP** drop-down list, click **Create Template**. The BGP template form displays. The top of the form contains fields for naming the template, and the bottom contains fields for defining BGP parameters.
8. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
9. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

Configure Basic BGP Parameters

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

Parameter Name	Description
Shutdown*	Click No to enable BGP for the VPN.
AS number*	Enter the local AS number.
Router ID	Enter the BGP router ID in decimal four-part dotted notation.

Parameter Name	Description
Propagate AS Path	Click On to carry BGP AS path information into OMP.
Internal Routes Distance	Enter a value to apply as the BGP route administrative distance for routes coming from one AS into another. Range: 0 through 255 Default: 200
Local Routes Distance	Specify the BGP route administrative distance for routes within the local AS. By default, a route received locally from BGP is preferred over a route received from OMP. Range: 0 through 255 Default: 200
External Routes Distance	Specify the BGP route administrative distance for routes learned from other sites in the overlay network. Range: 0 through 255 Default: 20

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

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

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

Configure Unicast Address Family

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

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

Parameter	Option	Sub-Option	Description	
Redistribute	Click Redistribute > New Redistribute .			
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.		
	Protocol	Choose the protocols from which to redistribute routes into BGP, for all BGP sessions. Options are:		
		static	Redistribute static routes into BGP.	
		connected	Redistribute connected routes into BGP.	
		ospf	Redistribute Open Shortest Path First routes into BGP.	
		omp	Redistribute Overlay Management Protocol routes into BGP.	
		nat	Redistribute Network Address Translation routes into BGP.	
		natpool-outside	Redistribute outside NAT routes into BGP.	
	At a minimum, choose the following:			
<ul style="list-style-type: none"> • For service-side BGP routing, choose OMP. By default, OMP routes are not redistributed into BGP. • For transport-side BGP routing, choose Connected, and then under Route Policy, specify a route policy that has BGP advertise the loopback interface address to its neighbors. 				
Route Policy	Enter the name of the route policy to apply to redistributed routes.			
Click Add to save the redistribution information.				
Network	Click Network > New Network .			
	Mark as Optional Row	Check Mark as Optional Row to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.		
	Network Prefix	Enter a network prefix, in the format <i>prefix/length</i> to be advertised by BGP.		
	Click Add to save the network prefix.			

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

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

To change the AS number, perform the following steps:

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

Configure BGP Neighbors

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



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

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

Parameter Name	Options	Sub-Options	Description
Address Family	Click On and select the address family. Enter the address family information. The software supports only the BGP IPv4 unicast address family.		
	Address Family	Select the address family. The software supports only the BGP IPv4 unicast address family.	
	Maximum Number of Prefixes	Specify the maximum number of prefixes that can be received from the neighbor. Range: 1 through 4294967295 Default: 0	
		Threshold	Specify the threshold at which to generate a warning message or restart the BGP connection. The threshold is a percentage of the maximum number of prefixes. You can specify either a restart interval or a warning only.
	Restart Interval	Specify the duration to wait for restarting the BGP connection. <i>Range:</i> 1 through 65535 minutes	
	Warning Only	Click On to display a warning message without restarting the BGP connection.	
	Route Policy In	Click On and specify the name of a route policy that will have the prefixes from the neighbour.	
Route Policy Out	Click On and specify the name of a route policy that will have the prefixes sent to the neighbour.		
Shutdown	Click On to enable the connection to the BGP neighbor.		

Configure MPLS Interface

Table 8: Feature History

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

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

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

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

To configure an MPLS interface, do the following:

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

Configure Label Range

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

Sample configuration:

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

Configure Route Targets

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

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

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

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

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

Configure Advanced Neighbor Parameter

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

Parameter Name	Description
Next-Hop Self	Click On to configure the router to be the next hop for routes advertised to the BGP neighbor.
Send Community	Click On to send the local router's BGP community attribute to the BGP neighbor.
Send Extended Community	Click On to send the local router's BGP extended community attribute to the BGP neighbor.
Negotiate Capability	Click On to allow the BGP session to learn about the BGP extensions that are supported by the neighbor.
Source Interface Address	Enter the IP address of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor.
Source Interface Name	Enter the name of a specific interface of the neighbor that BGP is to use for the TCP connection to the neighbor, in the format ge port/slot .
EBGP Multihop	Set the time to live (TTL) for BGP connections to external peers. Range: 0 to 255 Default: 1
Password	Enter a password to use to generate an MD5 message digest. Configuring the password enables MD5 authentication on the TCP connection with the BGP peer. The password is case-sensitive and can be up to 25 characters long. It can contain any alphanumeric characters, including spaces. The first character cannot be a number.
Keepalive Time	Specify the frequency at which keepalive messages are advertised to a BGP peer. These messages indicate to the peer that the local router is still active and should be considered available. Specify the keepalive time for the neighbor to override the global keepalive time. Range: 0 through 65535 seconds Default: 60 seconds (one-third the hold-time value)
Hold Time	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local router then terminates the BGP session to that peer. Specify the hold time for the neighbor to override the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive timer)
Connection Retry Time	Specify the number of seconds between retries to establish a connection to a configured BGP neighbor peer that has gone down. Range: 0 through 65535 seconds Default: 30 seconds

Parameter Name	Description
Advertisement Interval	For the BGP neighbor, set the minimum route advertisement interval (MRAI) between when BGP routing update packets are sent to that neighbor. Range: 0 through 600 seconds Default: 5 seconds for IBGP route advertisements; 30 seconds for EBGP route advertisements

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

Change the Scope of a Parameter Value

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

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

Configure Advanced BGP Parameters

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

Parameter Name	Description
Hold Time	Specify the interval after not receiving a keepalive message that the local BGP session considers its peer to be unavailable. The local device then terminates the BGP session to that peer. This hold time is the global hold time. Range: 0 through 65535 seconds Default: 180 seconds (three times the keepalive timer)

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

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

Configure BGP Using CLI

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

```

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

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

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

```



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

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

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

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

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

Verify BGP Redistribute Route in OMP

```
Device#show sdwan omp routes 10.0.0.0/8
```

```
-----
omp route entries for vpn 100 route 10.0.0.0/8
-----
```

```

RECEIVED FROM:
peer          172.16.0.0
path-id       470777
label         1002
status        C,I,R
loss-reason   not set
lost-to-peer  not set
lost-to-path-id not set
Attributes:
  originator   10.0.0.1
  type         installed
  tloc         172.16.0.1, mpls, ipsec
  ultimate-tloc not set
  domain-id    not set
  overlay-id   1
  site-id      1
  preference   not set
  tag          10000 <=====
  origin-proto eBGP
  as-path      not set
  unknown-attr-len not set
```

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

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

```

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

```

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

Configure Service-Side Routing

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

1. Configure a VPN.

```

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

```

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

2. Configure BGP to run in the VPN:

a. Configure the local AS number:

```

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

```

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

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

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

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

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

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

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

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

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

  exit-address-family
```

Example of configuring route targets:

```
vrf config

vrf definition 1
  rd 1:1

!
address-family ipv4

  route-target export 200:1

  route-target import 100:1
```

```

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

```

Redistribute BGP Routes and AS Path Information

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

```

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

```

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

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

```

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

```

```

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

```

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

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

```

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

```

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

```

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

```

Configure OSPF

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



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

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

1. Create an OSPF feature template to configure OSPF parameters. OSPF can be used for transport-side routing to enable communication between the Cisco Catalyst SD-WAN devices when the router is not directly connected to the WAN cloud.
2. Create a VPN feature template to configure VPN parameters for transport-side OSPF routing (in VPN 0). See the VPN help topic for more information.

Create an OSPF Template

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



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

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, select the type of device for which you are creating the template. To create a template for VPN 0 or VPN 512:
 - a. Click **Transport & Management VPN** located directly beneath the **Description** field, or scroll to the **Transport & Management VPN** section.
 - b. Under **Additional VPN 0 Templates**, click **OSPF**.
 - c. From the **OSPF** drop-down list, click **Create Template**. The OSPF template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OSPF parameters.
6. To create a template for VPNs 1 through 511, and 513 through 65530:
 - a. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
 - b. Click the **Service VPN** drop-down list.
 - c. Under **Additional VPN Templates**, click **OSPF**.

- d. From the **OSPF** drop-down list, click **Create Template**. The OSPF template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OSPF parameters.
7. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

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

Table 9:

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

Configure Basic OSPF

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

Table 10:

Parameter Name	Description
Router ID	Enter the OSPF router ID in decimal four-part dotted notation. This is the unique 32-bit identifier associated with the OSPF router for Link-State Advertisements (LSAs) and adjacencies.

Parameter Name	Description
Distance for External Routes	Specify the OSPF route administration distance for routes learned from other domains. <i>Range: 0 through 255 Default: 110</i>
Distance for Inter-Area Routes	Specify the OSPF route administration distance for routes coming from one area into another. <i>Range: 0 through 255 Default: 110</i>
Distance for Intra-Area Routes	Specify the OSPF route administration distance for routes within an area. <i>Range: 0 through 255 Default: 110</i>

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

Redistribute Routes into OSPF

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

Table 11:

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

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

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

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

Configure OSPF To Advertise a Maximum Metric

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

Table 12:

Parameter Name	Description
Type	Choose a type: <ul style="list-style-type: none"> • Administrative—Force the maximum metric to take effect immediately through operator intervention. • On-Startup—Advertise the maximum metric for the specified time.

Parameter Name	Description
Advertisement Time	If you selected On-Startup , specify the number of seconds to advertise the maximum metric after the router starts up. <i>Range:</i> 0, 5 through 86400 seconds <i>Default:</i> 0 seconds (the maximum metric is advertised immediately when the router starts up)

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

Configure OSPF Areas

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

Table 13:

Parameter Name	Description
Area Number	Enter the number of the OSPF area. <i>Range:</i> 32-bit number
Set the Area Type	Choose the type of OSPF area, Stub or NSSA.
No Summary	Click On to not inject OSPF summary routes into the area.
Translate	If you configured the area type as NSSA, choose when to allow Cisco Catalyst SD-WAN devices that are ABRs (area border routers) to translate Type 7 LSAs to Type 5 LSAs: <ul style="list-style-type: none"> • Always—Router always acts as the translator for Type 7 LSAs. That is no other router, even if it is an ABR, can be the translator. If two ABRs are configured to always be the translator, only one of them actually ends up doing the translation. • Candidate—Router offers translation services, but does not insist on being the translator. • Never—Translate no Type 7 LSAs.

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

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

Configure Interfaces in an OSPF Area

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

Table 14:

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

Parameter Name	Description
Hello Interval	Specify how often the router sends OSPF hello packets. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 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. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 40 seconds (4 times the default hello interval)
LSA Retransmission Interval	Specify how often the OSPF protocol retransmits LSAs to its neighbors. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 5 seconds
Interface Cost	Specify the cost of the OSPF interface. <i>Range:</i> 1 through 65535

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

Table 15:

Parameter Name	Description
Designated Router Priority	Set the priority of the router to be elected as the designated router (DR). The router with the highest priority becomes the DR. If the priorities are equal, the node with the highest router ID becomes the DR or the backup DR. <i>Range:</i> 0 through 255 <i>Default:</i> 1
OSPF Network Type	Choose the OSPF network type to which the interface is to connect: <ul style="list-style-type: none"> • Broadcast network—WAN or similar network. • Point-to-point network—Interface connects to a single remote OSPF router. • Non-broadcast—Point-to-multipoint. <i>Default:</i> Broadcast
Passive Interface	Click On or Off to specify whether to set the OSPF interface to be passive. A passive interface advertises its address, but does not actively run the OSPF protocol. <i>Default:</i> Off
Authentication	Specify the authentication and authentication key on the interface to allow OSPF to exchange routing update information securely.
• Authentication Type	Choose the authentication type: <ul style="list-style-type: none"> • Simple authentication—Password is sent in clear text. • Message-digest authentication—MD5 algorithm generates the password.

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

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

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

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

Configure an Interface Range for Summary LSAs

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

Table 16:

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

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

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

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

Configure Other OSPF Properties

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

Table 17:

Parameter Name	Description
Reference Bandwidth	Specify the reference bandwidth for the OSPF auto-cost calculation for the interface. <i>Range: 1 through 4294967 Mbps Default: 100 Mbps</i>
RFC 1538 Compatible	By default, the OSPF calculation is done per RFC 1583. Click Off to calculate the cost of summary routes based on RFC 2328.
Originate	Click On to generate a default external route into an OSPF routing domain: <ul style="list-style-type: none"> • Always—Click On to always advertise the default route in an OSPF routing domain. • Default metric—Set the metric used to generate the default route. <i>Range: 0 through 16777214 Default: 10</i> • Metric type—Select to advertise the default route as an OSPF Type 1 external route or an OSPF Type 2 external route.
SPF Calculation Delay	Specify the amount of time between when the first change to a topology is received until performing the SPF calculation. <i>Range: 0 through 600000 milliseconds (60 seconds) Default: 200 milliseconds</i>
Initial Hold Time	Specify the amount of time between consecutive SPF calculations. <i>Range: 0 through 600000 milliseconds (60 seconds) Default: 1000 milliseconds</i>
Maximum Hold Time	Specify the longest time between consecutive SPF calculations. <i>Range: 0 through 600000 Default: 10000 milliseconds (60 seconds)</i>
Policy Name	Enter the name of a localized control policy to apply to routes coming from OSPF neighbors.

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

Configure OSPF Using CLI

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

Configure Basic Service-Side OSPF

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



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

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

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

Configure OMP

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

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



- Note**
- Route advertisements in OMP are done either by applying the configuration at the global level or at the specific VRF level. For more information about route advertisements in OMP, see the *Configure OMP Advertisements* section in this topic.
 - Cisco IOS XE Catalyst SD-WAN devices use VRFs in place of VPNs. However, the following steps still apply to configure Cisco IOS XE Catalyst SD-WAN devices through Cisco SD-WAN Manager. When you complete the configuration, the system automatically maps the VPN configurations to VRF configurations.

Create OMP Template

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



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

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you're creating the template.
6. To create a custom template for OMP, choose the **Factory_Default_OMP_Template** and click **Create Template**. The OMP template form is displayed. The top of the form contains fields for naming the template, and the bottom contains fields for defining OMP parameters. You may need to click an operation or the plus sign (+) to display more fields.
7. In the **Template Name** field, enter a name for the template. The name can be up to 128 characters and can contain only alphanumeric characters.
8. In the **Template Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.

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

Table 18:

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

Configure Basic OMP Options

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

Table 19:

Parameter Name	Description
Graceful Restart for OMP	Ensure that Yes is selected to enable graceful restart. By default, graceful restart for OMP is enabled.
Overlay AS Number	Specify a BGP AS number that OMP advertises to the router's BGP neighbors.
Graceful Restart Timer	Specify how often the OMP information cache is flushed and refreshed. A timer value of 0 disables OMP graceful restart. <i>Range:</i> 0 through 604800 seconds (168 hours, or 7 days) <i>Default:</i> 43200 seconds (12 hours)
Number of Paths Advertised Per Prefix	Specify the maximum number of equal-cost routes to advertise per prefix. s advertise routes to Cisco Catalyst SD-WAN Controllers, and the controllers redistributes the learned routes, advertising each route-TLOC tuple. A Cisco IOS XE Catalyst SD-WAN device can have up to four TLOCs, and by default advertises each route-TLOC tuple to the Cisco Catalyst SD-WAN Controller. If a local site has two Cisco IOS XE Catalyst SD-WAN devices, a Cisco Catalyst SD-WAN Controller could potentially learn eight route-TLOC tuples for the same route. If the configured limit is lower than the number of route-TLOC tuples, the best route or routes are advertised. <i>Range:</i> 1 through 16 <i>Default:</i> 4
ECMP Limit	Specify the maximum number of OMP paths received from the Cisco Catalyst SD-WAN Controller that can be installed in the Cisco IOS XE Catalyst SD-WAN device's local route table. By default, a Cisco IOS XE Catalyst SD-WAN device installs a maximum of four unique OMP paths into its route table. <i>Range:</i> 1 through 16 <i>Default:</i> 4
Send Backup Paths (on Cisco Catalyst SD-WAN Controllers only)	Click On to have OMP advertise backup routes to Cisco IOS XE Catalyst SD-WAN devices. By default, OMP advertises only the best route or routes. If you configure to send backup paths, OMP also advertises the first non-best route in addition to the best route or routes.
Shutdown	Ensure that No is chosen to enable to the Cisco SD-WAN overlay network. Click Yes to disable OMP and disable the Cisco SD-WAN overlay network. OMP is enabled by default.
Discard Rejected (on Cisco Catalyst SD-WAN Controllers only)	Click Yes to have OMP discard routes that have been rejected on the basis of policy. By default, rejected routes aren't discarded.

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

Configure OMP Timers

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

Table 20:

Parameter Name	Description
Advertisement Interval	Specify the time between OMP Update packets. <i>Range: 0 through 65535 seconds Default: 1 second</i>
Hold Time	Specify how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed. <i>Range: 0 through 65535 seconds Default: 60 seconds</i> Starting from Cisco Catalyst SD-WAN Manager Release 20.12.1, the default hold time is 300 seconds.
EOR Timer	Specify how long to wait after an OMP session has gone down and then come back up to send an end-of-RIB (EOR) marker. After this marker is sent, any routes that weren't refreshed after the OMP session came back up are considered to be stale and are deleted from the route table. <i>Range: 1 through 3600 seconds (1 hour) Default: 300 seconds (5 minutes)</i>

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

Configure OMP Advertisements



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

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

Table 21:

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

Click **Save**.

Configure OMP Using CLI

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

OMP support in Cisco SD-WAN includes the following:

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

Configure OMP Graceful Restart

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

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

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

To disable OMP graceful restart:

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

The graceful restart timer is set up independently on each OMP peer; that is, it's set up separately on each Cisco IOS XE Catalyst SD-WAN device and Cisco Catalyst SD-WAN Controller. To illustrate what this means, let's consider a vSmart controller that uses a graceful restart time of 300 seconds, or 5 minutes, and a Cisco IOS XE Catalyst SD-WAN device that is configured with a timer of 600 seconds (10 minutes). Here, Cisco Catalyst SD-WAN Controller retains the OMP routes learned from that device for 10 minutes—the graceful restart timer value that is configured on the device and that the device has sent to Cisco Catalyst SD-WAN Controller during the setup of the OMP session. The Cisco IOS XE Catalyst SD-WAN device retains the routes it learns from the vSmart controller for 5 minutes, which is the default graceful restart time value that is used on the Cisco Catalyst SD-WAN Controller and that the controller sent to the device, also during the setup of the OMP session.

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

Advertise Routes to OMP

Table 22: Feature History

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

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

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



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

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

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

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

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

```

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

```



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

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

```

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

```

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

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

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

```

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

```

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

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



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

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

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

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

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

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

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

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

Configure the Number of Advertised Routes

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

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

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

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

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

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

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

Configure the Number of Installed OMP Paths

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

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

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

Configure the OMP Hold Time

The OMP hold time determines how long to wait before closing the OMP connection to a peer. If the peer doesn't receive three consecutive keepalive messages within the hold time, the OMP connection to the peer is closed. The default OMP hold time is 60 seconds but it can be configured to up to 65,535 seconds.

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

To modify the OMP hold time interval:

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

The hold time can be in the range 0 through 65535 seconds.

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

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

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

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

Configure the OMP Update Advertisement Interval

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

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

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

Configure the End-of-RIB Timer

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

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

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

Mapping Multiple BGP Communities to OMP Tags

Table 23: Feature History

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

For more information on the **show sdwan omp routes** command, refer [show sdwan omp routes](#).

Configure OSPFv3

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

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

Create an OSPFv3 Template

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



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

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

Basic Configuration

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

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

Configure Redistribute Routes into OSPFv3

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

Table 24: Redistribution Parameters

Parameter Name	Value	Description
Mark as Optional Row		Click Optional to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
Protocol *		Choose the protocols from which to redistribute routes into OSPFv3, for all OSPFv3 sessions.
	bgp	Redistribute BGP routes into OSPFv3.
	connected	Redistribute connected routes into OSPFv3.
	nat-route	Redistribute NAT routes into OSPFv3.
	omp	Redistribute OMP routes into OSPFv3.
	eigrp	Redistribute EIGRP routes into OSPFv3.
	lisp	Redistribute LISP routes into OSPFv3.
	isis	Redistribute IS-IS routes into OSPFv3.
	ospf	Redistribute OSPF routes into OSPFv3. Note Redistribute of OSPF is supported only for IPv4 address family.
	static	Redistribute static routes into OSPFv3.
Route Policy *		Enter the name of the route policy to apply to redistributed routes.

Click **Save**.

Configure OSPFv3 To Advertise a Maximum Metric

To configure OSPFv3 to advertise a maximum metric so that other devices do not prefer the Cisco IOS XE Catalyst SD-WAN device as an intermediate hop in their Shortest Path First (SPF) calculation, choose **Maximum Metric (Router LSA) > Add New Router LSA** and configure the following parameters:

Table 25:

Parameter Name	Description
Type	Choose a type: <ul style="list-style-type: none"> • Administrative—Force the maximum metric to take effect immediately through operator intervention. • On-Startup—Advertise the maximum metric for the specified time after the startup.

Parameter Name	Description
Advertisement Time	If you chose On-Startup , specify the number of seconds to advertise the maximum metric after the router starts up. <i>Range:</i> 0, 5 through 86400 seconds <i>Default:</i> 0 seconds (the maximum metric is advertised immediately when the router starts up)

Click **Save**.

Configure OSPFv3 Areas

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

Table 26:

Parameter Name	Description
Area Number	Enter the number of the OSPFv3 area. <i>Range:</i> 32-bit number
Set the Area Type	Choose the type of OSPFv3 area. The options are: <ul style="list-style-type: none"> • normal • stub - no external routes • nssa - not-so-stubby area, allows external routes
No Summary	If you configured the area type as stub or NSSA, click On to prevent OSPFv3 summary routes from being injected into the area.
Translate	If you configured the area type as NSSA, choose when you allow devices configured as area border routers (ABR) to translate Type 7 LSAs to Type 5 LSAs: <ul style="list-style-type: none"> • Always—The router always acts as the translator for Type 7 LSAs. That is no other router, even if it is an ABR, can be the translator. If two ABRs are configured to always be the translator, only one of them actually ends up doing the translation. • Candidate—The router offers translation services, but does not insist on being the translator. • Never—The router never become the NSSA translator for Type 7 LSAs.

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

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

Parameter Name	Description
Dead Interval	Specify the time interval within which the Cisco IOS XE Catalyst SD-WAN device must receive an OSPF hello packet from its neighbor. If no packet is received, the Cisco IOS XE Catalyst SD-WAN device assumes that the neighbor is down. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 40 seconds (4 times the default hello interval)
LSA Retransmission Interval	Specify how often the OSPF protocol retransmits LSAs to its neighbors. <i>Range:</i> 1 through 65535 seconds <i>Default:</i> 5 seconds
Interface Cost	Specify the cost of the OSPF interface. <i>Range:</i> 1 through 65535

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

Parameter Name	Description
Address	Enter the IP address and prefix length, in the format <i>prefix/length</i> for the IP or IPv6 address twice to be consolidated and advertised. The address type is dependent on the address family.
Cost	Specify a number for the Type 3 summary LSA. OSPFv3 uses this metric during its SPF calculation to determine the shortest path to a destination. <i>Range:</i> 0 through 16777215
No Advertise	Click On to not advertise the Type 3 summary LSAs or Off to advertise them.

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

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

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

Configure Advanced OSPFv3 Properties

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

Table 27:

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

Parameter Name	Description
Originate	Click On to generate a default external route into an OSPF routing domain: <ul style="list-style-type: none"> • Always—Click On to always advertise the default route in an OSPF routing domain. • Default metric—Set the metric used to generate the default route. <i>Range: 0 through 16777214 Default: 10</i> • Metric type—Choose the metric type, OSPF Type 1 external route or an OSPF Type 2 external route to advertise the default route.
SPF Calculation Delay (milliseconds)	Specify the time between when the first change to a topology is received until performing the SPF calculation. <i>Range: 0 through 600000 milliseconds (60 seconds) Default: 200 milliseconds</i>
Initial Hold Time (milliseconds)	Specify the time between consecutive SPF calculations. <i>Range: 0 through 600000 milliseconds (60 seconds) Default: 1000 milliseconds</i>
Maximum Hold Time (milliseconds)	Specify the longest time between consecutive SPF calculations. <i>Range: 0 through 600000 Default: 10000 milliseconds (60 seconds)</i>
Policy Name	Enter the name of a localized control policy to apply to the routes installed by OSPFv3 into the global Route Information Base (RIB).
Filter	Filter inhibit OSPFv3 routes that do not match the policy from being installed in the global RIB.

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

Configure OSPFv3 Using CLI

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

```

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

[metric-type {1|2}]

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

```

```

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

address-family ipv6 unicast vrf <vpn-id>
    router-id <ipv4-address-format>
    auto-cost reference-bandwidth <1-4294967>
    default-information originate [always] [route-map <route-map-name>] [metric <1-16777214>]
                                [metric-type {1|2}]
    distance <1-254>
    distance ospf {external <1-254> | intra-area <1-254> | inter-area <1-254>}
    timers throttle spf <1-600000> <1-600000> <1-600000>
    redistribute {bgp <1-4294967295> | connected | eigrp <vpn-id>| isis <vpn-id>| lisp | omp
|
                                static}
                                [route-map <route-map-name>]
    max-metric router-lsa [on-startup <5-86400>]
    table-map <route-map-name> [filter]
    area <1-4294967295> stub [no-summary]
    area <1-4294967295> nssa [no-summary] [translate type7 always]
    area <1-4294967295> range <ipv6-prefix>
                                ! 2001:DB8::/48
                                [not-advertise | advertise] [cost
<1-16777214>]
exit-address-family

```

OSPFv3 Table-Map Configuration

```

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

```

OSPFv3 IPsec Interface Authentication

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

```

interface GigabitEthernet2
    vrf forwarding 1

```

```

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

```

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

```

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

```

Configure EIGRP

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

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

Create an EIGRP Template

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



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

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

Basic Configuration

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

Parameter Name	Description
Autonomous System ID *	Enter the local AS number. <ul style="list-style-type: none"> • Range: 1-65,535 • Default: None

Configure IP4 Unicast Address Family

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

Table 28: Redistribution Parameters

Parameter Name	Value	Description
Mark as Optional Row		Click Optional to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
Protocol *		Select the protocols from which to redistribute routes into EIGRP, for all EIGRP sessions.
	bgp	Redistribute Border Gateway Protocol (BGP) routes into EIGRP.
	connected	Redistribute connected routes into EIGRP.
	nat-route	Redistribute network address translation (NAT) routes into EIGRP.
	omp	Redistribute Overlay Management Protocol (OMP) routes into EIGRP.
	ospf	Redistribute Open Shortest Path First (OSPF) routes into EIGRP. Note You can set metric values for redistribution using the CLI add-on feature template from Cisco IOS XE Catalyst SD-WAN Release 16.12.1b and later. Use the following command: <pre>redistribute ospf 1 metric 1000000 1 1 1 1500</pre> For more information, see CLI Add-on Feature Templates .
	static	Redistribute static routes into EIGRP.
Route Policy *		Enter the name of the route policy to apply to redistributed routes.
Click Add to save the redistribution information.		

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

Table 29: Configure Network

Parameter Name	Description
Mark as Optional Row	Click Optional to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables. See Create a Template Variables Spreadsheet.
Network Prefix *	Enter the network prefix you want EIGRP to advertise in the format of <i>prefix/mask</i> .
Click Add to save the network prefix.	

Configure Advanced Parameters

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

Table 30: Advanced Parameters

Parameter Name	Description
Hold Time (seconds)	Set the interval after which EIGRP considers a neighbor to be down. The local router then terminates the EIGRP session to that peer. This acts as the global hold time. <ul style="list-style-type: none"> • Range: 0 through 65,535 • Default: 15 seconds
Hello Interval (seconds)	Set the interval at which the router sends EIGRP hello packets. <ul style="list-style-type: none"> • Range: 0 through 65,535 • Default: 5 seconds
Route Policy Name	Enter the name of an EIGRP route policy.

Configure Route Authentication Parameters

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

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

Parameter	Option	Description
MD5	MD5 Key ID	Enter an MD5 key ID to compute an MD5 hash over the contents of the EIGRP packet using that value.
	MD5 Authentication Key	Enter an MD5 authentication key to use an encoded MD5 checksum in the transmitted packet.
	Authentication Key	A 256-byte unique piece of information that is used to compute the HMAC and is known both by the sender and the receiver of the message.
Click Add to save the authentication parameters.		



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

Configure Interface Parameters

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

Table 31: Interface Parameters

Parameter Name	Description
Mark as Optional Row	Click Optional to mark this configuration as device-specific. To include this configuration for a device, enter the requested variable values when you attach a device template to a device, or create a template variables spreadsheet to apply the variables.
Interface name	Enter the interface name(s) on which EIGRP should run.
Shutdown	No (the default) enables the interface to run EIGRP. Yes disables the interface.
Click Add to save the interfaces.	

Configure EIGRP Using CLI

Configure EIGRP on Cisco IOS XE Catalyst SD-WAN Devices

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

```

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

```

```

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

```

Example: Advertise EIGRP Routes to OMP

```

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

```

Verify EIGRP Configuration Using CLI

Configuration on Cisco IOS XE Catalyst SD-WAN Devices

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

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

```

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

```

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

```

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

```

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

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

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

Configure Routing Information Protocol (RIPv2) Using the CLI

You can configure RIPv2 using [CLI device templates](#) and [CLI Add-on feature templates](#).

This section provides information about RIPv2 configuration on Cisco IOS XE Catalyst SD-WAN devices.



Note

- Initial VRF routing table and address family submode configurations are required to verify RIPv2 configurations using **show ip protocols** command.
- These commands can be run in any order.

- Configure the RIPv2 routing process.

Enable a RIPv2 routing process and enter router configuration mode:

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

- Configure the RIPv2 VRF-aware support.

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

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

- Specify the RIP version.

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

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

- Configure RIP routes summarization

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

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

- Validate the source IP address.

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

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

- Configure interpacket delay.

Configure interpacket delay for outbound RIP updates, in milliseconds:

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

- Redistribute the routes into the RIP routing process.

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

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

- Filter the RIP-routing updates.

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

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

- Configure the RIP parameters.

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

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

Define a neighboring device with which to exchange routing information:

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

Apply an offset list to routing metrics:

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

Adjust routing protocol timers:

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

- Customize a RIP.

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

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

- Configure a route tag.

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

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

- Configure the traffic.

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

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

Configuration Example

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

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

Verify RIPv2 Configurations Using the CLI

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

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

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

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

```
Gateway of last resort is 10.0.5.13 to network 10.10.10.10
```

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

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

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

```
Gateway of last resort is not set
```

```
10.0.0.14/32 is subnetted, 1 subnets
```

```
R 10.14.14.14 [120/1] via 10.20.25.18, 00:00:18, GigabitEthernet5
```

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

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

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

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

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

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

Configure RIPng Using the CLI

You can configure RIPng using [CLI device templates](#) and [CLI Add-on feature templates](#).

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



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

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


```
Device(config)# ipv6 rip vrf-mode enable
```
 - b. Enable the forwarding of IPv6 unicast datagrams:

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

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



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

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

- 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)#
```

- 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
```

- Configure a route tag.

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

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

- Create an entry in the IPv6 prefix list:

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

- 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]
```

- 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]
```

- Configure the interface.

- Enable the specified IPv6 RIPng routing process on an interface:



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

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

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



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

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

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



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

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

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



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

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

Configuration Example for RIPng

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

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

```

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

```

Verify RIPng Configurations Using the CLI

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

```

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

```



CHAPTER 4

Multicast Overlay Routing

- [Multicast Overlay Routing, on page 75](#)
- [Supported Protocols, on page 77](#)
- [Traffic Flow in Multicast Overlay Routing, on page 79](#)
- [Configure Multicast Overlay Routing, on page 80](#)
- [Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN, on page 102](#)

Multicast Overlay Routing



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

Table 32: Feature History

Feature Name	Release Information	Description
Multicast over L3 TLOC Extension	Cisco IOS XE Release 17.3.2 Cisco vManage Release 20.3.1	This feature enables support for transport location (TLOC) which allows addition of the peers transport to avoid the extra cost of additional IP and allows the use of dynamic load balance across multiple transports.
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. The multicast routing protocols like, IPv4 Multicast, IGMPv3, PIM SSM, PIM ASM, Auto RP and Static RP distribute data (for example, audio/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.

The Cisco IOS XE Catalyst SD-WAN multicast overlay software extends Protocol Independent Multicast Source-Specific Multicast (PIM-SSM) over the Cisco Catalyst SD-WAN overlay using Overlay Management Protocol (OMP). Protocol Independent Multicast Sparse-Mode (PIM-SM) is deployed in the customer VPNs, and the Cisco IOS XE MVPN is used to integrate PIM in customer VPNs and OMP in the overlay. The OMP replicator is used in overlay multicast to optimize the multicast distribution tree across the overlay topology. The Cisco IOS XE Catalyst SD-WAN router supports IGMPv2 and IGMPv3 reports and advertises receiver's multicast interest to remote Cisco Catalyst SD-WAN routers using OMP. Depending on the level of optimization required, the Cisco Catalyst SD-WAN routers join or prune to or from the replicators, and replicators use OMP to relay the join or prune to the Cisco Catalyst SD-WAN router providing overlay connectivity to the PIM-RP or source.

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.

Multicast Overlay Supported Features

- 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
- 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

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

Limitations of Multicast Configuration

Multicast overlay routing does not support the following 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. In case data policy is configured, then only required traffic is matched and not multicast traffic.

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.

PIM

Cisco IOS XE Catalyst SD-WAN overlay multicast supports PIM version 2 (defined in RFC 4601), with some restrictions.

On the service side, the Cisco IOS XE Catalyst SD-WAN software supports native multicast. A router appears as a native PIM router and establishes PIM neighborhood with other PIM routers at a local site. A Cisco IOS XE SD-WAN router supports a directly connected local source, referred as a first hop router (FHR). Receivers residing downstream of a router can join multicast streams by exchanging IGMP membership reports directly with the device, and no other routers are required. Additionally, the Cisco Catalyst SD-WAN router can act as the PIM-RP for the local site.

On the transport side, PIM-enabled Cisco IOS XE Catalyst SD-WAN routers originate multicast service routes (called multicast autodiscover routes), sending them using OMP to the Cisco Catalyst SD-WAN Controllers. The multicast autodiscover routes indicate whether the router is a replicator and the local threshold. Each PIM router also conveys information learned from the PIM join messages sent by local-site multicast-enabled routers, including multicast group state, source information, and RPs. These routes assist Cisco IOS XE Catalyst SD-WAN routers in performing optimized joins across the overlay when joining existing multicast sources.

Cisco IOS XE Catalyst SD-WAN routers support both PIM source-specific mode (SSM) and ASM (Any Source Multicast) mode.

Rendezvous Points

The root of a PIM multicast shared tree resides on a router configured to be a rendezvous point (RP). In the Cisco Catalyst SD-WAN solution, RPs can be Cisco Catalyst SD-WAN routers or non-Cisco Catalyst SD-WAN routers that reside in the local site.

Cisco IOS XE Catalyst SD-WAN supports the following modes of RP discovery:

- Static RP
- Auto-RP
- Auto-RP Proxy

Dynamic RP-group mappings are propagated in the Cisco IOS XE Catalyst SD-WAN solution using Auto-RP. ACLs can be used to control or map certain group ranges to a specific RP. With this information, each PIM router has the ability to forward joins to the correct RP for the group that a downstream IGMP client is attempting to join. Auto-RP updates are propagated to downstream PIM routers if such routers are present in the local site and across the overlay to the remote sites that belong to the same VPN. While using Auto-RP, Replicator Node should be configured as the Auto-RP mapping agent.

Another RP selection model called bootstrap router (BSR) was introduced after Auto-RP in PIM-SM version 2. Auto-RP is a Cisco proprietary protocol, whereas PIM BSR is part of the PIM version 2 specification. BSR

performs similarly to Auto-RP in that it uses candidate routers for the RP function and for relaying the RP information for a group.

Replicators

For efficient use of WAN bandwidth, strategic Cisco IOS XE Catalyst SD-WAN routers can be deployed and configured as replicators throughout the overlay network. Replicators mitigate the requirement for a Cisco Catalyst SD-WAN router with local sources or the PIM-RP to replicate a multicast stream once for each receiver. As discussed above, replicators advertise themselves, using OMP multicast-autodiscover routes, to the Cisco Catalyst SD-WAN Controllers in the overlay network. The controllers then forward the replicator location information to the PIM-enabled Cisco IOS XE Catalyst SD-WAN routers that are in the same VPN as the replicator.

A replicator Cisco IOS XE Catalyst SD-WAN router receives streams from multicast sources, replicates them, and forwards them to other Cisco Catalyst SD-WAN routers with multicast receivers in the same VPN. The details of the replication process are discussed below, in the section Multicast Traffic Flow through the Overlay Network. A replicator is typically a Cisco IOS XE Catalyst SD-WAN router located at a colo-site or another site with a higher-speed connection to the WAN transport network.

Multicast Service Routes

Cisco IOS XE Catalyst SD-WAN routers send multicast service routes to the Cisco Catalyst SD-WAN Controller using OMP. From these routes, the controller processes and forwards joins for requested multicast groups towards the source address or PIM-RP as specified in the original PIM join message that resulted in a Cisco Catalyst SD-WAN router advertising the OMP multicast service route. The source address can be either the IP address of an RP if the originating router is attempting to join the PIM shared tree or the IP address of the actual source of the multicast stream if the originating router is attempting to join the source tree.

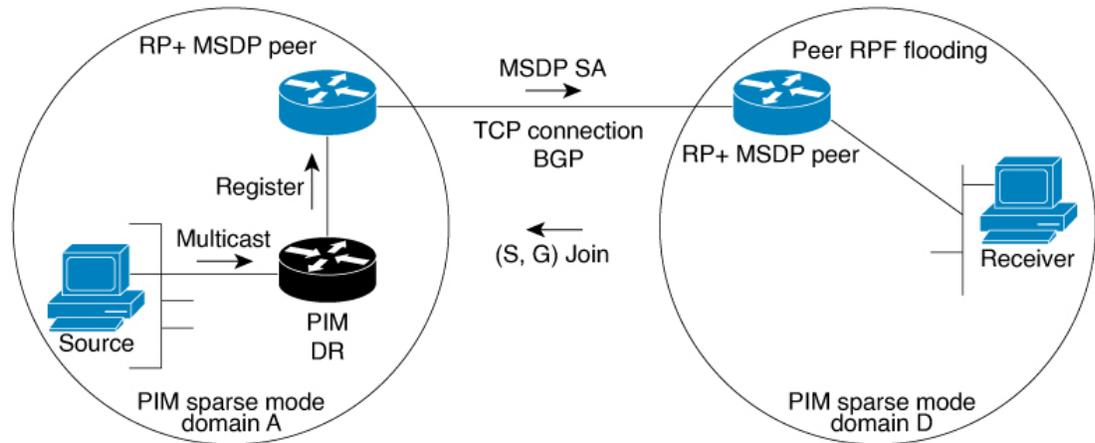
IGMP

Cisco IOS XE Catalyst SD-WAN routers support the Internet Group Management Protocol (IGMP) V2 and V3 protocol. IGMP is used by IPv4 hosts and routers to indicate their interest and in receiving multicast traffic for particular multicast groups. IGMP v3 report is used to indicate interest for a particular multicast group traffic from a specific source. From these membership reports, Cisco IOS XE Catalyst SD-WAN routers originate the corresponding PIM join or OMP service route advertisements.

MSDP

Multicast Source Discovery Protocol (MSDP) is a method of connecting multiple PIM-SM domains, and it is used to discover multicast sources in other PIM domains. When MSDP is configured in a network, rendezvous points (RP) exchange source information with RPs in other domains by maintaining MSDP peer relationships with MSDP-enabled routers in other domains. This peering relationship occurs over a TCP connection. A Cisco IOS XE Catalyst SD-WAN device can be configured as a RP so that it discover active sources outside of its domain.

Figure 4: MSDP



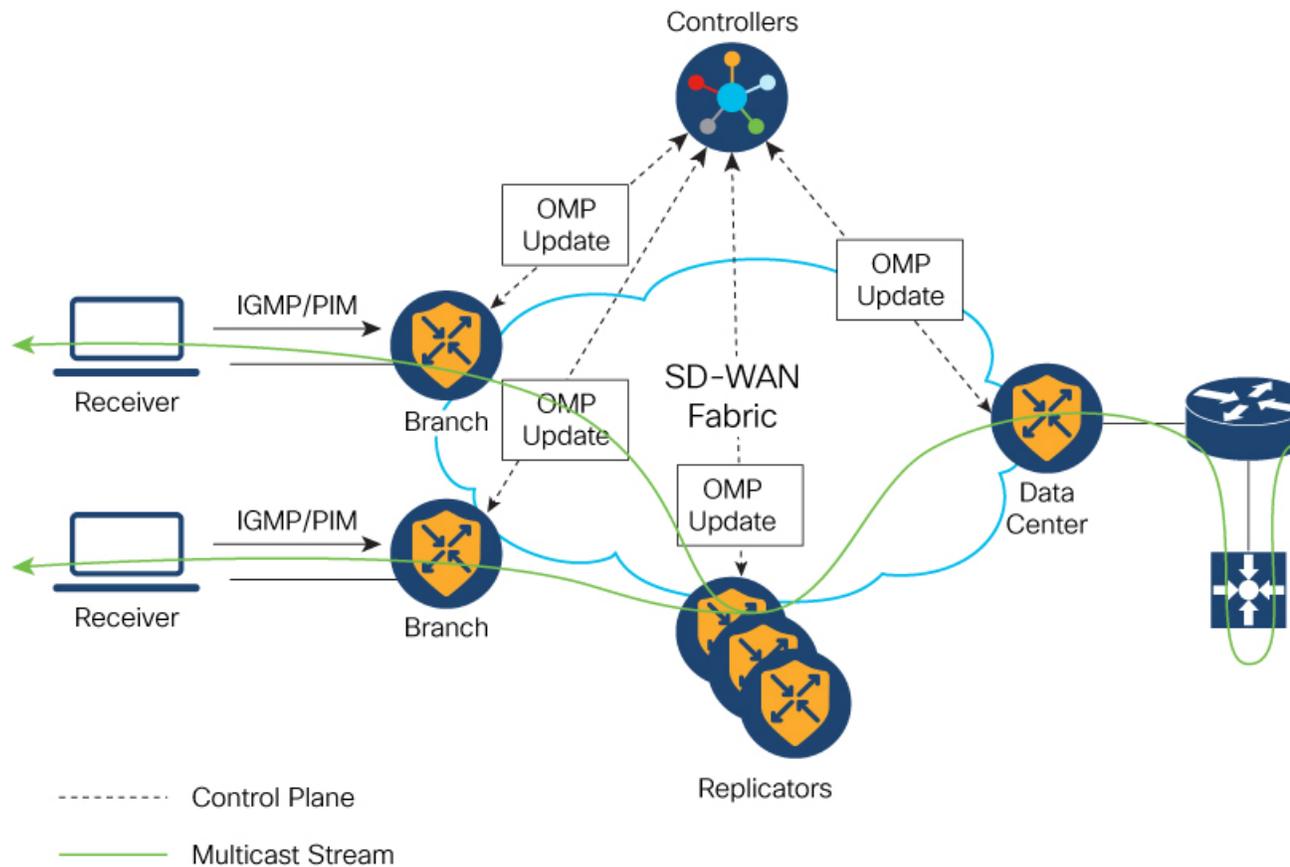
Here is an illustration of the sequence of events that occur when MSDP is implemented:

1. When a PIM designated router (DR) in domain A registers a source with its RP in domain A, the RP sends a Source Active (SA) message to all of its RP MSDP peers. The SA message identifies the source address, the group that the source is sending to, and the address or the originator ID of the RP, if configured.
2. The RP MSDP peer in domain B when it receives the SA message sends the SA message to all of its peers downstream.
3. The RP MSDP peer in domain B checks if there are any receivers of the advertised groups in its domain. If there are receivers in the group, the RP MSDP peer in domain B sends an (S, G) join toward the source. As a result, a connection is established between domain A and domain B. As multicast packets arrive at the RP, they are then forwarded down to the receivers in the RP's domain. When the receivers receiving the multicast traffic learns of the source outside the PIM-SM domain (through the arrival of a multicast packet from the source), it can then send a PIM join toward the source and join source's domain to receive the multicast traffic.

Traffic Flow in Multicast Overlay Routing

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

Figure 5: Multicast Overlay Routing Topology



Configure Multicast Overlay Routing

For any Cisco IOS XE SD-WAN routers to be able to participate in the multicast overlay network, you must configure PIM on those routers.

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.



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. The Cisco IOS XE Catalyst SD-WAN multicast overlay implementation supports IOS XE standard or extended access lists.

2. At least one replicator is mandatory for overlay multicast configuration.

3. You can optionally configure IGMP to allow individual hosts on the service side to join multicast groups within a particular VPN.

Configure Multicast

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

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



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

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
6. Click **Service VPN** in the **Service VPN** section.
7. Click the **Service VPN** drop-down list.
8. Under **Additional VPN Templates**, click **Multicast**.
9. To enable **Local Replicator** on the device, choose **On** (otherwise keep it **Off**).
10. To configure replicator, choose the **Threshold**. (Optional, keep it default if you are not configuring replicator).
11. Save feature template.
12. Attach feature template to device template.
13. 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 the value.

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.

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates > Configuration Groups**.
2. Click ... adjacent to the configuration group name and choose **Edit**.
3. Click **Service Profile**.

4. Click **Add Feature**.
5. From the feature drop-down list, choose **Multicast**.

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

- Protocol Independent Multicast (PIM)
- Internet Group Management Protocol (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 33: 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 34: PIM

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

Field	Description
ACL	<p>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.</p> <p>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.</p> <p>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.</p>
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.

Field	Description
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.
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 35: 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.
Source Address	Enter a source address to join a multicast group.
Add	Click Add to add the IGMP for the group.

Table 36: 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.

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

Configure Multicast Using the CLI

To configure multicast, perform the following:

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

Sample multicast configuration:

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

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.

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

(Optional) You can configure the same standard and extended ACL values in Cisco SD-WAN Manager, which generates the following example configurations:

```
ip pim vrf 1 bsr-candidate Loopback0 32 100 accept-rp-candidate 101
ip pim vrf 1 rp-candidate Loopback0 group-list 27 interval 30 priority 0
```



Note The example configurations are based on the example CLI add-on configuration shown in the procedure.

1. To configure an ACL for multicast, [Create a CLI add-on feature template and attach it to the device template.](#)

This section provides an example configuration.

```
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

```

2. From the **Configuration > Templates** window, choose **Feature**.
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](#).
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.
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.
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:

1. Create a PIM feature template to configure PIM parameters.
2. Optionally, create an IGMP feature template to allow individual hosts on the service side to join multicast groups within a particular VPN. For more information, see [Configure IGMP](#).
3. Optionally, create a multicast feature template to configure a Cisco IOS XE Catalyst SD-WAN to be a multicast replicator.
4. Create a VPN feature template to configure parameters for the VPN that is running PIM.

Create a PIM Feature Template

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



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

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

6. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
7. Click the **Service VPN** drop-down list.
8. Under **Additional VPN Templates**, click **PIM**.
9. 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.
10. 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.
11. 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.
12. Click **Basic Configuration** and configure SSM – On/Off.
13. Configure access list (if already defined).
14. Configure RP option – Auto-RP or static RP.
15. Configure RP Announce settings.
16. Configure the interface name on the service side.
17. 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.

Table 37:

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

Configure Basic PIM

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

Table 38:

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**.

Configure PIM Interfaces

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:

Table 39:

Parameter Name	Description
Name	Enter the name of an interface that participates in the PIM domain, in the format ge slot /port .

Parameter Name	Description
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**.

Rendezvous Point Selection Process by a PIM BSR

Table 40: Feature History

Feature Name	Release Information	Description
Dynamic Rendezvous Point (RP) Selection by a PIM BSR	Cisco IOS XE Catalyst SD-WAN Release 17.5.1a Cisco vManage 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.

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.



Note Cisco Auto-RP cannot co-exist with PIM BSR. Cisco Auto-RP mode must be disabled with spt-only mode.

To avoid a single point of failure, you can configure several candidate BSRs in a PIM domain. A BSR is selected among the candidate BSRs automatically. The BSRs use bootstrap messages to discover which BSR has the highest priority. This router then announces to all PIM routers in the PIM domain that it is the BSR. Any router in the network can be 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 first hop router of the receiver learns about the source, it sends a join message directly to the source, creating a source-based distribution tree from the source to the receiver. This source tree does not include a RP unless the RP is located within the shortest path between the source and the receiver.



Note 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.



Note 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.

Features and Benefits

- 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.

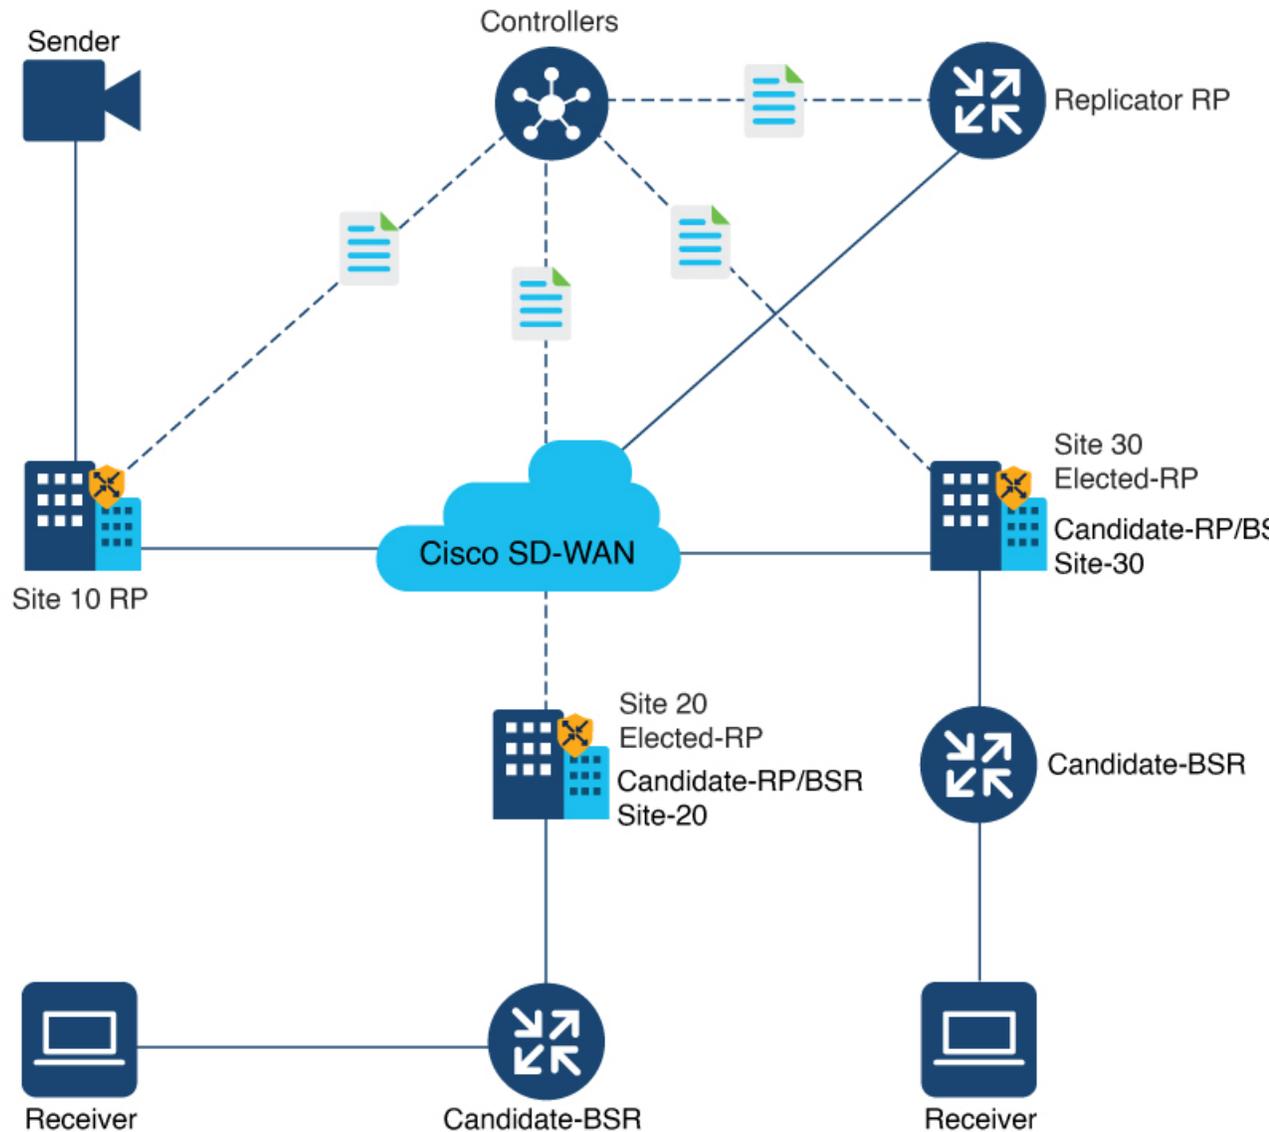
Restrictions for PIM BSR

- 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.

Sample Topology for RP Selection by a PIM BSR

The following is a sample topology for RP selection by a PIM BSR on Cisco IOS XE Catalyst SD-WAN devices.

Figure 6: Topology for PIM BSR Selection



Configure a PIM BSR

Prerequisites for Configuring a BSR Candidate

- Every Cisco Catalyst SD-WAN site must have its own RP.
- SPT-only mode must be enabled on all Cisco Catalyst SD-WAN sites.



Note 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.

Workflow

For a PIM BSR to elect the RP, configure the following in Cisco SD-WAN Manager:

1. Multicast feature template with **SPT Only** set to **On** for the selected Cisco IOS XE Catalyst SD-WAN device.
2. PIM feature template with an interface.
3. RP candidate.
4. BSR candidate.

Configure Shortest-Path Tree (SPT-Only) Mode for a Multicast Feature Template

In Cisco SD-WAN Manager, configure **SPT Only** mode to ensure that the RPs can communicate with each other using the shortest-path tree.



Note When configuring a BSR, configuration of **SPT Only** mode is mandatory.

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



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

3. Click **Add Template**.
4. From the **Select Devices** drop-down list, choose a Cisco IOS XE Catalyst SD-WAN device.
5. Under **Other Templates**, choose **Cisco Multicast**.
6. In the **Template Name** field, enter a name for the template.
7. In the **Description** field, enter a description of the template.
The description can be up to 2048 characters and can contain only alphanumeric characters.
8. Under the **Basic Configuration** section for **SPT Only**, choose **On**.
9. To enable the **Local Replicator** on the device, choose **On** (otherwise keep it set to **Off**).
10. 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).

11. Click **Save**.

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.

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



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

3. Click **Add Template**.
4. From the **Select Devices** drop-down list, choose a Cisco IOS XE Catalyst SD-WAN device.
5. Under **Other Templates**, choose **Cisco PIM**.
6. 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.
7. In the **Description** field, enter a description of the template. The description can be up to 2048 characters and can contain only alphanumeric characters.
8. Click **Interface**.
For information on how to configure a PIM interface, see [Configure PIM](#).
9. Click **New Interface**.
10. In the **Interface Name** field, specify an interface with a value.
11. In the **Query Interval (seconds)** field, the field auto-populates.
12. In the **Join/Prune Interval (seconds)** field, the field auto-populates.
13. Click **Add**.
14. 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.

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



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

3. Edit the PIM feature template that you created by clicking ... and then clicking **Edit**.
4. Click **Basic Configuration**.

5. Click **RP Candidate**.
6. Click **New RP Candidate**.
7. From the **Interface** drop-down list, choose the interface that you used for configuring the PIM feature template.
8. (Optional) In the **Access List** field, if you have configured the access list with a value, add the same value.
9. (Optional) In the **Interval** field, if you have configured the interval with a value, add the same interval value.
10. In the **Priority** field, specify a higher priority on the Cisco IOS XE Catalyst SD-WAN device than on the service-side device.
11. Click **Add**.
12. Click **Update** to save your configuration changes.

Configure the BSR Candidate

1. Repeat Step 1 through Step 4 from the *Configure the RP Candidate* section.
2. Click **BSR Candidate**.
3. In the **BSR Candidate** field, choose the same interface from the drop-down list that you used for configuring the PIM feature template.
4. (Optional) In the **Hash Mask Length** field, specify the hash mask length.
Valid values for hash mask length are from 0 – 32.
5. In the **Priority** field, specify a higher priority on the Cisco IOS XE Catalyst SD-WAN device than on the service-side device.
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.
7. Click **Update** to save your configuration changes.

CLI Configurations for PIM BSR Selection

Configure a BSR Candidate

1. Configure a Cisco IOS XE Catalyst SD-WAN device as a candidate BSR:

```
Device(config)# ip pim vrf 1 bsr-candidate Loopback 99
```



Note The Loopback interface is used only as an example here. Loopback is one of many interface types that can be used for configuring an RP candidate.

2. Use the `show ip pim vrf bsr-router` command to view information about the BSR:

```
Device# show ip pim vrf 1 bsr-router
PIMv2 Bootstrap information
This system is the Bootstrap Router (BSR)
  BSR address: 10.1.10.2 (?)
  Uptime:      15:46:38, BSR Priority: 100, Hash mask length: 32
  Next bootstrap message in 00:00:52
  Candidate RP: 10.1.10.2(Loopback0)
    Holdtime 75 seconds
    Advertisement interval 30 seconds
    Next advertisement in 00:00:18
  Group acl: 27
```

Configure an RP Candidate

1. Configure a Cisco IOS XE Catalyst SD-WAN device as a candidate RP for all multicast groups or selective groups:

```
Device(config)# ip pim vrf 1 rp-candidate Loopback 1 priority 0
```

or

```
Device(config)# ip pim vrf 1 rp-candidate Loopback 1 group-list acl1 priority 0
Device(config)# ip pim vrf 1 rp-candidate Loopback 2 group-list acl2 priority 0
```

2. Use the `show ip pim vrf 1 rp mapping` command to verify the RP mapping assignments:

```
Device# show ip pim vrf 1 rp mapping
PIM Group-to-RP Mappings
This system is a candidate RP (v2)
This system is the Bootstrap Router (v2)

Group(s) 224.0.0.0/4
  RP 10.1.10.2 (?), v2
    Info source: 10.1.10.2 (?), via bootstrap, priority 0, holdtime 75
    Uptime: 15:46:47, expires: 00:00:57
Group(s) 225.0.0.0/8
  RP 10.1.10.2 (?), v2
    Info source: 10.1.10.2 (?), via bootstrap, priority 0, holdtime 75
    Uptime: 15:46:47, expires: 00:00:57
  RP 10.1.10.1 (?), v2
    Info source: 10.1.10.1 (?), via bootstrap, priority 10, holdtime 75
    Uptime: 15:45:45, expires: 00:00:59
Group(s) 226.0.0.0/8
  RP 10.1.10.2 (?), v2
    Info source: 10.1.10.2 (?), via bootstrap, priority 0, holdtime 75
    Uptime: 15:46:55, expires: 00:00:49
  RP 10.1.10.1 (?), v2
    Info source: 10.1.10.1 (?), via bootstrap, priority 10, holdtime 75
    Uptime: 15:46:02, expires: 00:01:09
Group(s) 227.0.0.0/8
  RP 10.1.10.2 (?), v2
    Info source: 10.1.10.2 (?), via bootstrap, priority 0, holdtime 75
    Uptime: 15:47:13, expires: 00:00:59
  RP 10.1.10.1 (?), v2
    Info source: 10.1.10.1 (?), via bootstrap, priority 10, holdtime 75
    Uptime: 15:46:20, expires: 00:00:53
Group(s) 228.0.0.0/8
  RP 10.1.10.2 (?), v2
    Info source: 10.1.10.2 (?), via bootstrap, priority 0, holdtime 75
    Uptime: 15:47:31, expires: 00:01:13
```

Configure a Cisco IOS XE Catalyst SD-WAN Device as SPT-Only

1. Configure a Cisco IOS XE Catalyst SD-WAN device as spt-only:

```
Device(config)# sdwan multicast address-family ipv4 vrf 1  
spt-only
```

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):
```

System IP	SPT-Only Mode	Label	Received		Sent	
			(X,G) Join/Prune	(S,G) Join/Prune	(X,G) Join/Prune	(S,G) 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

Sample Multicast Configuration With SPT-Only

```
Device(config)# sdwan  
Device(config)# multicast  
Device(config)# address-family ipv4 vrf 1  
Device(config)# spt-only  
!
```

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

```

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.

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.

Navigate to the Template Window and Name the Template

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



Note In Cisco vManage Release 20.x.7 and earlier releases, **Device Templates** is titled **Device**.

3. Click **Create Template**.
4. From the **Create Template** drop-down list, choose **From Feature Template**.
5. From the **Device Model** drop-down list, choose the type of device for which you are creating the template.
6. Click **Service VPN** located directly beneath the **Description** field, or scroll to the **Service VPN** section.
7. Click the **Service VPN** drop-down list.
8. Under **Additional VPN Templates**, click **IGMP**.
9. 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.
10. Add interface name on the service side to enable IGMP.
11. (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.
12. (Optional) Enter group address to join and source address.

13. 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.
14. 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 to the left of the parameter field and select the value.

Table 41:

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

Configure Basic 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.

Table 42:

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

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



Note 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, on page 87](#).

Configure MSDP Using a CLI Template

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

This section provides example CLI configurations to configure MSDP.

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

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

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.

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.

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

3. Configure an MSDP Mesh Group.

Configures an MSDP mesh group and indicates that an MSDP peer belongs to that mesh group.



Note You can configure multiple mesh groups per device.

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



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.

Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN

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 vManage 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.

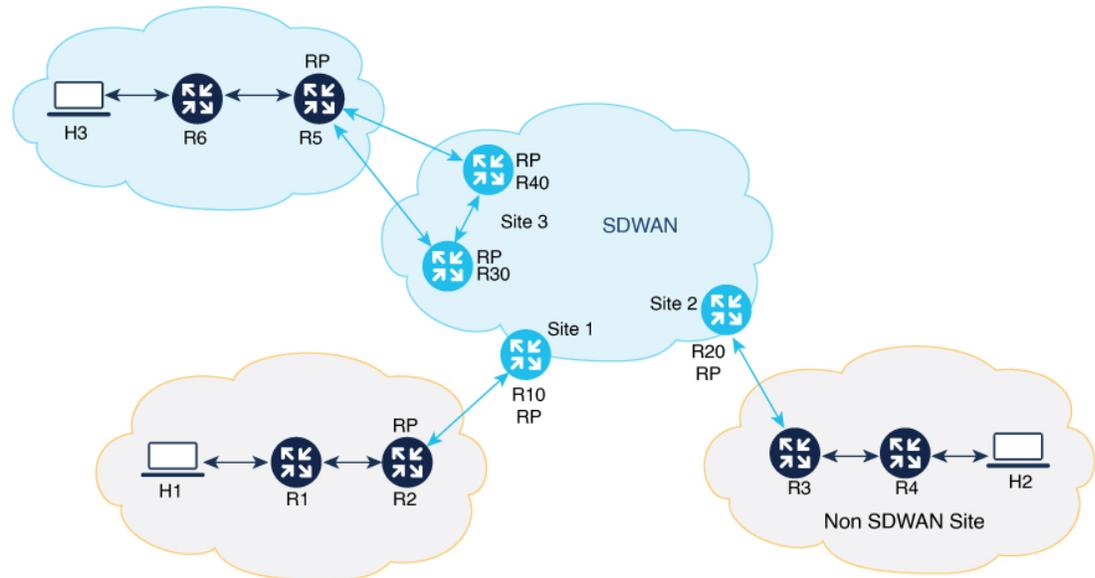
Information About Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN

MSDP facilitates interconnection of multiple Protocol Independent Multicast Sparse-Mode (PIM-SM) domains. 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. For more information about MSDP, see [MSDP, on page 78](#).

From Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, you can configure Cisco IOS XE Catalyst SD-WAN devices for MSDP interoperability with other devices. When Cisco IOS XE Catalyst SD-WAN devices are configured for MSDP interoperability, they convert Source Active (SA) messages received from MSDP peers into OMP routes, and vice-versa.

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 7: MSDP Interoperability



Single Homed Network

In the sample topology, MSDP interoperability is enabled on the Cisco IOS XE Catalyst SD-WAN device, R20, at site 2. R3 is configured as an RP for its PIM domain at the non-SD-WAN site. MSDP peering is established between R3 at the non-SD-WAN site and on R20 at site 2. When source H2 sends traffic to R4, R4 initiates a data registration with R3, and thereafter, R3 sends an MSDP SA message to R20. As MSDP interoperability is enabled in R20, R20 converts the received MSDP SA message to OMP SA routes, and then advertises them to all Cisco IOS XE Catalyst SD-WAN devices located at other sites through the Cisco SD-WAN Controller serving the Cisco IOS XE Catalyst SD-WAN devices. When the Cisco IOS XE Catalyst SD-WAN device R10 at site 1 receives this OMP SA route, R10 converts the OMP SA route into MSDP SA message and advertises the MSDP SA message to its MSDP peer R2 at the non-SD-WAN site. If R2 has any receivers interested in the group advertised in MSDP SA message, then R2 sends a (S,G) join towards the source. As a result, an inter-domain source tree is established across Cisco Catalyst SD-WAN. As multicast packets arrive at R2 (RP), they are then forwarded down its own shared tree to the group members in the RP's domain. R20 withdraws the advertised OMP SA route only when the MSDP SA message expires.

Dual-Homed Network

A dual home network is where there are two Cisco IOS XE Catalyst SD-WAN devices configured for MSDP interoperability. In the dual-homed Cisco Catalyst SD-WAN site 3, MSDP peering must be established between the Cisco IOS XE Catalyst SD-WAN devices R30, R40, and the non-SDWAN device R5. When the source registers its traffic with the RP R5, R5 sends a MSDP SA message to both R30 and R40. When R30 receives the MSDP SA message, it converts the MSDP SA message into OMP SA routes and then advertises to all the Cisco IOS XE Catalyst SD-WAN devices located at other sites, and to R40 within the same, site 3. MSDP SA filter must be configured between R30 and R40 to drop the SA message received from other Cisco IOS XE Catalyst SD-WAN devices and sites through the Overlay Management Protocol (OMP). The Cisco IOS XE Catalyst SD-WAN device R10 at site 1 receives two OMP SA routes for the same Source Group (S, G) and caches them both. R10 then converts the OMP SA route into MSDP SA message and advertises to its MSDP peer R2 at the non-SD-WAN site. If R2 has any receivers interested in the group advertised in MSDP

SA message, then R2 sends a (S,G) join towards the source. As a result, a inter-domain source tree is established across 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, on page 77](#).

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.

Prerequisites for Support for MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN

- For MSDP interoperability to work, you must enable shortest-path tree (SPT) SPT-only mode on a Cisco IOS XE Catalyst SD-WAN device, and the device must be selected as an RP. For more information, see the **Basic Configuration** section in [Configure Multicast Using Configuration Groups, on page 81](#).
- For MSDP interoperability, the peer devices must be set up in a mesh group.
- In a dual-homed setup, configure an MSDP SA filter on a Cisco IOS XE Catalyst SD-WAN device to drop MSDP SA messages from the other Cisco IOS XE Catalyst SD-WAN device.

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



Note You cannot configure the MSDP interoperability using the feature template or the configuration groups in Cisco SD-WAN Manager.

Perform the following tasks to configure MSDP interoperability on Cisco IOS XE Catalyst SD-WAN device:

1. Enable MSDP on Cisco IOS XE Catalyst SD-WAN device. For more information, see [Configure MSDP Using a CLI Template, on page 101](#).
2. Configure MSDP interworking using a CLI template. For more information see [Configure MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN Using a CLI Template, on page 105](#).

Configure MSDP to Interconnect Cisco SD-WAN and Non-SD-WAN Using a CLI Template

Use the CLI templates to configure the MSDP interoperability feature in Cisco Catalyst SD-WAN. For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

1. Enable MSDP on a Cisco IOS XE Catalyst SD-WAN device. For more information, see [Configure MSDP Using a CLI Template, on page 101](#)
2. 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
```

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

The following is a sample output from the `show platform software sdwan multicast remote-nodes vrf/` command, which shows if MSDP interoperability is enabled or not.

```
Device# show platform software sdwan multicast remote-nodes vrf 1
Multicast SDWAN Overlay Remote Nodes (* - Replicator, ^ - Delete Pending):
```

System IP	Mode	SPT-Only	MSDP	I-Work Label	Received		Sent	
					(X,G)	(S,G)	(X,G)	(S,G)
10.16.255.11	No	No	No	1003	0/0	0/0	0/0	1/0
10.16.255.15	No	No	No	1003	1/0	1/0	0/0	0/0
10.16.255.16	Yes	No	No	1003	1/0	1/0	0/0	0/0
10.16.255.21	Yes	Yes	Yes	1003	0/0	0/0	0/0	0/0

Monitor MSDP Configuration to Interconnect Cisco SD-WAN and Non-SD-WAN

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/ Downtime	Reset SA Count	Peer Name
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.

```
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

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 To resolve the problem, do the following:

Solution Check the MSDP peering status between the Cisco IOS XE Catalyst SD-WAN device and the device in non-SD-WAN.

Solution 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 5

Radio Aware Routing



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

Table 44: 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 is a mechanism that uses radio signals to interact with the routing protocol OSPFv3 to signal the appearance, disappearance, and link conditions of one-hop routing neighbors. In large mobile networks, connections to the routing neighbors are interrupted due to distance and radio obstructions. RAR addresses the challenges faced by merging IP routing and radio communications in mobile networks.

- [Supported Devices for RAR, on page 110](#)
- [Prerequisites for RAR, on page 110](#)
- [Benefits of RAR, on page 110](#)
- [Restrictions for RAR, on page 110](#)
- [Information about RAR, on page 111](#)
- [Configure RAR, on page 114](#)

Supported Devices for RAR

The following 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

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.

Benefits of RAR

The Radio Aware Routing feature offers the following 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.

Restrictions for RAR

The Radio Aware Routing feature has the following 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 (HA) technology is not supported.

Information about RAR

Radio-Aware Routing (RAR) is a mechanism that uses radio interfaces to interact with the Open Shortest Path First (OSPFv3) protocol to signal the appearance and link conditions of one-hop routing neighbors.

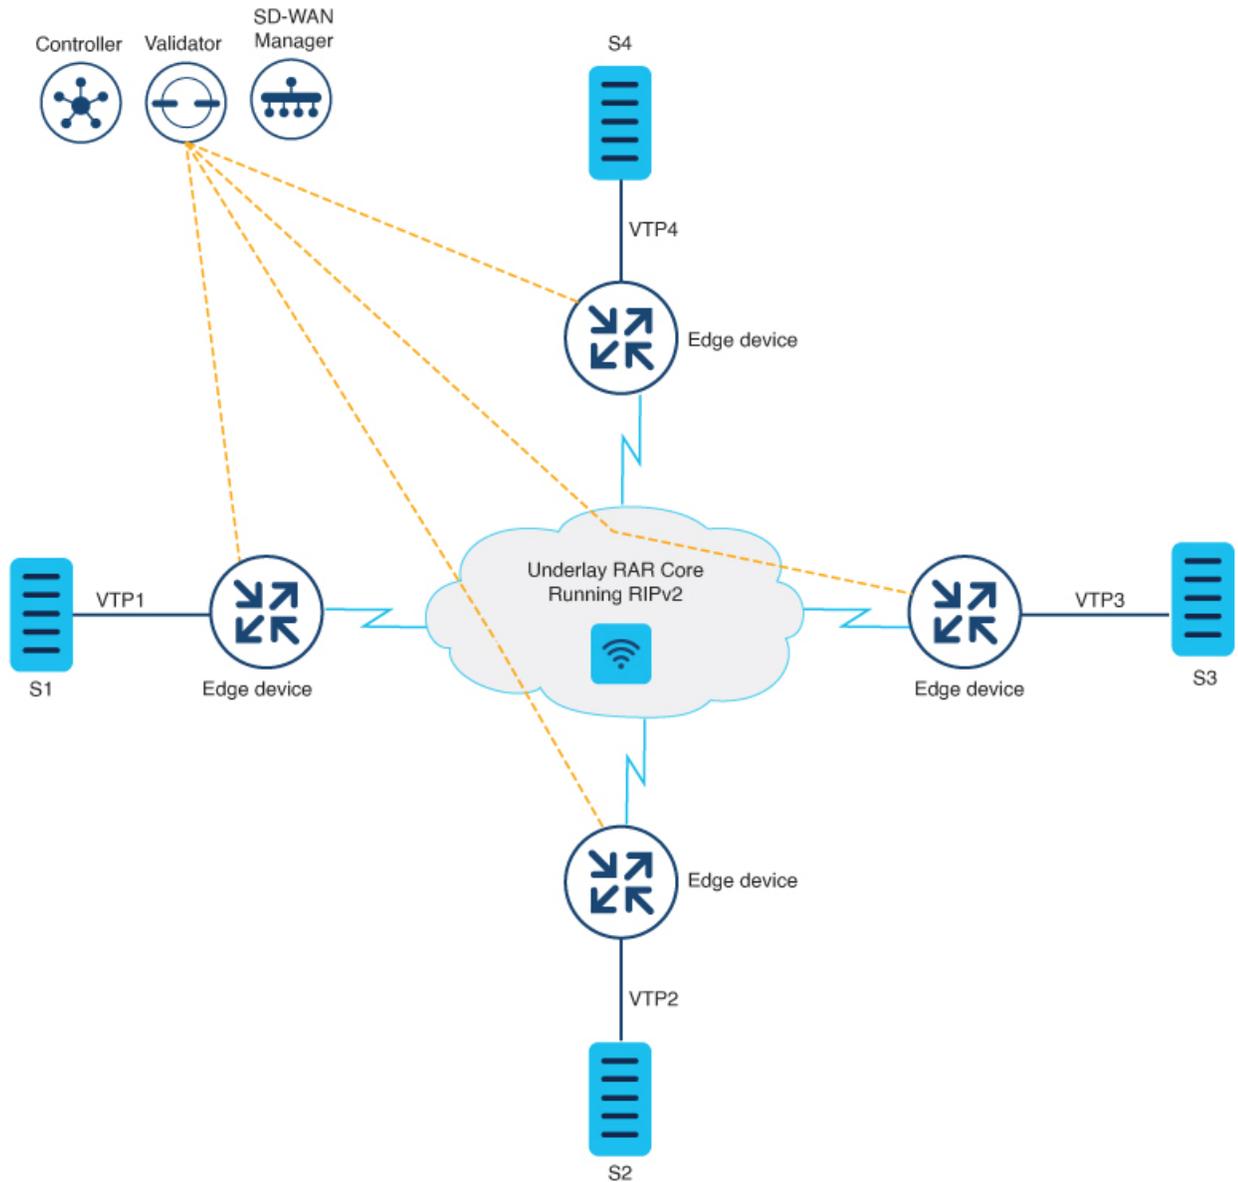
In large mobile network, distance and radio obstructions often interrupt the connections to the routing neighbors. When these signals do not reach the routing protocols, protocol timers are used to update the status of a neighbor. Routing protocols have lengthy timer, which is not recommended in mobile networks.

The connectivity between two Cisco IOS XE Catalyst SD-WAN devices happen over a PPPoE connection using variable bandwidth and a limited buffering. OSPFv3 and EIGRP are the supported routing protocols.

Overview of RAR

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

Figure 8: 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-WAN tunnels are established over the PPPoE-RAR underlay network.

357545

- Cisco SD-WAN Manager, Cisco SD-WAN Controller, and Cisco SD-WAN Validator connect through a radio connection in the deployment scenario.

Mobile Ad Hoc Networks (MANETs)

MANETs for device-to-radio communications address the challenges faced when merging IP routing and mobile radio communications in ad hoc networking applications. MANET-routing protocols provide signaling among MANET routers, including scope-limited flooding and point-to-point delivery of MANET routing protocol signaling in a network.

System Components of RAR

The Radio Aware Routing (RAR) feature is implemented using the MANET (Mobile ad hoc network) infrastructure comprising of different components such as 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

PPPoE extensions are used when the router communicates with the radio. In the Cisco IOS implementation of PPPoE, each individual session is represented by virtual access interface (connectivity to a radio neighbor) on which, QoS can be applied with these 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 the most of the setup to communicate between a router and a radio, VMI addresses the need to manage and translate events that higher layers (example, routing protocols) consume. In addition, VMI operates in the Bypass mode.

In Bypass mode, every Virtual Access Interface (VAI) representing a radio neighbor is exposed to 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 (OSPF) so that the routing protocols can leverage VMI for their optimum 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 helps in aggregating the multiple virtual access interfaces created from PPPoE. VMI presents a single multi access layer 2 broadcast capable interface. The VMI layer handles re-directs unicast routing protocol traffic to the appropriate P2P link (Virtual-Access interface), and replicates any Multicast/Broadcast traffic that needs to flow. Since the routing

protocol communicates to a single interface, the size of the topology database is reduced, without impacting the integrity of the network.

Configure RAR

To configure RAR using Cisco vManage, [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
  ppoe service manet_radio //note: Enter the ppoe 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
ppoe enable group PPPOE_RAR

interface vm1
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 vm1
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 the RAR Feature in Bypass Mode

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



Note Before you begin the RAR configuration, 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 protocol. By default, bypass mode does not appear 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
!
```

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
!
```

Configuration in Bypass Mode

- 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
```

- 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 the RAR Feature in Aggregate Mode

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



Note 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.

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

!

```

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
!
```

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 6

Route Leaking Between VPNs



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.

Table 45: 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	This feature enables you to leak routes bidirectionally between the global VRF and service VPNs. Route leaking allows service sharing and is beneficial in migration use cases because it allows bypassing hubs and provides migrated branches direct access to nonmigrated branches.
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, and 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.

Feature Name	Release Information	Description
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 the following: - 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-WAN devices - 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 VPN	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-WAN devices.

- [Supported Protocols, on page 120](#)
- [Restrictions for Route Leaking and Redistribution, on page 121](#)
- [Information About Route Leaking , on page 122](#)
- [Workflow to Configure Route Leaking Using Cisco SD-WAN Manager, on page 124](#)
- [Configure and Verify Route Leaking Using the CLI, on page 129](#)
- [Configure Route Redistribution Between Global VRF and Service VPNs Using the CLI, on page 135](#)
- [Verify Route Redistribution, on page 138](#)
- [Configure Route Leaking Between Service VPNs Using a CLI Template, on page 139](#)
- [Verify Route-Leaking Configurations Between Service VPNs Using the CLI, on page 140](#)
- [Configure VRRP Tracker for Tracking Leaked Service VPNs Using the CLI, on page 141](#)
- [Verify VRRP Tracking, on page 143](#)
- [Configuration Example for Route Leaking, on page 144](#)

Supported Protocols

The following protocols are supported for route leaking between the global VRF and service VPNs.

- Connected
- Static
- BGP
- OSPF
- EIGRP

The following protocols are the supported destination and source protocols for route redistribution between the service VPNs and global VRF.

Source Protocols

- Connected
- Static
- BGP
- OSPF
- EIGRP

Destination Protocols

- BGP
- OSPF
- EIGRP



Note The EIGRP protocol can be used only on service VPNs and not on the global VRF. Therefore, route leaking is supported only for routes from service VPNs to the global VRF.

Restrictions for Route Leaking and Redistribution

- The EIGRP protocol can be used only on service VRFs and 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.
- Service-side NAT isn't supported with route leaking between the global VRF and service VRFs.
- NAT isn't supported with transport VRF route leaking.
- IPv6 address family is not supported.
- 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.
- Inter-service VRF route leaking on Cisco IOS XE Catalyst SD-WAN devices with multitenancy is not supported.
- Overlay Management Protocol (OMP) routes do not participate in VRF route leaking to prevent overlay looping.
- Route leaking across different devices or sites using export policies in Cisco SD-WAN is not supported.
- Redistribution in EIGRP requires bandwidth, load, reliability, delay, and MTU settings to select the best path.
- Route replicate with **all** keyword is not recommended.
- Route leaking using centralized policy is not supported.

- While configuring route leaking for a VRF, the **route-replicate** command under the **global-address-family ipv4** command shouldn't have the keyword **all** specified as the protocol for the unicast option to prevent route looping.

```
global-address-family ipv4
  route-replicate from vrf <vrf> unicast all
```

- In this example, the keyword **all** should be replaced with specific protocol name as shown here:

```
global-address-family ipv4
  route-replicate from vrf <vrf> unicast connected
```

Information About Route Leaking

Route Leaking Between Global VRF and Service VPNs

The Cisco Catalyst SD-WAN solution lets you segment the network using 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).



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 VRF Route Leaking

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.9.1a, Cisco vManage Release 20.9.1.

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

To resolve routing-scalability challenges introduced when you use Cisco SD-WAN Controllers, you can leak routes between the VRFs at the edge device.

To configure the inter-service VRF route leaking feature using Cisco SD-WAN Manager, see [Configure Route Leaking Between Service VRFs](#).

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

- 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 a 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

- **Service Provider Central Services:** SP Central services under MPLS can be directly accessed without having to duplicate them for each VPN. This makes accessing central services easier and 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 VRFs is used where the VRF traffic goes through a zone-based firewall on the same branch router while being PCI compliant.

How Route Preference is Determined

If a route is replicated or leaked between the global VRF and service VPNs, the following rule determines the route preference.

For a device that receives route from two sources where both these routes use the same source VRFs and one of the routes is replicated, the non-replicated route is preferred.

If the mentioned rule doesn't apply, the following rules determine the route preference in this sequence:

1. Prefer the route with smaller administrative distance.
2. Prefer the route with smaller default administrative distance.
3. Prefer a non-replicated route over a replicated route.
4. Compare original VRF-names. Prefer the route with the lexicographically smaller VRF-name.
5. Compare original subaddress families. Prefer unicast routing over multicast routing.
6. Prefer the oldest route.

Workflow to Configure Route Leaking Using Cisco SD-WAN Manager

1. Configure and enable the Localized Policy and attach the Route Policy.
2. Configure and enable the Route Leaking feature between Global and Service VPN.
3. Configure and enable the Route Leaking feature between Service VPNs.
4. Attach the Service Side VPN Feature Template to the Device Template.

Configure Localized Route Policy

Configure Route Policy

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Policies**.
2. Select **Localized Policy**.
3. From the **Custom Options** drop-down, under Localized Policy, select **Route Policy**.
4. Click **Add Route Policy**, and select **Create New**.
5. Enter a name and description for the route policy.

6. In the left pane, click **Add Sequence Type**.
7. In the right pane, click **Add Sequence Rule** to create a single sequence in the policy. Match is selected by default.
8. Select a desired protocol from the **Protocol** drop-down list. The options are: IPv4, IPv6, or both.
9. Click a match condition.
10. On the left, enter the values for the match condition.
11. On the right enter the action or actions to take if the policy matches.
12. Click **Save Match and Actions** to save a sequence rule.
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**.
14. Click **Save Route Policy**.

Add the Route Policy

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

Attach the Localized Policy to the Device Template



Note The first step in utilizing the Localized Policy that was created previously is to attach it to the device template.

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

2. Click **Device Templates** and select the desired template.
3. Click **...**, and click **Edit**.
4. Click **Additional Templates**.
5. From the **Policy** drop-down, choose the **Localized Policy** that is created.
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.

7. Click **Next** and then **Configure Devices**.
8. Wait for the validation process and push configuration from Cisco SD-WAN Manager to the device.

Configure and Enable Route Leaking between Global and Service VPNs

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. To configure route leaking, click **Feature Templates**.



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

Do one of the following:

- To create a feature template:
 - a. Click **Add Template**. Choose a device from the list of devices. The templates available for the selected device display in the right pane.
 - b. Choose the **Cisco VPN** template from the right pane.



Note Route leaking can be configured on service VPNs only. Therefore, ensure that the number you enter in the **VPN** field under **Basic Configuration** is one of the following: 1—511 or 513—65527.

For details on configuring various VPN parameters such as basic configuration, DNS, Virtual Router Redundancy Protocol (VRRP) tracking, and so on, see [Configure a VPN Template](#). For details specific to the route leaking feature, proceed to Step c.

- c. Enter Template Name and Description for the feature template.
- d. Click **Global Route Leak** below the **Description** field.
- e. To leak routes from the global VRF, click **Add New Route Leak from Global VPN to Service VPN**.

1. In 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.
 2. In the **Route Policy Leak from Global to Service** drop-down list, choose **Global**. Next, choose one of the available route policies from the drop-down list.
 3. For the **Redistribute to protocol (in Service VPN)** field, click **Add Protocol**.
In the **Protocol** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
In the **Redistribution Policy** drop-down list, choose **Global**. Next, choose one of the available redistribution policies from the drop-down list.
 4. Click **Add**.
- f. To leak routes from the service VPNs to the global VRF, click **Add New Route Leak from Service VPN to Global VPN**.
1. In the **Route Protocol Leak from Service to Global** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
 2. In the **Route Policy Leak from Service to Global** drop-down list, choose **Global**. Next, choose one of the available route policies from the drop-down list.
 3. For the **Redistribute to protocol (in Global VPN)** field, click **Add Protocol**.
In the **Protocol** drop-down list, choose **Global** to choose a protocol. Otherwise, choose **Device-Specific** to use a device-specific value.
In the **Redistribution Policy** drop-down list, choose **Global**. Next, choose one of the available redistribution policies from the drop-down list.
 4. Click **Add**.
- g. Click **Save/Update**. The configuration does not take effect till the feature template is attached to the device template.
- h. To redistribute the leaked routes using Cisco SD-WAN Manager, use [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.

- To modify an existing feature template:
 - a. Choose a feature template you wish to modify.

- b. Click ... next to the row in the table, and click **Edit**.
- c. Click **Global Route Leak**.
- d. 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**.
The update route leak dialog box appears.
- e. Perform all operations from Step d of creating a feature template.
Perform all operations from Step c of creating a feature template.
- f. Click **Save Changes**.
- g. Click **Update**.

**Note**

- The configuration does not take effect till the Service VPN feature template is attached to the device template.

Configure Route Leaking Between Service VPNs

Minimum supported release: Cisco vManage Release 20.9.1

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.
3. Navigate to the **Cisco VPN** template for the device.

**Note**

To create a **VPN** template, see [Create VPN Template](#)

4. Click **Route Leak**.
5. Click **Route Leak between Service VPN**.
6. Click **Add New Inter Service VPN Route Leak**.
7. From the **Source VPN** drop-down list, choose **Global** to configure the service VPN from where you want to leak the routes. Otherwise, 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.)
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.

- 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.

This field is disabled if no route policies are available.

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

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.

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

This field is disabled if no route policies are available.

- Click **Add**.
- Click **Save**.

Attach the Service Side VPN Feature Template to the Device Template

- From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
- Click **Device Templates** and select the desired template.
- Click **...**, and click **Edit**.
- Click **Service VPN**.
- Click **Add VPN**. Select the Service VPN feature template listed in the Available VPN Templates pane. Click right-shift arrow and add the template to Selected VPN Templates list.
- Click **Next** once it moves from the left (Available VPN Templates) to the right side (Selected VPN Templates).
- Click **Add**.
- Click **Update**.
- Click **Next** and then **Configure Devices**.
- Finally, wait for the validation process and push configuration from Cisco SD-WAN Manager to the device.

Configure and Verify Route Leaking Using the CLI

Example: Leak Routes between Global VRF and Service VPNs

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 Configuration



Note In the output, leaked routes are represented by a + sign next to the route leaked. Example: C+ denotes that a connected route was leaked.

```
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
```

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.



Note In the output, leaked routes are denoted by a + sign next to the route leaked. Example: C+ denotes that a connected route was leaked.

```
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
```

Example: Filter Routes 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
  !
!
```

Verify Configuration



Note In this output, leaked routes are denoted by a + sign next to the route leaked. Example: C+ denotes that a connected route was leaked.

```
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
```

To monitor leaked routes, use the **show ip cef** command. The output shows replicated or 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, hwc]
path 7F6B3D9E7B70, share 1/1, type receive, for IPv4
receive for GigabitEthernet5
output chain:
receive
```

Example: Redistribute BGP Route into OSPF and EIGRP Protocols

These examples show how to replicate BGP route from global VRF to service VRF.

```
Device#config-transaction
Device(config)# vrf definition 2
```

```
Device(config-vrf)# address-family ipv4
Device(config-ipv4)# route-replicate from vrf global unicast bgp 1
Router(config-ipv4)# commit
```

Configure to Redistribute BGP Routes in Global VRF to EIGRP in Service VRF



Note The redistribution of BGP routes into other protocols is supported only if the `bgp redistribute-internal` configuration is present in the BGP route.

```
Device#config-transaction
Device(config)# router eigrp test
Device(config-router)# address-family ipv4 unicast vrf 2 autonomous-system 100
Device(config-router-af)# topology base
Device(config-router-af-topology)# redistribute vrf global bgp 1 metric 10000 100 200 1
1500
Device(config-ipv4)# commit
```

* Here we are redistributing BGP routes in global VRF to EIGRP in VRF 2.
 * Routes replication must be done before doing inter VRF redistribution.

Verify Configuration

View BGP Route is Present in Global VRF Before Configuring

```
Device#show ip route bgp
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

10.0.0.0/9 is subnetted, 1 subnets
B 172.16.255.1 [200/20] via 10.1.15.14, 00:00:25
Device#
```

* We have a BGP route in the global VRF.

View BGP Route is not Present in Service VRF Before Configuring

Use the `show ip route vrf <vrf id> [protocol]` command to view the BGP route in the service VRF table.

```
Device#show ip route vrf 2 bgp

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

Device#

* We do not have any BGP route in VRF 2.

View BGP Route After Configuring

Use the **show running config** [*configuration-hierarchy*] | **details** command to verify if the replication configuration exists.

```
Device#show running-config | section vrf definition 2
vrf definition 2
  rd 1:1
  route-target export 1:1
  route-target import 1:1
  !
  address-family ipv4
    route-replicate from vrf global unicast bgp 1
  exit-address-family
Device#
```

* We have successfully applied the route-replicate configuration.

* In our example we are replicating bgp 1 routes from global VRF to VRF 2.

View BGP Route From Global VRF is Replicated into Service VRF After Configuring

Use the **show ip route vrf** <*vrf id*> [**protocol**] command to view the BGP route in the service VRF table.

```
Device#show ip route vrf 2 bgp
```

```
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

```
      10.0.0.0/9 is subnetted, 1 subnets
B    +   172.16.255.1 [200/20] via 10.1.15.14, 00:04:01
Device#
```

* After route replication, we can see that the BGP route in the global VRF has been replicated into VRF 2.

* + sign indicates replicated routes.

View EIGRP Configuration Without BGP Redistribution Information

```
Device#show running-config | section router eigrp
router eigrp test
!
address-family ipv4 unicast vrf 2 autonomous-system 100
!
topology base
exit-af-topology
network 10.0.0.0
exit-address-family
Router#
```

View EIGRP Topology Table

Use the **show eigrp address-family ipv4 vrf<vrf-num>topology** command to view the BGP route in the service VRF table.

```
Device#show eigrp address-family ipv4 vrf 2 topology
EIGRP-IPv4 VR(test) Topology Table for AS(100)/ID(10.10.10.2)
Topology(base) TID(0) VRF(2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.0.0.0/8, 1 successors, FD is 1310720
   via Connected, GigabitEthernet2
```

Device#

* EIGRP 100 is running on VRF 2.

View EIGRP Route After BGP Redistribution

Use the **show eigrp address-family ipv4 vrf<vrf-num>topology** command to view the BGP route is redistributed into the EIGRP protocol.

```
Device#show eigrp address-family ipv4 vrf 2 topology
EIGRP-IPv4 VR(test) Topology Table for AS(100)/ID(10.10.10.2)
Topology(base) TID(0) VRF(2)
Codes: P - Passive, A - Active, U - Update, Q - Query, R - Reply,
       r - reply Status, s - sia Status

P 10.10.10.0/8, 1 successors, FD is 1310720
   via Connected, GigabitEthernet2
P 172.16.0.0/12, 1 successors, FD is 131072000
   via +Redistributed (131072000/0)
```

-Device#

* BGP route has been redistributed into EIGRP.

Configure Route Redistribution Between Global VRF and Service VPNs Using the CLI

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



Note 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
```

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
```

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.

- BGP protocol:

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

- 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]
```

- 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.

- BGP protocol:

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

- 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]
```

- 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
```

Verify Route Redistribution

Example 1:

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
vm5#
```

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
```

Example 2:

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

```

Configure Route Leaking Between Service VPNs Using a CLI Template

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.9.1a

For more information about using CLI templates, see [CLI Add-on Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

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

- 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]

```

- 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

```

Verify Route-Leaking Configurations Between Service VPNs Using the CLI

Minimum supported release: Cisco IOS XE Catalyst SD-WAN Release 17.9.1a

The following is a sample output from the **show ip route vrf** command that shows the routes that are replicated for the redistribution to VRF 2:

```

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

```

The following is a sample output from the **show ip cef vrf** command that shows the replicated routes from VRF 1:

```
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, hwc]
  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, hwc]
  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
```

Configure VRRP Tracker for Tracking Leaked Service VPNs Using the CLI

Configure a track.

1. 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}
```

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

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

3. Return to privileged EXEC mode

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

Configure VRRP version 2 (VRRPv2).

1. Configure an interface type such as, Gigabit Ethernet.

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

2. Associate a VRF instance with the Gigabit Ethernet interface.

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

3. Set a primary IP address for the Gigabit Ethernet interface.

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

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

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

5. Create a VRRP group and enter VRRP configuration mode.

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

6. Enable the support of VRRP version 2 simultaneously with VRRP version 3.

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

7. Set the priority level for VRRP.

```
Device(config-if-vrrp)# priority level
```

8. Configure interface list tracking as a single entity.

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

9. 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
```

10. Specify a primary IP address for VRRP.

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

Configure a VRF.

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

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

2. Set an address family IPv4 in vrf configuration mode.

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

3. 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
```

Verify VRRP Tracking

Example 1:

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
```

Example 2:

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 track 1
Track 1
  IP route 209.165.200.225 209.165.200.236 reachability
  Reachability is Down (no ip route)
  1 change, last change 1w1d
  VPN Routing/Forwarding table "vrrp"
  First-hop interface is unknown
rtr3#
```

Example 3:

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.

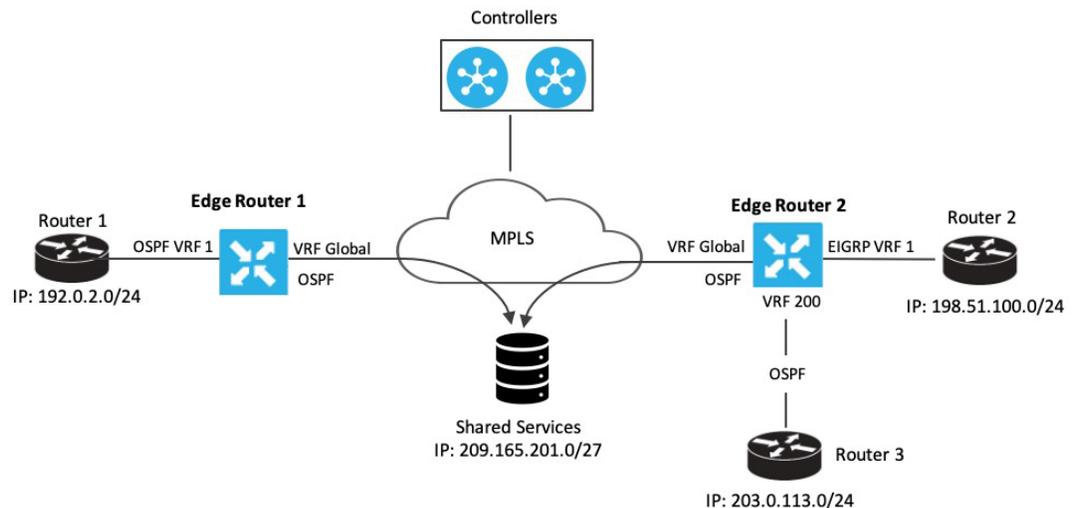
```
Device# show running-config interface GigabitEthernet 4
Building configuration...
Current configuration : 234 bytes
!
interface GigabitEthernet4
ip address 172.16.0.1 255.255.255.0
negotiation auto
vrrp 7 address-family ipv4
  priority 200
  vrrpv2
  track 5 decrement 5β-----priority decrement
 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

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 and clients who reside in one VPN can reach matching prefixes that exist 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 sits behind the Edge Router 2 on network that has EIGRP in VRF 1. Router 3 also sits behind Edge Router 2 and has OSPF running in VRF 200.



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
```




CHAPTER 7

BFD for Routing Protocols in Cisco Catalyst SD-WAN



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.

Table 46: Feature History

Feature Name	Release Information	Description
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 protocols in the Cisco Catalyst SD-WAN solution. BFD provides a consistent failure detection method to detect forwarding path failures at a uniform rate, therefore enabling faster reconvergence time.

- [Information About BFD for Routing Protocols](#), on page 148
- [Configure BFD for Routing Protocols](#), on page 151
- [Configure BFD for Routing Protocols Using CLI](#), on page 157
- [Monitor and Verify BFD Configuration](#), on page 159
- [Troubleshoot Common BFD Errors](#), on page 160

Information About BFD for Routing Protocols

Overview of BFD

In enterprise networks, the convergence of business-critical applications onto a common IP infrastructure is becoming more common. Given how critical data is, these networks are typically constructed with a high degree of redundancy. While such redundancy is desirable, its effectiveness is dependent upon the ability of individual network devices to quickly detect failures and reroute traffic to an alternate path. The detection times in existing protocols are typically greater than one second, and sometimes much longer. For some applications, this duration is too long to be useful. This is where Bi-directional Forwarding Detection (BFD) comes in.

BFD provides rapid failure detection times between forwarding engines, while maintaining low overhead. It also provides a single, standardized method of link/device/protocol failure detection at any protocol layer and over any media, thus enabling faster reconvergence of business-critical applications.

Benefits of Configuring BFD for Routing Protocols

- Fast failure detection times for all media types, encapsulations, topologies, and routing protocols
- Faster reconvergence of applications
- Consistent method of failure detection

How BFD Works in Cisco Catalyst SD-WAN

With the introduction of this feature, the Cisco Catalyst SD-WAN solution now has two types of BFDs that are distinct features that work independently without conflict.

- **BFD Support for Cisco Catalyst SD-WAN Routing Protocols (Legacy BFD):** This feature is termed as legacy BFD because is already available for Cisco IOS XE and is being extended to the Cisco Catalyst SD-WAN solution starting Cisco IOS XE Catalyst SD-WAN Release 17.3.1a.
- **Cisco Catalyst SD-WAN BFD:** This feature is specific to overlay BFD, which is an existing feature in Cisco Catalyst SD-WAN.

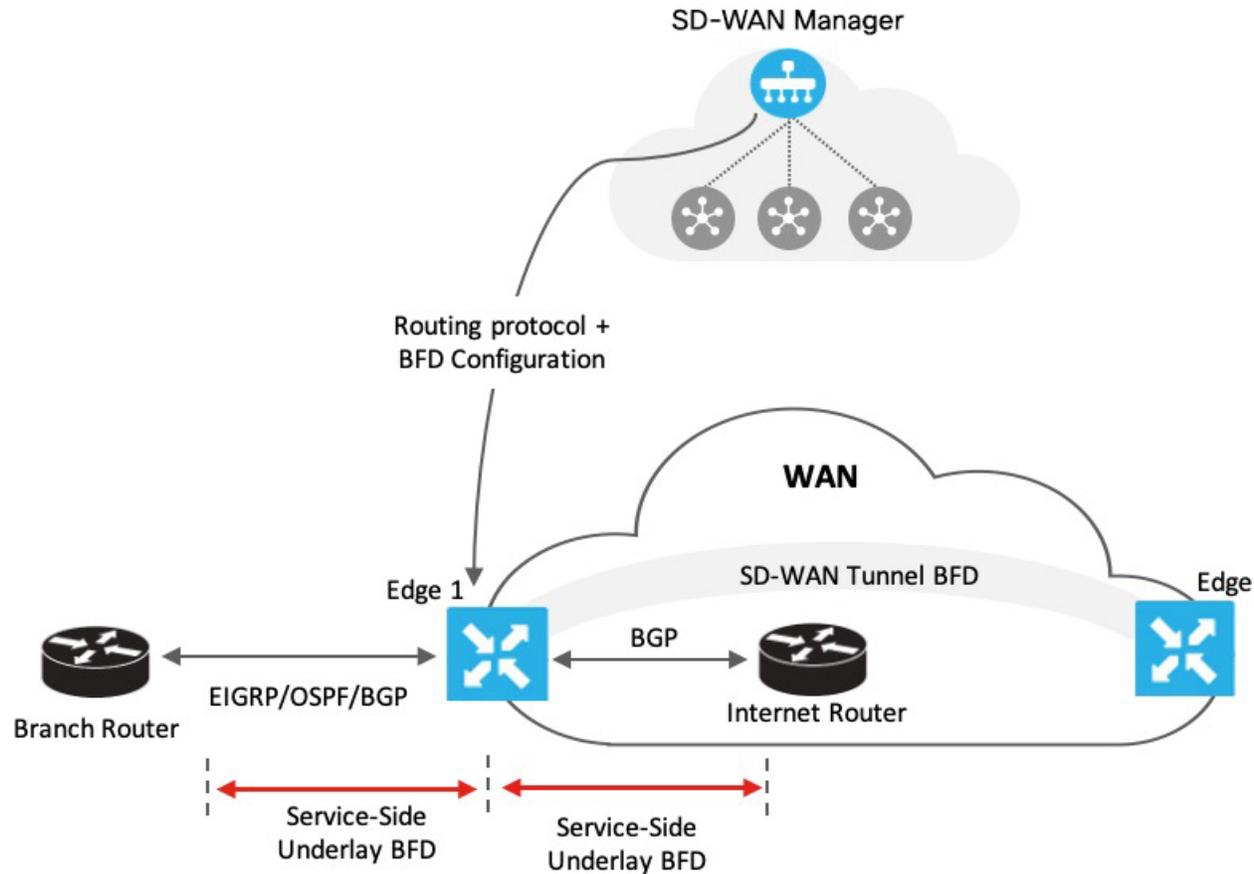
For more information on Cisco Catalyst SD-WAN BFD, see [Cisco Catalyst SD-WAN BFD](#).

Table 47: 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
<ul style="list-style-type: none"> • Runs on both, transport-side and service-side interfaces • The following protocols can be registered: BGP, OSPF, and EIGRP <ul style="list-style-type: none"> • 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 	<ul style="list-style-type: none"> • 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



Note The rest of this document is specific to BFD for Routing Protocols in Cisco Catalyst SD-WAN (legacy BFD).



As represented in the image, BFD is configured for a routing protocol through Cisco SD-WAN Manager. Cisco SD-WAN Manager then pushes this configuration to the edge router. In this example, let's assume that OSPF is configured to receive forwarding path detection failure messages from BFD. If there's a physical link failure, OSPF is prompted to shut down its neighbors and restore any routing information it may have advertised to or received from its remote neighbors.

Similarly, the router, Edge 1 is connected to the internet router through its transport interface. BFD is configured for BGP between the transport side of Edge 1 and the internet router. Here, BFD detects the health of the connection and reports any failures.

Supported Protocols and Interfaces

Supported Protocols

The following routing protocols in Cisco Catalyst SD-WAN can be configured to receive forwarding path detection failure messages from BFD:

- BGP
- EIGRP

- OSPF and OSPFv3

Supported Interfaces

- GigabitEthernet
- TenGigabitEthernet
- FiveGigabitEthernet
- FortyGigabitEthernet
- HundredGigabitEthernet
- SVI
- Subinterfaces

Limitations and Restrictions

The following restrictions apply to Cisco IOS XE Catalyst SD-WAN devices in controller mode.

- Only single-hop BFD is supported.
- BFD is not supported for static routes.
- To change the BFD session modes between software mode and hardware mode, you need to remove all existing BFD configuration and reconfigure it.
- BFD is only supported for BGP, EIGRP, OSPF, and OSPFv3.
- BFD for routing protocols in Cisco Catalyst SD-WAN cannot be monitored through Cisco SD-WAN Manager. Use CLI show commands for monitoring BFD for Cisco Catalyst SD-WAN routing protocols.
- Once a BFD session is established, BFD session modes (echo to no echo, and vice-versa; or software to hardware, and vice-versa) don't update immediately after changing the BFD template parameters in Cisco SD-WAN Manager. The BFD mode change takes effect only after the session flaps at least once.

Configure BFD for Routing Protocols

Cisco SD-WAN Manager does not provide an independent template to configure BFD for routing protocols. However, supported protocols can be registered or deregistered to receive BFD packets by adding configurations using the CLI add-on template in Cisco SD-WAN Manager. Use the CLI add-on template to configure the following:

- Add a single-hop BFD template with parameters such as timer, multiplier, session mode, and so on.
- Enable the BFD template under interfaces. Only one BFD template can be added per interface.
- Enable or disable BFD for the supported routing protocols. The configuration to enable or disable BFD is different for each of the supported routing protocols: BGP, EIGRP, OSPF, and OSPFv3.

Enable BFD for Routing Protocols

Configure BFD for Service-Side BGP

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



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

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template and to enable BFD for service-BGP as shown in the following example.

```
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
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these parameters is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default). Once created, the BFD template is enabled under an interface (GigabitEthernet1, in this example).



Note To modify a BFD template enabled on an interface, you need to remove the existing template first, modify it, and then enable it on the interface again.



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

7. Click **Save**.
8. [Attach Feature Template to Device Template](#)



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

9. [Attach the device template to the device.](#)

Configure BFD for Transport-Side BGP

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



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

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template and to enable BFD for transport-BGP as shown in the following example:

```
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
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these parameters is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default). Once created, the BFD template is enabled under an interface (GigabitEthernet1, in this example). In this example, GigabitEthernet1 is also the source of the SD-WAN tunnel. Allowing service under the tunnel interface of GigabitEthernet1 ensures that BGP and BFD packets pass over the tunnel.



Note To modify a BFD template enabled on an interface, you need to remove the existing template first, modify it, and then enable it on the interface again.



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

7. Click **Save**.
8. [Attach Feature Template to Device Template](#)



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

9. [Attach the device template to the device](#).

Configure BFD for Service-Side EIGRP

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



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

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
6. Enter the CLI configuration to add a single-hop BFD template 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
```

Understanding the CLI Configuration

In this example, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default).

Once created, the BFD template is enabled under an interface (GigabitEthernet5, in this example).



Note To modify a BFD template enabled on an interface, you first need to remove the existing template, modify it, and enable it on the interface again.

7. Click **Save**.
8. [Attach Feature Template to Device Template](#)



Note For the configuration to take effect, the device template must have an EIGRP feature template attached to it.

9. [Attach the device template to the device](#).

Configure BFD for Service-Side OSPF and OSPFv3

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



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

3. Click **Add Template**.
4. Choose a device from the device list.
5. Choose the **CLI Add-on Template** under **Other Templates**.
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
```

Understanding the CLI Configuration

In these examples, a single hop BFD template is created specifying the minimum and maximum interval and the multiplier. Specifying these is mandatory. In addition, you have the option to also specify other BFD parameters such as echo mode (enabled by default), and BFD dampening (off by default).

Once created, the BFD template is enabled under an interface (GigabitEthernet5, in this example).



Note To modify a BFD template enabled on an interface, you first need to remove the existing template, modify it, and enable it on the interface again.

7. Click **Save**.
8. [Attach the CLI Add-on Template with this configuration to the device template.](#)



Note For the configuration to take effect, the device template must have an OSPF feature template attached to it.

9. [Attach the device template to the device.](#)

Attach Feature Template to Device Template

After creating a CLI add-on template to enable BFD, attach the template to the device template for the configuration to take effect. Follow this procedure to attach the configuration to a device template. Ensure that the device template you attach the feature template to already has the relevant feature template (BGP, OSPF, EIGRP) attached to it.

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



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

3. Click **Create Template** and choose **From Feature Template** from the drop-down options.
4. From the **Device Model** drop-down options, choose a device. Enter a name and description for the template.
5. Click **Create**.
6. Click **Additional Templates**.
7. In the **CLI Add-on Template** field, choose the CLI add-on template you configured to enable BFD for routing protocols.
8. Click **Create**.

Next: [Attach device template to device](#)

Configure BFD for Routing Protocols Using CLI

To configure BFD for BGP, EIGRP, OSPF, and OSPF3 using device CLI, follow the steps in this topic.

Create BFD Template

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
```



Note The CLI configuration for creating a BFD template remains the same irrespective of the protocol you configure it for.

Enable BFD for Service-Side BGP

This example shows that BGP is configured, BFD is enabled on the interface under VRF 1, and then on 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
```

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

```

Enable BFD for EIGRP

This example shows that EIGRP is configured, BFD is enabled on the interface under VRF 1, and then on 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
!

```

Enable BFD for OSPFv3

This example shows that OSPFv3 is configured, BFD is enabled on the interface under VRF 1, and then on service-side EIGRP.

```

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

```

Monitor and Verify BFD Configuration

This sections provides a list of commands that you can run to verify your BFD configuration.

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

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
```

CONTROLLER		PEER		SITE		DOMAIN		PEER		PRIV	
GROUP	PEER	PEER	PEER	ID	ID	PRIVATE IP	PUB	IP	IP	UP	DOWN
TYPE	PROT	SYSTEM	IP	ID	ID	LOCAL	COLOR	PROXY	STATE	UPTIME	PORT
		PUBLIC	IP			PORT					ID
vsmart	dtls	172.16.255.19	100	1	10.0.5.19	12355	lte	No	up	0:12:45:44	12355
		10.0.5.19									0
vsmart	dtls	172.16.255.20	200	1	10.0.12.20	12356	lte	No	up	0:15:59:45	12356
		10.0.12.20									0
vmanage	dtls	172.16.255.22	200	0	10.0.12.22	12346	lte	No	up	0:15:59:45	12346
		10.0.12.22									0
		< ----	up								

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
.....
```




CHAPTER 8

Cisco Catalyst SD-WAN BFD



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.

Table 48: Feature History

Feature Name	Release Information	Description
Automatically Suspend Unstable Cisco Catalyst SD-WAN BFD Sessions	Cisco IOS XE Catalyst SD-WAN Release 17.10.1a Cisco Catalyst SD-WAN Control Components Release 20.10.1	With this feature, you can automatically suspend an unstable Cisco Catalyst SD-WAN Bidirectional Forwarding Detection (BFD) session based on flap-cycle parameters or on Service-Level Agreement (SLA) parameters. You can also monitor the suspended BFD sessions and manually reset suspended BFD sessions. With this feature, you can automatically suspend an unstable Cisco Catalyst SD-WAN Bidirectional Forwarding Detection (BFD) session based on flap-cycle parameters or on Service-Level Agreement (SLA) parameters.

- [Information About Cisco Catalyst SD-WAN BFD, on page 164](#)
- [Information About Automatically Suspending BFD Sessions, on page 164](#)

- [Restrictions for Automatically Suspending BFD Sessions, on page 166](#)
- [Configure Automatic Suspension of BFD Sessions Using a CLI Template, on page 167](#)
- [Verify Automatic Suspension of BFD Sessions, on page 168](#)

Information About Cisco Catalyst SD-WAN BFD

Within Cisco Catalyst SD-WAN, there are the following types of BFD:

- **Cisco Catalyst SD-WAN BFD**

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

- Is enabled by default and cannot be disabled
- Is typically enabled for the Cisco Catalyst SD-WAN Overlay Management Protocol (OMP)
- Besides link failures, Cisco Catalyst SD-WAN BFD also measures latency, loss, jitter, and other link statistics used by application-aware routing

For more information on Cisco Catalyst SD-WAN BFD for measuring latency, loss, and jitter used by application-aware routing, see [Application-Aware Routing](#).

- **BFD Support for Routing Protocols in Cisco Catalyst SD-WAN**

This type of BFD supports BGP, OSPF, and EIGRP routing protocols in Cisco Catalyst SD-WAN.

For more information on BFD for routing protocols, see [BFD for Routing Protocols in Cisco Catalyst SD-WAN](#).

Information About Automatically Suspending BFD Sessions

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

BFD sessions may experience flapping, meaning that the BFD session enters a down state and then returns to an up state. This can occur when one device that is part of the BFD session becomes unavailable and then returns to being available. When a BFD session flaps, applications running on that tunnel are disrupted. The unstable BFD session can be brought up, but due to the unstable connection, the BFD session can quickly become disrupted again. With this feature, you avoid the impact of application traffic getting steered unnecessarily from one overlay path to another path because of an unstable BFD session.

To avoid the cycle of BFD session flaps, Cisco Catalyst SD-WAN provides an automatic suspension mechanism for suspending BFD sessions based on the following parameters:

- **Flap cycle**

A flap cycle is defined only as the following:

- BFD session is in the up state
- BFD session is in the down state
- BFD session is coming back up

- **SLA threshold**

An SLA threshold is the threshold by which the BFD session is added to the suspended list. An SLA threshold is a threshold value for a traffic metric, such as loss, latency, or jitter. If one of these metrics indicates that traffic performance has degraded to a point defined by a threshold, the BFD session state changes to suspended. These thresholds reflect the level of traffic performance specified in the SLA.



Note An SLA threshold is an optional configuration. If you configure a SLA threshold, configure higher metrics for loss, latency, and jitter, so the SLA threshold does not conflict with the SLA parameters as defined in the SLA classes. For more information on SLA classes, see the [Cisco Catalyst SD-WAN Policies Configuration Guide](#).

Benefits of Automatically Suspending BFD Sessions

- Supports manual removal of the affected circuit or tunnel interface from the BFD suspended list.
- Provides monitoring of a suspended tunnel.

How Automatically Suspending BFD Sessions Works

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

Configure the following BFD session parameters using a Cisco SD-WAN Manager device CLI template or a CLI add-on template:

Table 49: BFD Session Flap Cycle and SLA Parameters

Field	Description
enable-lr	Enable last resort upon BFD suspension. For more information on enabling a last resort on a tunnel interface, see last-resort-circuit .
duration	Duration of time for which the BFD session remains in the suspended state.
flapping-window	Time frame or window to detect the BFD session flap.
flap-count	Number of BFD session flaps after which the BFD session is suspended. The recommended flap-count is 3.
thresholds	SLA threshold triggering a BFD session to be suspended.

BFD Session Suspension 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.

For a BFD session in the suspended state, the following occurs:

1. If a session reflows or exceeds the threshold parameters defined, the session is moved back to suspended state and the duration is reset again.
2. 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.
3. 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](#).

Regular SLA measurement and echo response or path maximum transmission unit (PMTU) control traffic only is sent across the suspended BFD session.



Note Data traffic is not sent across the overlay network when a BFD session is in the suspended state.



Note This feature does not manipulate the state of the BFD session.



Note 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.

Restrictions for Automatically Suspending BFD Sessions

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

- For a Cisco IOS XE Catalyst SD-WAN device with a single TLOC, automatic suspension of a BFD session may cause BFD sessions to be dropped.
- 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.
- Cisco SD-WAN Manager feature templates do not support configuration of automatic suspension of BFD sessions.

Support is provided only for configuring BFD automatic suspension using a device CLI or a CLI add-on template.

- If duplicated traffic is sent on a different BFD session, the duplicated traffic may get routed through a BFD suspended session.

Configure Automatic Suspension of BFD Sessions Using a CLI Template

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

For more information about using CLI templates, see [CLI Add-On Feature Templates](#) and [CLI Templates](#).



Note By default, CLI templates execute commands in global config mode.

1. Enable BFD automatic suspension with or without last resort.

```
auto-suspend
  enable-lr

auto-suspend
  no enable-lr
```



Note 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](#).

2. Configure the following flap parameters:

```
duration sec
flapping-window sec
flap-count flap-count
```



Note When using SLA-based BFD automatic suspension, **duration** should be more than the number of the **bfd multiplier** x the **bfd poll interval**. We recommend that you configure BFD automatic suspension duration to be more than 30 minutes.

3. (Optional) Configure SLA parameters.

```
thresholds
  color
  all
  jitter jitter-value
  latency latency-value
  loss loss-value
  !
```

Prior to enabling SLA thresholds, configure BFD session flapping parameters and duration.

Here is a complete configuration example for configuring BFD automatic suspension with last resort enabled.

```
auto-suspend
  enable-lr
  duration 3600
```

```

flapping-window 300
flap-count      1
thresholds
  color
  all
  latency 10
  loss    10
  jitter  10

```



Note 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](#).

Verify Automatic Suspension of BFD Sessions

Minimum supported releases: Cisco IOS XE Catalyst SD-WAN Release 17.10.1a and Cisco Catalyst SD-WAN Control Components Release 20.10.1

The following sample output from the **show sdwan bfd sessions suspend** command displays the total suspend count, indicating the number of times that the BFD session has been suspended:

```

Device# show sdwan bfd sessions suspend
SYSTEM IP      STATE  SOURCE TLOC  REMOTE TLOC  DST PUBLIC  DST PUBLIC  RE-SUSPEND  SUSPEND  TOTAL  SUSPEND
            COLOR  COLOR          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

```

The following columns are added for analyzing BFD session suspension metrics: **RE-SUSPEND COUNT**, **SUSPEND TIME LEFT**, **TOTAL COUNT**, and **SUSPEND DURATION**.

The following sample output from the **show sdwan bfd sessions alt** command displays if a suspended flag has been added to a BFD session and other BFD session metrics:

```

Device# show sdwan bfd sessions alt
*Sus = Suspend
*NA = Flag Not Set
SYSTEM IP      SITE ID  STATE  SOURCE TLOC  REMOTE TLOC  DST PUBLIC  DST PUBLIC  BFD-LD  FLAGS  UPTIME
            COLOR  COLOR          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

```

The following columns are added for BFD suspension: **BFD-LD** and **FLAGS**.

Local discriminator (LD) is a unique identifier for all BFD sessions. The value for LD must be a nonzero value. LD is an internal value that Cisco Technical Assistance Center (TAC) uses for troubleshooting BFD sessions.

A BFD session flag, **Sus**, is added for identifying BFD sessions that are suspended.

The following sample output displays the BFD sessions for which the **Sus** flag is added to the BFD session:

```

Device# show sdwan bfd history
SYSTEM IP      SITE ID  COLOR  STATE  DST PUBLIC  DST PUBLIC  ENCAP  TIME  RX  TX  DEL  FLAGS
            IP           PORT        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]

```

The following sample output displays a BFD session summary, including which BFD sessions are up, down, flapped, or that have been 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
```

The following fields are added for BFD session suspension: `sessions-flap`, `sessions-up-suspended`, and `sessions-down-suspended`.



CHAPTER 9

Cisco Catalyst SD-WAN Controller Route Filtering by TLOC Color



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.

Table 50: Feature History

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 Controllers can reduce the number of routes that they advertise to routers in the network, to exclude routes that are not relevant to a particular device. The filtering to reduce the number of routes is based on the colors of TLOCs on each device. For example, a route to a public TLOC is not relevant to a router that only has private TLOCs. Advertising fewer routes helps to avoid reaching the send path limit for routers in the network.

- [Information About Cisco SD-WAN Controller Route Filtering by TLOC Color, on page 172](#)
- [Supported Devices for Cisco SD-WAN Controller Route Filtering by TLOC Color, on page 174](#)
- [Prerequisites for Cisco SD-WAN Controller Route Filtering by TLOC Color, on page 174](#)
- [Restrictions for Cisco SD-WAN Controller Route Filtering by TLOC Color, on page 175](#)
- [Configure Cisco SD-WAN Controller Route Filtering by TLOC Color Using a CLI Template, on page 175](#)
- [Monitor Cisco SD-WAN Controller Route Filtering by TLOC Color, on page 177](#)

Information About Cisco SD-WAN Controller Route Filtering by TLOC Color

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

Using route filtering, Cisco SD-WAN Controllers can reduce the number of routes that they advertise to routers in the network, to exclude 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.

Benefits

Advertising fewer routes offers the following benefits:

- Avoids reaching the send path limit:

Cisco SD-WAN Controller route filtering by TLOC color helps to avoid reaching the send path limit for routers in the network. For example, the send path limit might be set to 32, but Cisco SD-WAN Controllers might have more than 32 routes for a particular prefix to advertise to a device. Filtering out irrelevant routes helps to avoid reaching the limit.

- Prioritizes relevant routes:

If the send path limit is set to a low value, X, and if there are many routes to advertise, a Cisco SD-WAN Controller might advertise X irrelevant routes to a device, reaching the send path limit before advertising any relevant routes. This could result in a routing failure. Advertising only relevant routes prevents this possible failure.

Default Behavior

Cisco SD-WAN Controller route filtering by TLOC color is disabled by default.

Logic

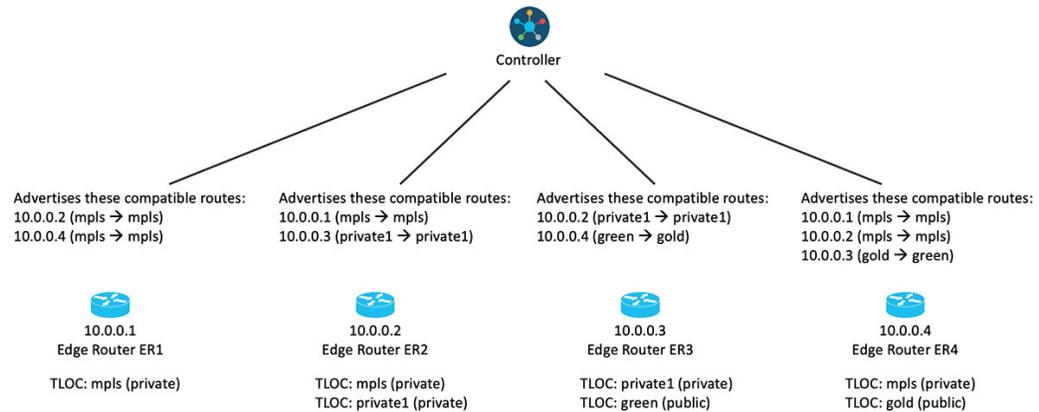
Cisco SD-WAN Controllers apply the following logic when determining whether routes are compatible:

- 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 has only TLOCs with private colors, Cisco SD-WAN Controllers do not advertise public routes to the device. Similarly, if a router has only TLOCs with public colors, Cisco SD-WAN Controllers do not advertise private routes to the device. The following illustration provides a more detailed example:

Figure 9: Cisco Catalyst 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.

Override

You can override the default logic if necessary and do one of the following:

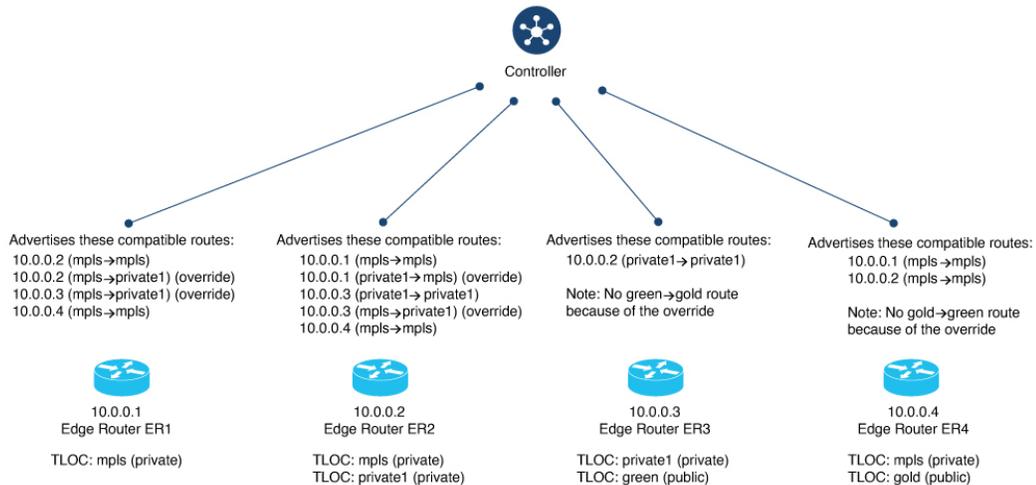
- Configure two TLOC colors to be compatible even if they are incompatible by default.
- Configure two TLOC colors to be incompatible even if they are compatible by default.

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, on page 176](#).

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 10: Cisco Catalyst SD-WAN Controller Route Filtering by TLOC Color, With the Feature Enabled and Overrides



Updating Cisco SD-WAN Controller of Changes

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, on page 176](#).

Supported Devices for Cisco SD-WAN Controller Route Filtering by TLOC Color

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

Cisco IOS XE Catalyst SD-WAN devices

Prerequisites for Cisco SD-WAN Controller Route Filtering by TLOC Color

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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. We do not support scenarios in which route filtering by TLOC color is enabled on some Cisco SD-WAN Controllers and disabled on others within the same network.

Configure Cisco SD-WAN Controller Route Filtering by TLOC Color Using a CLI Template

The following sections describe how to configure Cisco SD-WAN Controller route filtering by TLOC color.

Enable Route Filtering Using a CLI Template

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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.

The following configuration applies to a Cisco SD-WAN Controller.

1. Enter OMP mode.
`omp`
2. Enter filter-route configuration mode.
`filter-route`
3. Enable route filtering.
`outbound tloc-color`

Example

```
omp
  filter-route
    outbound tloc-color
!
```

Configure the Update Interval for Route Filtering by TLOC Color Using a CLI Template

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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.

The following configuration applies to a Cisco IOS XE Catalyst SD-WAN device.

1. Enter OMP configuration mode.

```
omp
```

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

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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

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.

Override Default TLOC Color Compatibility for Cisco SD-WAN Controller Route Filtering

The following configuration applies to a Cisco SD-WAN Controller.

1. Enter system mode.

```
system
```

2. Enter TLOC color compatibility mode.

```
tloc-color-compatibility
```

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
```

Example

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 vml
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

The following sections describe how to monitor Cisco SD-WAN Controller route filtering by TLOC color.

View TLOC Colors for a Device

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

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. When applying route filtering, the controllers use this TLOC color information to determine which routes are relevant to a device.

The following example shows the TLOC colors that the peer device 10.0.0.15 is advertising—in this case, lte and 3g.

```
vmanage#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

Minimum releases: Cisco IOS XE Catalyst SD-WAN Release 17.11.1a, Cisco Catalyst SD-WAN Control Components Release 20.11.1

To check the compatibility of TLOC colors, use the **request support omp tloc-color-compatible** command.

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-compatible 3g lte
Checking compatibility for colors:3g and lte
TLOC colors: 3g and lte are compatible
```

The following examples requests information about whether the 3g and mpls TLOC colors are compatible. They are incompatible:

```
vsmart# request support omp tloc-color-compatible 3g mpls
Checking compatibility for colors:3g and mpls
TLOC colors: 3g and mpls are incompatible
```



CHAPTER 10

Transport Gateway



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.

- [Transport Gateway, on page 180](#)
- [Information About Transport Gateways, on page 180](#)
- [Restrictions for Transport Gateways, on page 186](#)
- [Use Cases for Transport Gateways, on page 187](#)
- [Configure a Router as a Transport Gateway Using Cisco SD-WAN Manager, on page 189](#)
- [Configure a Router as a Transport Gateway Using a CLI Template, on page 190](#)
- [Configure the Transport Gateway Path Preference, on page 190](#)
- [Configure the Site Type for a Router Using Cisco SD-WAN Manager, on page 192](#)
- [Configure the Site Type for a Router Using a CLI Template, on page 192](#)
- [Verify the Site Type of a Router Using the CLI, on page 193](#)
- [Verify a Transport Gateway Configuration Using the CLI, on page 193](#)

Transport Gateway

Table 51: Feature History

Feature Name	Release Information	Description
Transport Gateway	Cisco Catalyst SD-WAN Manager Release 20.12.1 Cisco IOS XE Catalyst SD-WAN Release 17.12.1a	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: <ul style="list-style-type: none"> • 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

Information About Transport Gateways

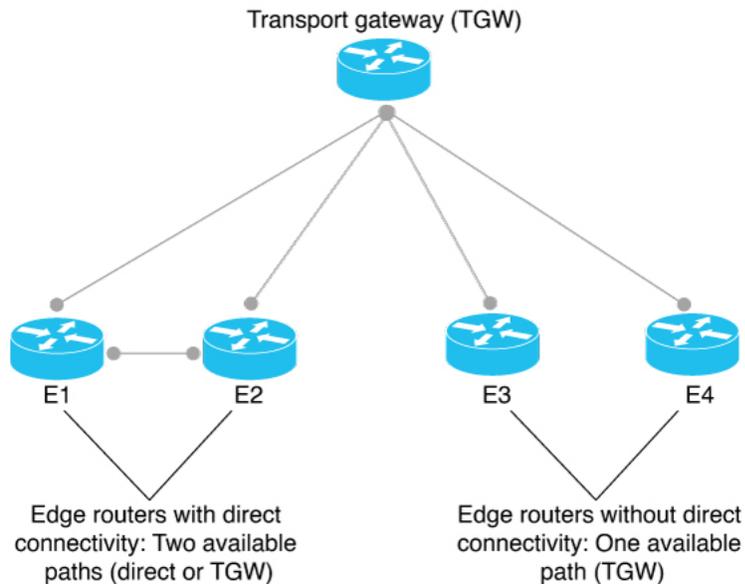
A transport gateway connects routers that may or may not have direct connectivity. A common use case for transport gateways is to provide connectivity between routers in disjoint networks, such as between a physical LAN and a cloud-based network.

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, thereby providing 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 11: Transport Gateway



Hub-and-Spoke Topology

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, on page 195](#).

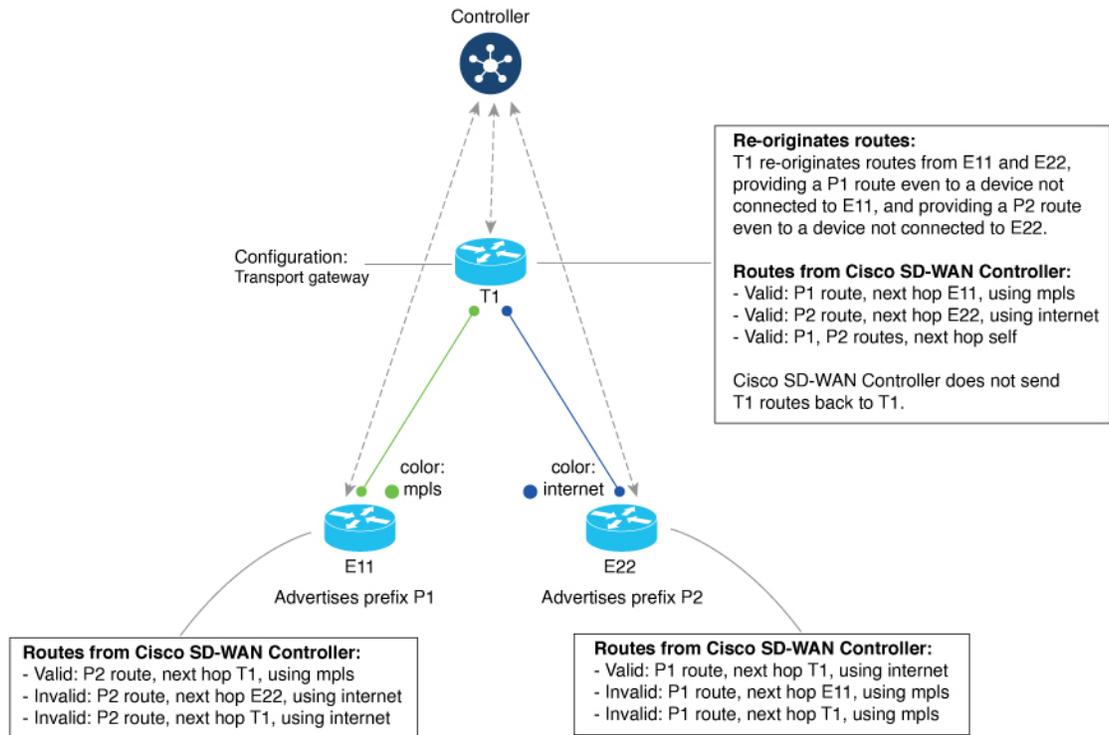
Re-originating Routes

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.

Figure 12: 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 [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](#), on page 192.

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.

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, on page 190](#).

Best Path Logic

Router Configuration		Resulting Best Path Behavior
Transport Gateway Path Behavior	Specify Site Type(s)	
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.

Multiple Transport Gateway Options

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.

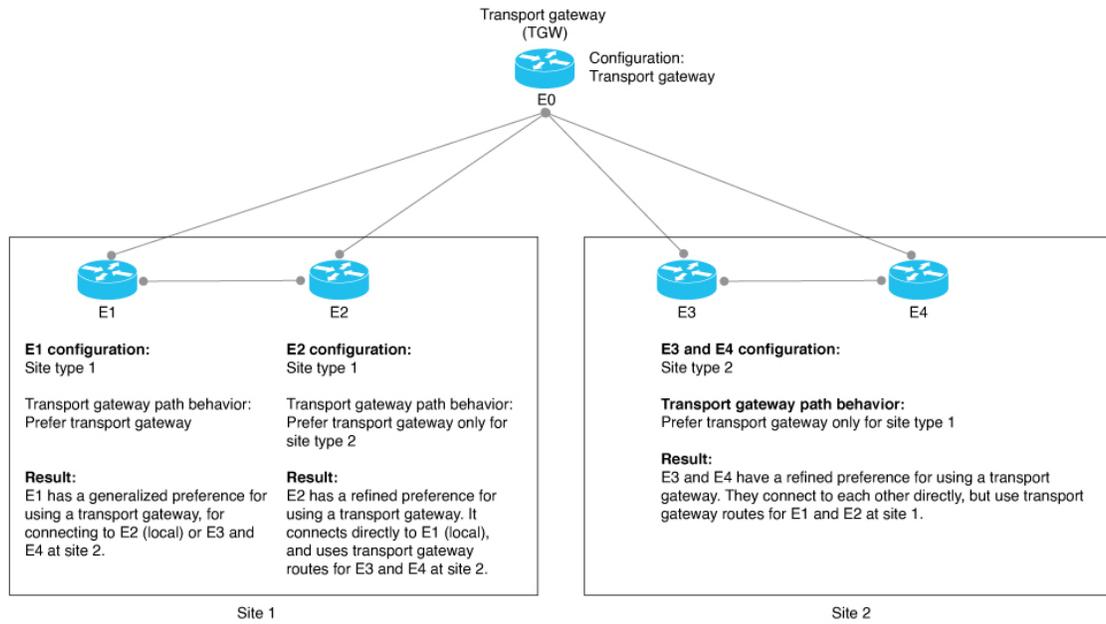
Configuration Overview

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, on page 189](#).
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, on page 190](#). 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, on page 192](#).
- 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.

Figure 13: Edge Routers and Transport Gateway Path Preference



The devices in the illustration are configured as follows:

Device	Configuration
E0	<ol style="list-style-type: none"> 1. Configure as a transport gateway. <ul style="list-style-type: none"> • By feature template: In a Cisco System template, use the Transport Gateway field. • By CLI add-on template: <pre>system transport-gateway enable</pre>

Device	Configuration
E1	<ol style="list-style-type: none"> 1. Configure the site type as type-1. <ul style="list-style-type: none"> • By feature template: In a Cisco System template, use the Site Type field. • By CLI add-on template: <pre style="margin-left: 20px;">system site-type type-1</pre> 2. For best path, configure a preference for transport gateway routes. <ul style="list-style-type: none"> • 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: <pre style="margin-left: 20px;">omp best-path transport-gateway prefer</pre>
E2	<ol style="list-style-type: none"> 1. Configure the site type as type-1. <ul style="list-style-type: none"> • By feature template: In a Cisco System template, use the Site Type field. • By CLI add-on template: <pre style="margin-left: 20px;">system site-type type-1</pre> 2. For best path, configure a preference for transport gateway routes for traffic to type-2 devices. <ul style="list-style-type: none"> • 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: <pre style="margin-left: 20px;">omp best-path transport-gateway prefer transport-gateway-settings type-2</pre>

Device	Configuration
E3 and E4	<ol style="list-style-type: none"> Configure the site type as type-2. <ul style="list-style-type: none"> By feature template: In a Cisco System template, use the Site Type field. By CLI add-on template: <pre>system site-type type-2</pre> For best path, configure a preference for transport gateway routes for traffic to type-1 devices. <ul style="list-style-type: none"> 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: <pre>omp best-path transport-gateway prefer transport-gateway-settings type-1</pre>

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 ensure that a route that is re-originated by one transport gateway is not advertised to another transport gateway. By preventing the advertising of a transport gateway route to another transport gateway, the Cisco SD-WAN Controller helps to prevent any potential 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 <i>Cisco SD-WAN Systems and Interfaces Configuration Guide, Cisco IOS XE Release 17.x</i> .

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.

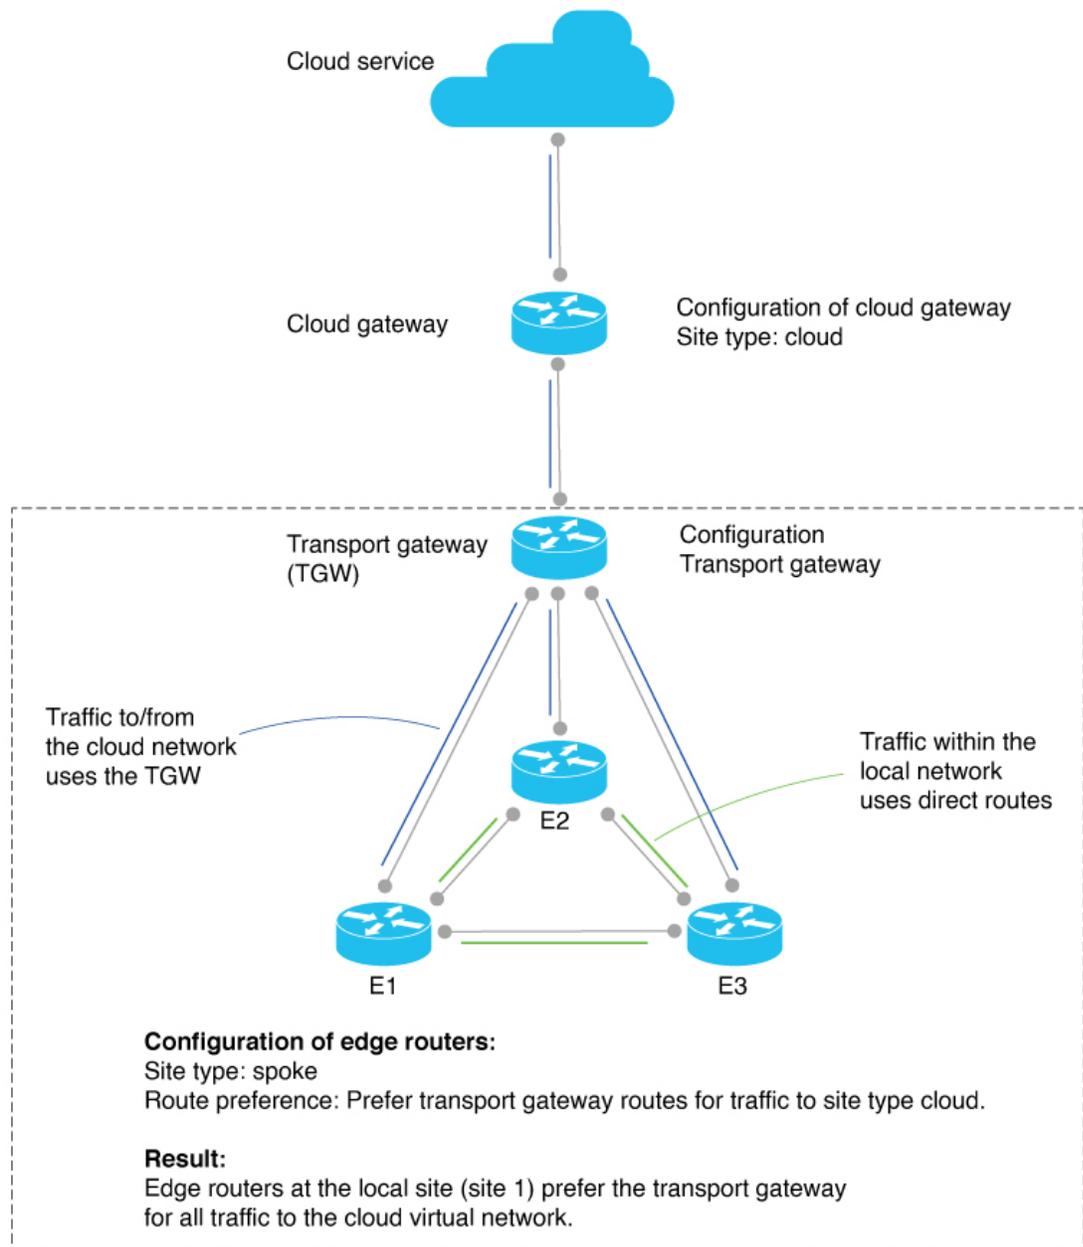
To create a transport gateway to bridge the local and cloud networks, network administrators configure the devices as follows:

Intent	Devices to Configure	Configuration
Configure the cloud gateway router with site type cloud.	Cloud gateway router	<ol style="list-style-type: none"> Configure the site type as cloud. <ul style="list-style-type: none"> By feature template: In a Cisco System template, use the Site Type field. By CLI template: <pre>system site-type cloud</pre>
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	<ol style="list-style-type: none"> Enable as a transport gateway. <ul style="list-style-type: none"> By feature template: In a Cisco System template, use the Transport Gateway field. By CLI template: <pre>system transport-gateway enable</pre>

Intent	Devices to Configure	Configuration
<p>Traffic within the local network uses direct routes, not transport gateway routes.</p> <p>Traffic from the local network to the cloud uses a transport gateway route.</p>	<p>Edge routers in the local network</p>	<ol style="list-style-type: none"> 1. Use a transport gateway route for all cloud-destined traffic. <ul style="list-style-type: none"> • By feature template: In an OMP template, use the Transport Gateway Path Behavior field. Choose the Prefer Transport Gateway Path option. • By CLI template: <pre>omp best-path transport-gateway prefer transport-gateway-settings cloud</pre> 2. Configure the site type as spoke. <ul style="list-style-type: none"> • By feature template: In a Cisco System template, use the Site Type field. • By CLI template: <pre>system site-type spoke</pre>

The following illustration shows the topology and configuration:

Figure 14: Transport Gateway Topology and Configuration



Configure a Router as a Transport Gateway Using Cisco SD-WAN Manager

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

2. Click **Feature Templates**.
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**.
4. In **Basic Configuration**, in the **Transport Gateway** field, choose **On**.
5. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure a Router as a Transport Gateway Using a CLI Template

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.

1. Enter system configuration mode.

```
system
```

2. Enable transport gateway functionality.

```
transport-gateway enable
```



Note To disable transport gateway functionality, use the **no** form of the command.

Example

```
system
transport-gateway enable
```

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 Cisco SD-WAN Manager

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

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.

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, on page 182](#).
6. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure the Transport Gateway Path Preference Using a CLI Template

Do the following on a device to configure it to use a transport gateway:

1. Enter sdwan configuration mode.
2. Enter system OMP configuration mode.
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.

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, on page 182](#).

```
omp best-path transport-gateway-settings site-types site-types
```

Option	Description
<i>site-types</i>	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 **omp best-path transport-gateway prefer** in the previous step.

Example

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 Site Type for a Router Using Cisco SD-WAN Manager

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.
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**.
4. In **Basic Configuration**, click **Site Type** and choose a type from the drop-down list .
5. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure the Site Type for a Router Using a CLI Template

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.

1. Enter system configuration mode.

```
system
```

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
```

Example

The following example configure a router with site types cloud and branch.

```
system
  site-type cloud branch
```

Verify the Site Type of a Router Using the CLI

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.

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

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 11

Hub-and-Spoke



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.

- [Hub-and-Spoke, on page 195](#)
- [Information About Hub-and-Spoke, on page 196](#)
- [Restrictions for Hub-and-Spoke, on page 206](#)
- [Use Cases for Hub-and-Spoke, on page 206](#)
- [Configure a Hub-and-Spoke Topology, on page 207](#)
- [Verify a Hub-and-Spoke Configuration, on page 208](#)

Hub-and-Spoke

Table 52: Feature History

Feature Name	Release Information	Description
Hub-and-Spoke Configuration	Cisco Catalyst SD-WAN Manager Release 20.12.1 Cisco IOS XE Catalyst SD-WAN Release 17.12.1a	Hub-and-spoke configuration simplifies the process of configuring a hub-and-spoke topology, making complex centralized control policy unnecessary. Instead, the configuration requires only a few simple configurations: a single command each on (a) the Cisco SD-WAN Controllers serving a network, (b) a router that serves as a hub, and (c) the routers that operate as spokes.

Information About Hub-and-Spoke

The hub-and-spoke topology is fundamental to networking, but configuring this topology can be complex, requiring expertise, and in a Cisco Catalyst SD-WAN environment, it can require lengthy centralized control policy configuration steps.

From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, a new configuration method speeds hub-and-spoke configuration without requiring complex control policy. Briefly put, this method involves configuring the Cisco SD-WAN Controllers 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	<p>Enable hub-and-spoke configuration in the network.</p> <p>See the following:</p> <ul style="list-style-type: none"> • Configure a Cisco Catalyst SD-WAN Controller to Enable Hub-and-Spoke Using Cisco SD-WAN Manager, on page 207 • Configure a Cisco SD-WAN Controller to Enable Hub-and-Spoke Using a CLI Template, on page 208 <p>The CLI template method uses the topology hub-and-spoke enable command.</p>
2. Configure a router as a transport gateway to function as a hub.	Router designated as hub	<p>Enable transport gateway functionality on the router.</p> <p>See Configure a Router as a Transport Gateway, for Hub-and-Spoke, on page 208.</p> <p>The CLI template method uses the transport-gateway enable command.</p>
3. Configure routers to function as spokes.	Routers designated as spokes	<p>Configure the device site type as spoke.</p> <p>See Configure the Site Type for a Router, for Hub-and-Spoke, on page 208.</p> <p>The CLI template method uses the site-type command.</p>

Result

This configuration results in the following:

- Cisco 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.

Example: Hub-and-Spoke Connectivity

The detailed example in this section shows how connectivity between devices in the network changes when you convert a full-mesh network to a hub-and-spoke topology. The following table shows information about the devices in the example, and the color coding used in the numerous illustrations that follow in this example section.

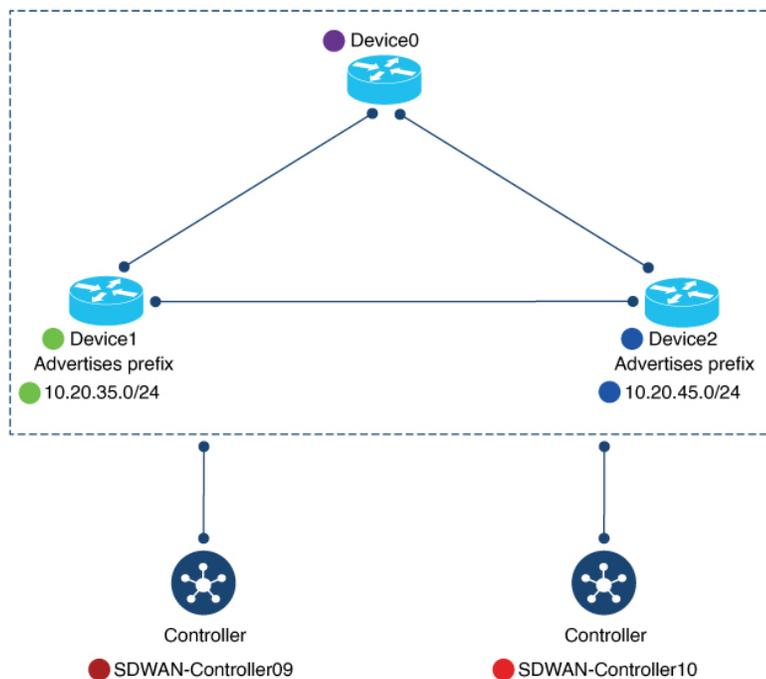
Table 53: Devices, IP Addresses, Roles, Interfaces, and Prefixes

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

Device	Intended Role	Interfaces	Prefixes
SDWAN-Controller10 172.16.255.20 Color in illustration: Red	Cisco SD-WAN Controller	Not applicable	Not applicable

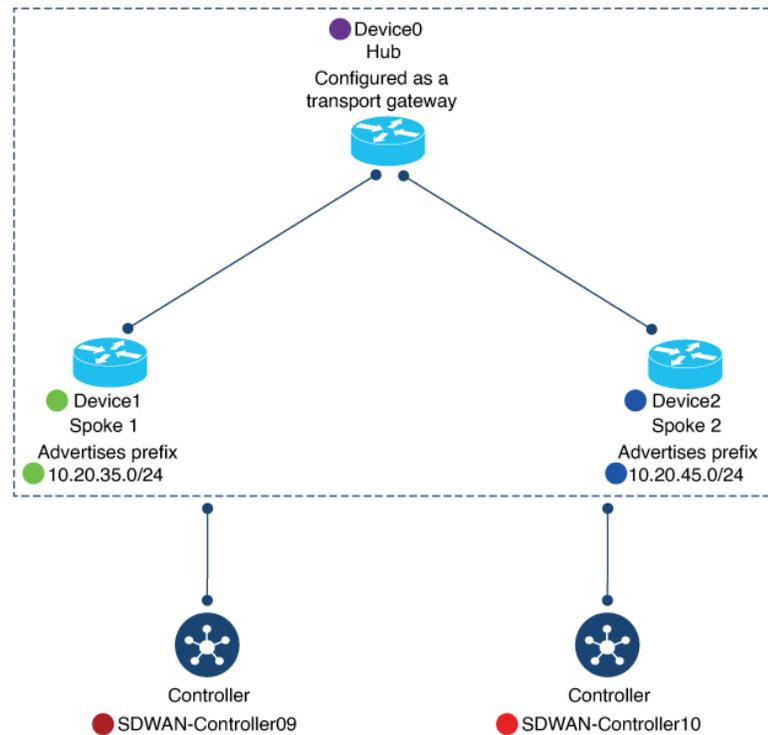
The following figure shows the initial state of the network, with full-mesh connectivity before configuring hub-and-spoke.

Figure 15: Before: Full-Mesh Connectivity



The following figure shows the network connectivity after configuring hub-and-spoke.

Figure 16: After: Hub-and-Spoke Connectivity



Device0 (Hub) Before and After

This section shows the connectivity for Device0 (hub) before and after configuring hub-and-spoke. It includes information about:

- BFD sessions
- OMP routes
- IP routes

BFD Sessions

Before configuring hub-and-spoke, on Device0 (future hub), the **show sdwan bfd sessions** command shows that it has BFD sessions with both Device1 (Spoke1) and Device2 (Spoke1).

After configuring hub-and-spoke, Device0 (hub) retains the same BFD sessions with both Device1 (Spoke1) and Device2 (Spoke2).

Figure 17: Hub: BFD Sessions Before and After

Before

```
Device0-future-hub#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC	IP
172.16.255.45	2500	up	3g	lte	10.0.20.15		10.0.6.45
172.16.255.35	1500	up	3g	lte	10.0.20.15		10.5.1.35
172.16.255.45	2500	up	lte	lte	10.1.15.15		10.0.6.45
172.16.255.35	1500	up	lte	lte	10.1.15.15		10.5.1.35

BFD sessions with Device1 (green)
and Device2 (blue)

After

```
Device0-Hub#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC	IP
172.16.255.45	2500	up	3g	lte	10.0.20.15		10.0.6.45
172.16.255.35	1500	up	3g	lte	10.0.20.15		10.5.1.35
172.16.255.45	2500	up	lte	lte	10.1.15.15		10.0.6.45
172.16.255.35	1500	up	lte	lte	10.1.15.15		10.5.1.35

BFD sessions with Device1 (green)
and Device2 (blue)

OMP Routes

Before configuring hub-and-spoke, on Device0 (future hub), 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 configuring hub-and-spoke on Device0 (hub), the Device1 (Spoke1) prefix and the Device2 (Spoke2) prefix are reachable through the hub itself (**FROM PEER** column shows **0.0.0.0**).

Figure 18: Hub: OMP Routes Before and After

Before

```
Device0-future-hub#show sdwan omp route vpn 1
```

TENANT	VPN	PREFIX	FROM PEER	PATH ID	LABEL	STATUS	ATTRIBUTE TYPE	TLOC IP	COLOR

Device1 prefix									
0	1	10.20.35.0/24	172.16.255.19	13	1003	C,I,R	installed	172.16.255.35	lte
			172.16.255.20	21	1003	C,R	installed	172.16.255.35	lte
0	1	10.20.45.0/24	172.16.255.19	46	1003	C,I,R	installed	172.16.255.45	lte
			172.16.255.20	17	1003	C,R	installed	172.16.255.45	lte

Annotations: Device1 prefix (green), Device2 prefix (blue), via Device1 (green), via Device2 (blue)

After

```
Device0-hub#show sdwan omp route vpn 1
```

TENANT	VPN	PREFIX	FROM PEER	PATH ID	LABEL	STATUS	ATTRIBUTE TYPE	TLOC IP	COLOR

Device1 prefix									
0	1	10.20.35.0/24	0.0.0.0	10737	1003	C,Red,R,	installed	172.16.255.15	lte
			0.0.0.0	41894	1003	TGW-R			
			0.0.0.0	10737	1003	C,Red,R,	installed	172.16.255.15	3g
			0.0.0.0	41895	1003	TGW-R			
			172.16.255.19	8	1003	C,I,R	installed	172.16.255.35	lte
			172.16.255.20	8	1003	C,R	installed	172.16.255.35	lte
0	1	10.20.45.0/24	0.0.0.0	10737	1003	C,Red,R,	installed	172.16.255.15	lte
			0.0.0.0	41894	1003	TGW-R			
			0.0.0.0	10737	1003	C,Red,R,	installed	172.16.255.15	3g
			0.0.0.0	41895	1003	TGW-R			
			172.16.255.19	9	1003	C,I,R	installed	172.16.255.45	lte
			172.16.255.20	9	1003	C,R	installed	172.16.255.45	lte

Annotations: Device1 prefix (green), Device2 prefix (blue), via Hub (green), via Device2 (blue)

IP Routes

Before configuring hub-and-spoke, on Device0 (future hub), 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 configuring hub-and-spoke, for Device0 (hub), this remains the same.

Figure 19: Hub: IP Routes Before and After

Before

```
Device0-hub#show ip route vrf 1
```

```
m 10.20.35.0/24 [251/0] via 172.16.255.35, 09:20:11, Sdwan-system-intf
m 10.20.45.0/24 [251/0] via 172.16.255.45, 09:20:11, Sdwan-system-intf
```

Annotations: Device1 prefix (green) via Device1 (green), Device2 prefix (blue) via Device 2 (blue)

After

```
Device0-hub#show ip route vrf 1
```

```
m 10.20.35.0/24 [251/0] via 172.16.255.35, 10:14:26, Sdwan-system-intf
m 10.20.45.0/24 [251/0] via 172.16.255.45, 10:14:26, Sdwan-system-intf
```

Annotations: Device1 prefix (green) via Device1 (green), Device2 prefix (blue) via Device 2 (blue)

Device1 (Spoke1) Before and After

This section shows the connectivity for Device1 (Spoke1) before and after configuring hub-and-spoke. It includes information about:

- BFD sessions
- OMP routes
- IP routes

BFD Sessions

Before configuring hub-and-spoke, on Device1 (future Spoke1), the **show sdwan bfd sessions** command shows BFD sessions with both Device0 (future hub) and Device2 (future Spoke2).

After configuring hub-and-spoke, Device1 (Spoke1) has only BFD sessions with the hub, not with other spokes—no BFD session with Spoke2 in this example.

Figure 20: Spoke1: BFD Sessions Before and After

Before

```
Device1-future-spoke1#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC IP
172.16.255.45	2500	up	lte	lte	10.5.1.35	10.0.6.45
172.16.255.15	500	up	lte	3g	10.5.1.35	10.0.20.15
172.16.255.15	500	up	lte	lte	10.5.1.35	10.1.15.15

BFD sessions with Device2 (blue)
and Hub (purple)

After

```
Device1-spoke1#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC IP
172.16.255.15	500	up	lte	3g	10.5.1.35	10.0.20.15
172.16.255.15	500	up	lte	lte	10.5.1.35	10.1.15.15

BFD sessions only with Hub (purple)

OMP Routes

Before configuring hub-and-spoke, on Device1 (future Spoke1), 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 configuring hub-and-spoke, Device1 (Spoke1) can reach the Device2 (Spoke2) prefix only through the hub.

Figure 21: Spoke1: OMP Routes Before and After

Before

```
Device1-future-spoke1#show sdwan omp route vpn 1
```

TENANT	VPN	PREFIX	FROM PEER	PATH ID	LABEL	STATUS	ATTRIBUTE TYPE	TLOC IP	COLOR
0	1	10.20.45.0/24	172.16.255.19	43	1003	C,I,R	installed	172.16.255.45	lte
			172.16.255.20	21	1003	C,R	installed	172.16.255.45	lte

Device2 prefix

via Device2

After

```
Device1-spoke1#show sdwan omp route vpn 1
```

TENANT	VPN	PREFIX	FROM PEER	PATH ID	LABEL	STATUS	ATTRIBUTE TYPE	TLOC IP	COLOR
0	1	10.20.45.0/24	172.16.255.19	10	1003	C,I,R	installed	172.16.255.15	lte
			172.16.255.19	11	1003	C,I,R	installed	172.16.255.15	3g
			172.16.255.20	10	1003	C,R	installed	172.16.255.15	lte
			172.16.255.20	11	1003	C,R	installed	172.16.255.15	3g

Device2 prefix

via Hub

IP Routes

Before configuring hub-and-spoke, on Device1 (future Spoke1), the **show ip route vrf 1** command shows that it can reach the Device2 prefix directly through Device2.

After configuring hub-and-spoke, Device1 (Spoke1) can reach the Device2 (Spoke2) prefix only through the hub.

Figure 22: Spoke1: IP Routes Before and After

Before

```
Device1-future-spoke1#show ip route vrf 1
```

m	10.20.45.0/24	[251/0]	via 172.16.255.45,	06:03:36,	Sdwan-system-intf
---	---------------	---------	--------------------	-----------	-------------------

Device2 prefix (blue) via Device2 (blue)

After

```
Device1-spoke1#show ip route vrf 1
```

m	10.20.45.0/24	[251/0]	via 172.16.255.15,	10:14:58,	Sdwan-system-intf
---	---------------	---------	--------------------	-----------	-------------------

Device2 prefix (blue) via Hub (purple)

Device2 (Spoke2) Before and After

This section shows the connectivity for Device2 (Spoke2) before and after configuring hub-and-spoke. It includes information about:

- BFD sessions
- OMP routes

- IP routes

The changes for Device2 before and after configuring hub-and-spoke mostly mirror the changes for Device1.

BFD Sessions

Before configuring hub-and-spoke, on Device2 (future Spoke2), the **show sdwan bfd sessions** command shows BFD sessions with both Device0 (future hub) and Device1 (future Spoke1).

After configuring hub-and-spoke, Device2 (Spoke2) has only BFD sessions with the hub, not with other spokes—no BFD session with Spoke1 in this example.

Figure 23: Spoke2: BFD Sessions Before and After

Before

```
Device2-future-spoke2#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC	IP
172.16.255.35	1500	up	lte	lte	10.0.6.45		10.5.1.35
172.16.255.15	500	up	lte	3g	10.0.6.45		10.0.20.15
172.16.255.15	500	up	lte	lte	10.0.6.45		10.1.15.15

BFD sessions with Device1 (green) and Hub (purple)

After

```
Device2-spoke2#show sdwan bfd sessions
```

SYSTEM IP	SITE ID	STATE	SOURCE TLOC COLOR	REMOTE TLOC COLOR	SOURCE IP	DST PUBLIC	IP
172.16.255.15	500	up	lte	3g	10.0.6.45		10.0.20.15
172.16.255.15	500	up	lte	lte	10.0.6.45		10.1.15.15

BFD sessions only with Hub (purple)

OMP Routes

Before configuring hub-and-spoke, on Device2 (future Spoke2), the **show sdwan omp route vpn 1** command shows that it can reach the Device1 (Spoke1) prefix directly through Device1. This is evident because the **TLOC IP** column shows the system IP of Device1.

After configuring hub-and-spoke, Device2 (Spoke2) can reach the Device1 (Spoke1) prefix only through the hub.

- Reducing 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 you to apply network services, such as firewall policy, to all inter-spoke traffic.

Configuring a hub-and-spoke topology using the process described here simplifies the configuration process, avoiding complex centralized control policy.

Restrictions for Hub-and-Spoke

Restriction	Description
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.

Use Cases for Hub-and-Spoke

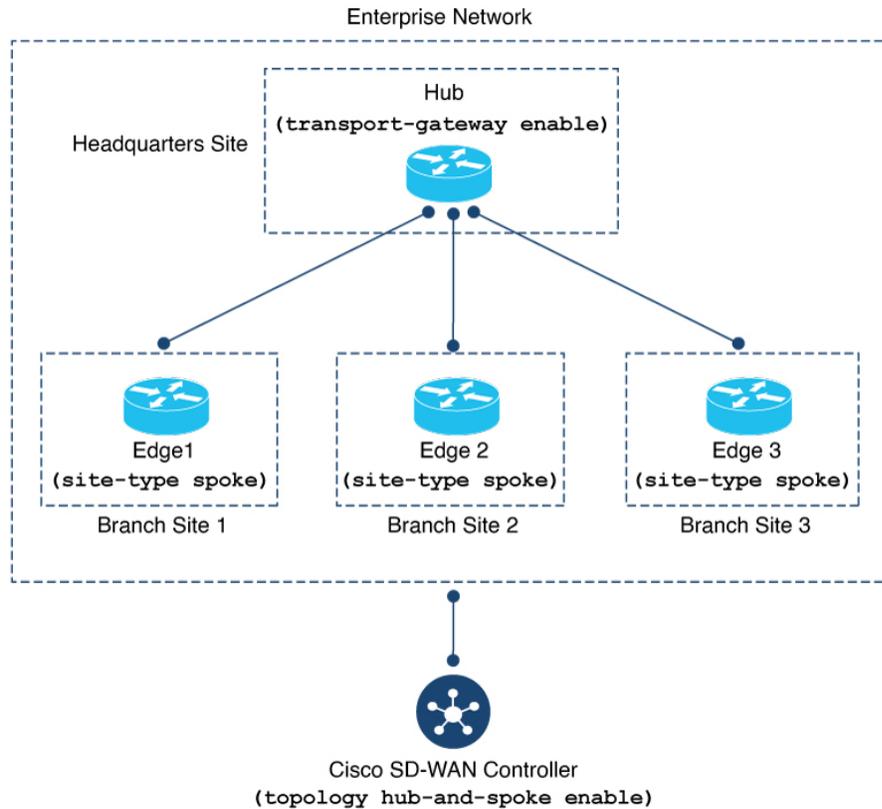
In this use case, an organization's network includes the following elements:

- A single device at the organization's headquarters site that runs numerous network services, such as an enterprise firewall. Network administrators have chosen to designate this as a hub device.
- Three branch sites, each with an edge router serving the site.

Network administrators have chosen to configure a hub-and-spoke topology to route all traffic flowing between branch sites through the hub at the headquarters site. This enables them to apply the centralized network services to all traffic between branch sites.

They configure a hub-and-spoke topology as shown in the following illustration:

Figure 26: Hub-and-Spoke Topology



Configure a Hub-and-Spoke Topology

The sections that follow describe procedures for configuring a hub-and-spoke topology using transport gateways.

Configure a Cisco Catalyst SD-WAN Controller to Enable Hub-and-Spoke Using Cisco SD-WAN Manager

1. From the Cisco SD-WAN Manager menu, choose **Configuration** > **Templates**.
2. Click **Feature Templates**.
3. Do one of the following:
 - To create a new System template for Cisco SD-WAN Controllers, click **Add Template**, choose **vSmart**, and click **System**.
 - To edit an existing Cisco SD-WAN Controller System template, locate a template of type **vSmart System** in the table of existing feature templates, click **...** adjacent to the template, and choose **Edit**.
4. In the **Topology** field, choose **Hub and Spoke**.

5. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure a Cisco SD-WAN Controller to Enable Hub-and-Spoke Using a CLI Template

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.

1. Enter system configuration mode.

```
system
```

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

```
system
 topology hub-and-spoke enable
```

Configure a Router as a Transport Gateway, for Hub-and-Spoke

Hub-and-spoke configuration makes use of 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 189](#)
- [Configure a Router as a Transport Gateway Using a CLI Template, on page 190](#)

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 192](#)
- [Configure the Site Type for a Router Using a CLI Template, on page 192](#)

Verify a Hub-and-Spoke Configuration

Hub-and-spoke configuration makes use of transport gateways and the site type parameter, which are described in the [Transport Gateway](#).

- For information about verifying a transport gateway configuration, see [Verify a Transport Gateway Configuration Using the CLI, on page 193](#).

- For information about verifying the site type, see [Verify the Site Type of a Router Using the CLI](#), on page 193.
- 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: [Example: Hub-and-Spoke Connectivity](#), on page 197

Verify that a Cisco Catalyst SD-WAN Controller Has Enabled Hub-and-Spoke Configuration

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
```




CHAPTER 12

Symmetric Routing



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

- [Symmetric Routing, on page 212](#)
- [Information About Symmetric Routing, on page 212](#)
- [Configuration Overview, on page 220](#)
- [Supported Scenarios, on page 229](#)
- [Prerequisites for Symmetric Routing, on page 236](#)
- [Restrictions for Symmetric Routing, on page 237](#)
- [Configure Symmetric Routing, on page 237](#)
- [Verify Symmetric Routing, on page 241](#)
- [Monitor RIB Metric Translation, on page 244](#)

Symmetric Routing

Table 54: Feature History

Feature Name	Release Information	Description
Symmetric Routing	Cisco Catalyst SD-WAN Control Components Release 20.12.1	You can use affinity groups, affinity group preference, and translation of RIB metrics to ensure symmetric routing of traffic flows across devices in a network. Symmetric routing accommodates various network topologies, including Multi-Region Fabric.
	Cisco IOS XE Catalyst SD-WAN Release 17.12.1a	To support symmetric routing beyond the overlay network, transport gateways can translate RIB metrics to control plane protocols such as BGP and OSPF. This extends the path preference configuration to routers outside of the overlay network, such as routers in a data center LAN.

Information About Symmetric Routing

Symmetric routing refers to traffic flows between two endpoints, that use the same route for traffic in both directions. Some networking functionality requires symmetric routing in order to operate correctly, such as Cisco Network Based Application Recognition (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 such that the preferred route between two endpoints will be consistent for traffic in both directions. This ensures symmetric routing for traffic flows between those endpoints. In some scenarios, you can even ensure that symmetric routing for traffic flows that extend to a device outside of the Cisco Catalyst SD-WAN overlay network.

Assumption That Routers Remain Operational

All of this applies to a situation when a router does not become inoperable during a traffic flow. If a router that is part of the path of a traffic flow becomes inoperable, traffic must change routes, which can temporarily cause asymmetric routing of the traffic flow.

Benefits of Symmetric Routing Configuration

Before Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, configuring symmetric routing has involved the following:

- In the overlay network: Complex and error-prone control policies to set up hop-by-hop routing for traffic in both directions to ensure symmetric routing.
- In service-side routing: Complex route-maps to set up path symmetry for traffic in both directions.

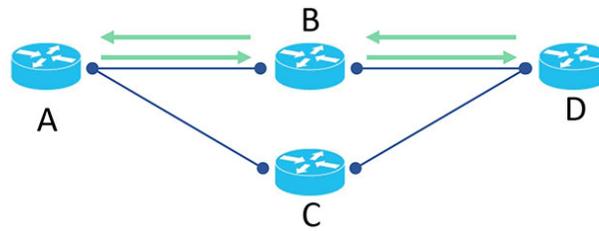
From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, you can use affinity groups and preferences, and OMP metric redistribution, to achieve symmetric routing. The sections below describe the details and supported scenarios.

Mechanisms for Ensuring Symmetric Routing

In a network managed by Cisco Catalyst SD-WAN, the Overlay Management Protocol (OMP) maintains control plane tasks. This includes applying a best-path algorithm to determine each next hop for traffic between two endpoints. OMP considers numerous parameters when comparing the various available next hops. For information, see [Unicast Overlay Routing](#) in the *Cisco Catalyst SD-WAN Routing Configuration Guide, Cisco IOS XE Release 17.x*.

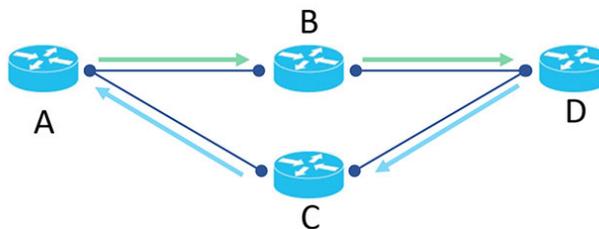
For return traffic to choose the same path, we need to ensure that for each hop, the best path calculation will favor the same route in both directions. For example, the following figure depicts a flow from A to D. The first hop is from A to B, followed by B to D. We need to ensure that for a given traffic flow, in the reverse direction, the first hop is from D to B, followed by B to A.

Figure 27: Symmetric Flow



If the reverse traffic (from D to A) uses D to C as its first hop, the traffic flow would be asymmetric, as shown in the following figure:

Figure 28: Asymmetric Flow



Mechanisms

For topologies using transport gateways as routing hubs, or Multi-Region Fabric networks, Cisco Catalyst SD-WAN uses the following mechanisms to ensure that devices choose the same path between two endpoints, for traffic in both directions:

Mechanism	Description
Affinity group	<p>Affinity groups enable you to specify the order of preference for choosing among multiple next hops for a traffic flow. For information about router affinity, see Router Affinity in the <i>Cisco Catalyst SD-WAN Multi-Region Fabric (also Hierarchical SD-WAN) Configuration Guide</i>.</p> <p>Related configuration procedures:</p> <ul style="list-style-type: none"> • Configure an Affinity Group or Affinity Group Preference on a Device, Using Cisco SD-WAN Manager • Configure an Affinity Group on a Router Using the CLI <p>Uses the affinity-group <i>group-id</i> command.</p>
Derived affinity group	<p>From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, when border routers in a Multi-Region Fabric topology, or transport gateways serving Multi-Region Fabric subregions, re-originate routes, they assign a derived affinity group to the route. This is one part of the overall mechanism of ensuring that return traffic uses the same gateway or border router as the forward traffic.</p> <p>Border routers use the derived affinity attribute instead of affinity groups to determine a preferred route within the core region. A lower derived affinity value has a higher preference. For example, if border router BR1 has two border routers, BR2 and BR3 available as a next hop, BR1 chooses the one with a lower derived affinity group value as computed by the border router.</p> <p>Note As described in Prerequisites for Symmetric Routing, on page 236, to ensure symmetric routing, border routers and transport gateways require (a) an affinity group number or (b) affinity groups per VRF for all VRFs that the devices handle.</p>
Affinity group for a specific VRF range	<p>You can configure a router to have different affinity groups for different VRF ranges. Per-VRF affinity groups provide granular control of route preferences according to VRF.</p> <p>Related configuration procedures:</p> <ul style="list-style-type: none"> • Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager, on page 238 • Configure Router Affinity Groups for Specific VRFs Using a CLI Template, on page 238 <p>Uses the affinity-per-vrf <i>affinity-group vrf-range vrf-range</i> command.</p>

Mechanism	Description
Affinity preference order	<p>Together with affinity groups, this provides control of route preferences for the next hop. When you configure the affinity preference order manually, a device prefers routes with an affinity group that occurs earlier in the preference order.</p> <p>From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, you can configure automatic affinity preference order. When you use this, a device prefers routes with a lower affinity group number. In this case affinity group numbers are not treated as arbitrary tags, but instead signify route priority, where a lower affinity group means higher priority.</p> <p>Note From Cisco IOS XE Catalyst SD-WAN Release 17.12.1a, devices tag vRoutes (routes in the Cisco Catalyst SD-WAN overlay network) with the affinity preference order attribute, as follows:</p> <ul style="list-style-type: none"> • If you manually configure an affinity preference order for a device, the device tags vRoutes with the preference order you have configured, with a maximum of eight affinity groups (the first eight in the list). • If you configure auto affinity preference order, the device tags vRoutes with a value used internally by Cisco Catalyst SD-WAN that indicates the automatic preference order. • If you manually configure an affinity preference order for a device, and also configure auto affinity preference order, the device tags vRoutes with a value used internally by Cisco Catalyst SD-WAN that indicates the automatic preference order, as in the previous option. (For information about the use case for configuring the affinity preference order manually and using auto simultaneously, see Configure a Router to Use Automatic Affinity Group Preference Using Cisco SD-WAN Manager, on page 237.)
Affinity preference order (continued)	<p>Related configuration procedures:</p> <ul style="list-style-type: none"> • (Manual configuration) Configure an Affinity Group or Affinity Group Preference on a Device, Using Cisco SD-WAN Manager • (Manual configuration) Configure Affinity Group Preference on a Router Using the CLI Uses the affinity-group preference list command. • (Automatic configuration) Configure a Router to Use Automatic Affinity Group Preference Using Cisco SD-WAN Manager, on page 237 • (Automatic configuration) Configure a Router to Use Automatic Affinity Group Preference Using a CLI Template, on page 239 Uses the affinity-group preference-auto command.

Mechanism	Description
Redistribution of OMP metrics into service-side routing protocols	<p>In a network topology that includes routers managed by Cisco Catalyst SD-WAN and routers not managed by Cisco Catalyst SD-WAN, you can propagate routing information base (RIB) metrics from OMP to the service-side portion of the network, which can use the border gateway protocol (BGP) or the open shortest path first (OSPF) protocol. This ensures that service-side routers can prioritize the same routes for return traffic to enable routing symmetry even across different control planes. For information, see Translating OMP Metrics for Devices Outside of the Overlay Network, on page 216.</p> <p>Related configuration procedure:</p> <ul style="list-style-type: none"> • Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template, on page 239 <p>Uses the redistribute omp translate-rib-metric command.</p>

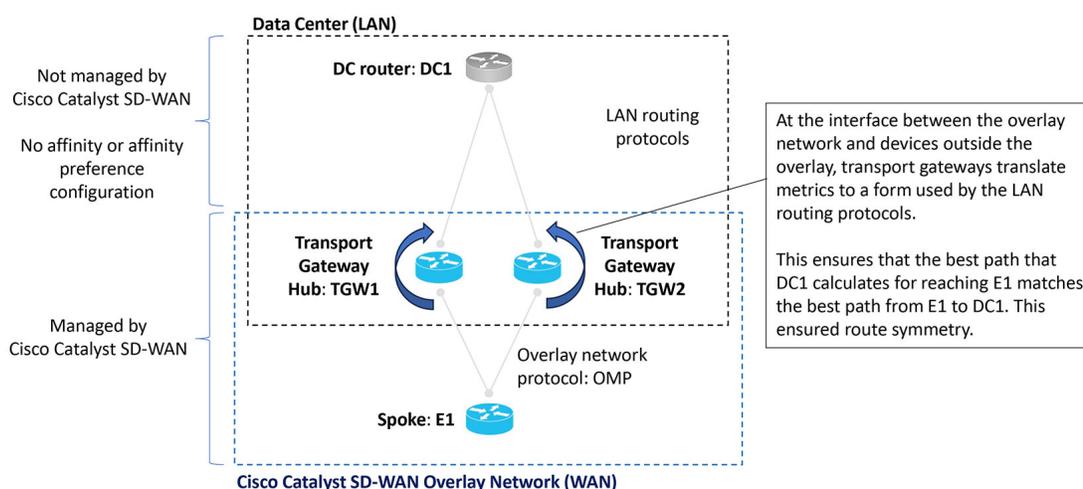
Translating OMP Metrics for Devices Outside of the Overlay Network

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 29: Translating OMP Metrics



Related Topics

[Translating OMP Metrics to BGP Attributes, on page 217](#)

[Translating OMP Metrics to an OSPF Metric, on page 219](#)

[Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template, on page 239](#)

[Monitor RIB Metric Translation, on page 244](#)

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 (Note about terminology: Among the OMP metrics, there is one specifically called OMP.)
- OMP AS-PATH

...to derive three BGP attributes:

- BGP MED
- BGP LOCAL_PREF
- BGP AS_PATH

For information about viewing the OMP metrics for a route and the resulting BGP attributes, see [Monitor RIB Metric Translation, on page 244](#).

The translation from OMP to BGP is as follows:

Table 55: Translation 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	<p>Two possibilities:</p> <ul style="list-style-type: none"> • If the propagate-aspath command is used: <ol style="list-style-type: none"> (a) 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. (b) 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: <p>A list of its own local AS value configured for the router, repeated (OMP route metric) times, prepending the value a maximum of 13 times.</p>



Note 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). Omitting this causes a router to treat the AS-Path metric as empty.

The router includes these BGP attributes with routes that it re-originates to a device in a LAN that is outside of the overlay network and is using BGP.

BGP Attributes without RIB Metric Translation

The following table shows combinations of OMP metrics, and the BGP attributes that the router derives when RIB metric translation is not enabled.

Table 56: 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 with RIB Metric Translation

The following table shows combinations of OMP metrics, and the BGP attributes that the router derives when RIB metric translation is enabled.

Table 57: Translation from OMP to BGP with RIB Metric Translation Enabled

	OMP Metrics: Example Combinations		Translation to BGP Attributes: propagate-aspath enabled and translate-rib-metric enabled		
Example	OMP route metric	OMP AS-PATH	BGP MED	BGP LOCAL_PREF	BGP AS_PATH
1	0	100 101	0	255	100 101 (Nothing prepended because the OMP route metric is 0)

	OMP Metrics: Example Combinations		Translation to BGP Attributes: propagate-aspath enabled and translate-rib-metric enabled		
Example	OMP route metric	OMP AS-PATH	BGP MED	BGP LOCAL_PREF	BGP AS_PATH
2	1	100 101	1	254	100 100 101 (1 repetition of the initial value prepended to the list)
3	2	100 101	2	253	100 100 100 101 (2 repetitions of the initial value prepended to the list)
4	10	(empty) In this example, the local AS value is 200 .	10	245	200 200 200 200 200 200 200 200 200 200 (10 repetitions of the router AS value)
5	14	100 101	14	241	100 100 100 100 100 100 100 100 100 100 100 100 100 100 101 (13 repetitions, the maximum, of the initial value prepended to the list)

Translating OMP Metrics to an OSPF Metric

If you do not configure a router to translate RIB metrics, the router uses a default OSPF metric when redistributing routes to a device outside of the Cisco Catalyst SD-WAN overlay network. The default OSPF metric is 16777214 (hexadecimal FFFFFFFE).

When you enable a router to translate RIB metrics, the router assigns the OMP route metric value as the OSPF metric. For example, if the OMP route metric is 10, the OSPF metric will also be 10.

For information about viewing the OMP metrics for a route and the resulting BGP metrics, see [Monitor RIB Metric Translation, on page 244](#).

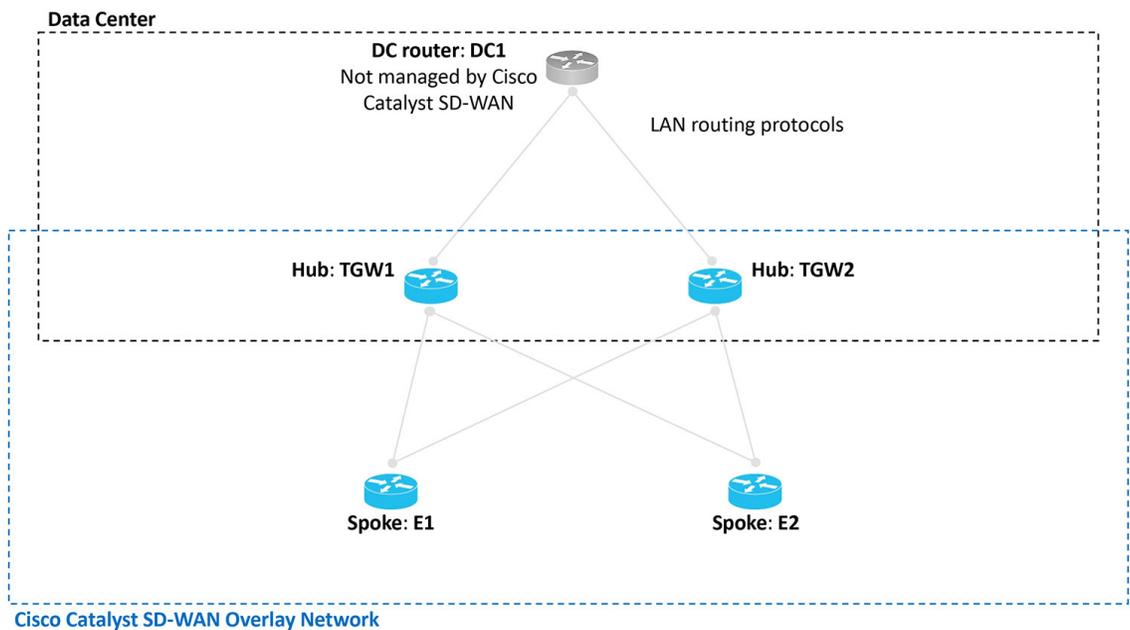
Configuration Overview

An overview of the configuration workflow is helpful for understanding the scenarios in which Cisco Catalyst SD-WAN supports symmetric routing. The following figures show 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).

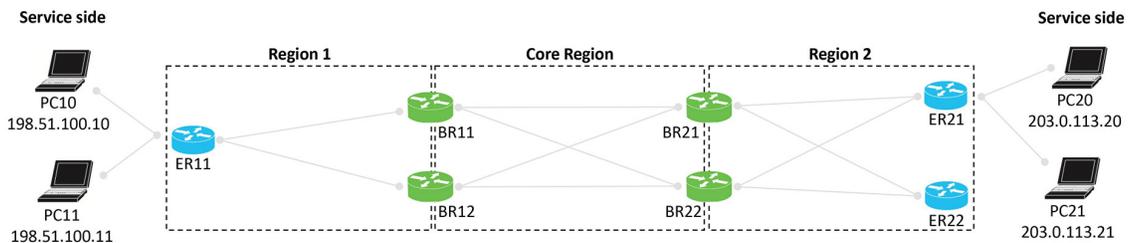
Figure 30: Transport Gateway Scenario with a Data Center LAN



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 31: Multi-Region Fabric Scenario



Configuration Overview

The following steps 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	<p>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.</p> <p>With automatic affinity preference order, a spoke device or edge router prefers paths tagged with a lower affinity group number.</p> <p>For configuration instructions, see Configure a Router Affinity Group or Affinity Group Preference, on page 238.</p>
2. Configure affinity groups	Transport gateways Border routers in a Multi-Region Fabric scenario	<p>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.</p> <p>For example, if a device has a VRF range of 1 to 10, you can configure a device as follows:</p> <ul style="list-style-type: none"> • System-level affinity group 10 • Affinity groups per VRF: Affinity group 20 for VRF6 through VRF 10 <p>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.</p> <p>For configuration instructions, see Configure a Router Affinity Group or Affinity Group Preference, on page 238.</p>

Configuration Step	Devices	Description
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 , on page 216. For configuration instructions, see Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template , on page 239.

The following illustrations 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 32: Transport Gateway Scenario with a Data Center LAN, Showing a Configuration for Symmetric Routing

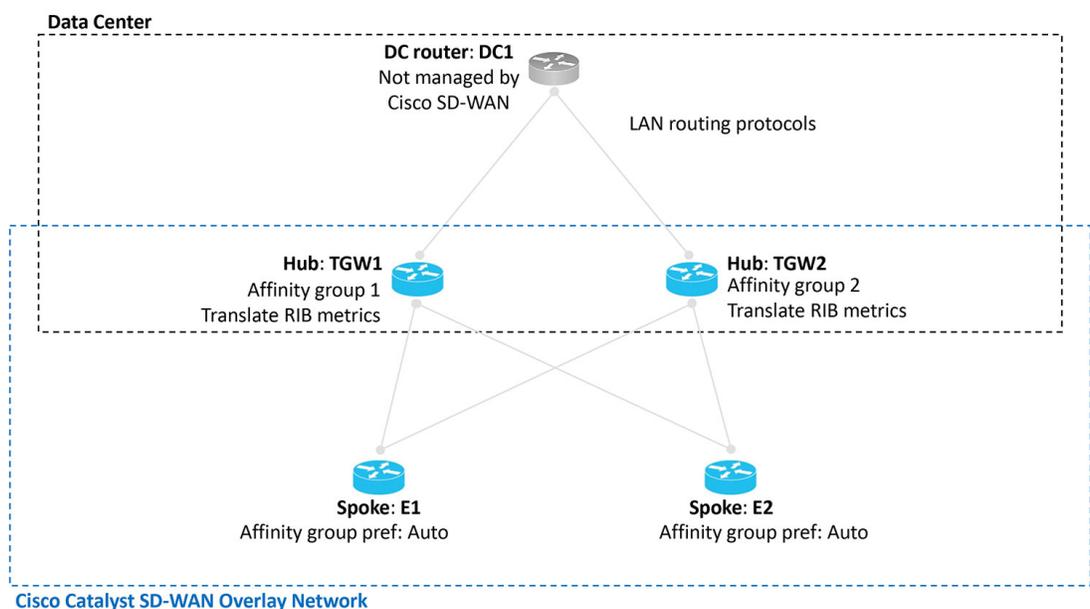
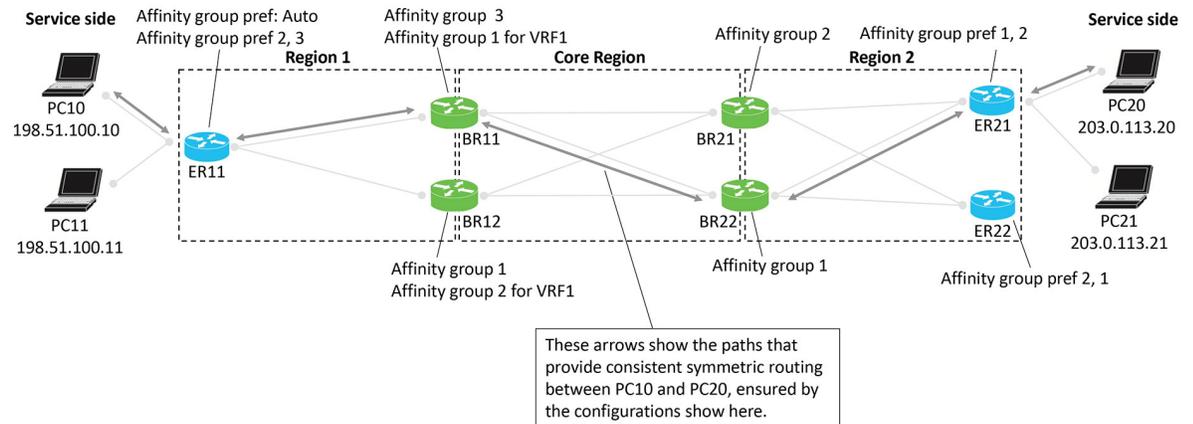


Figure 33: Multi-Region Fabric Scenario, Showing a Configuration for Symmetric Routing

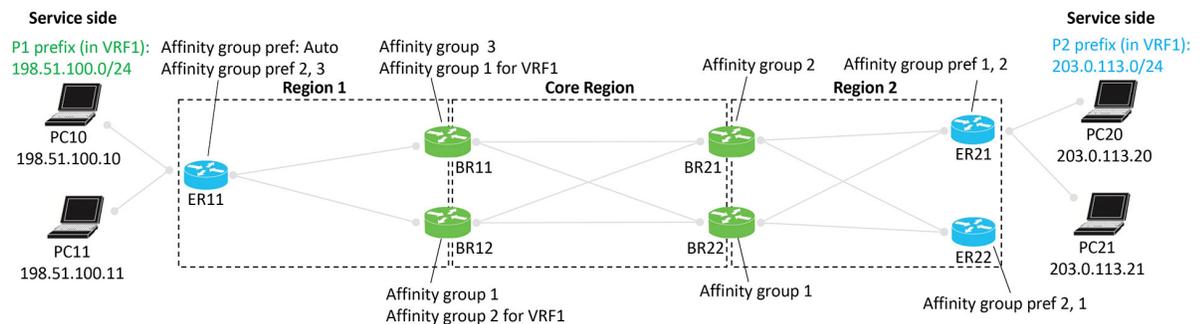


Example of Configuration for Symmetric Routing and the Mechanism

The following comprehensive example shows an approach to configuring border routers and edge routers in a Multi-Region Fabric environment to provide symmetric routing between the PC devices served by edge router ER11 in Region 1, and the PC devices served by ER21 in Region 2. Specifically, the example focuses on traffic between PC10 and PC20.

The step-by-step illustrations show how route re-origination and path preference result in a preference for the same path through multiple hops for traffic in both directions.

Figure 34: Multi-Region Fabric Scenario, Configuration for Symmetric Routing



Advertising P1 Routes

Edge router ER11 advertises P1 routes. The process of re-originating these routes to border routers, and eventually to ER21 and ER22 occurs from left to right in the illustration. In the process, border routers assign affinity groups and derived affinity groups when re-originating routes.

Routers in the network choose preferred routes as follows:

- Outside the core region: According to affinity group preference

- In the core region: According to the lowest derived affinity group (dag) value

Figure 35: Edge Router ER11 Advertises P1 Routes

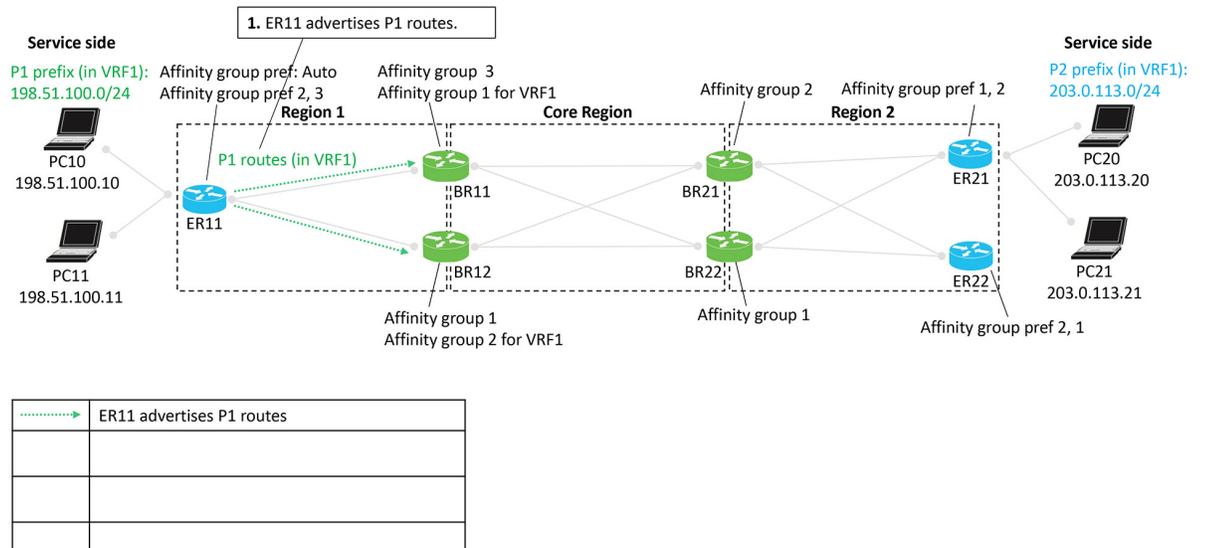


Figure 36: Border Routers BR11 and BR12 Re-Originate the P1 Routes

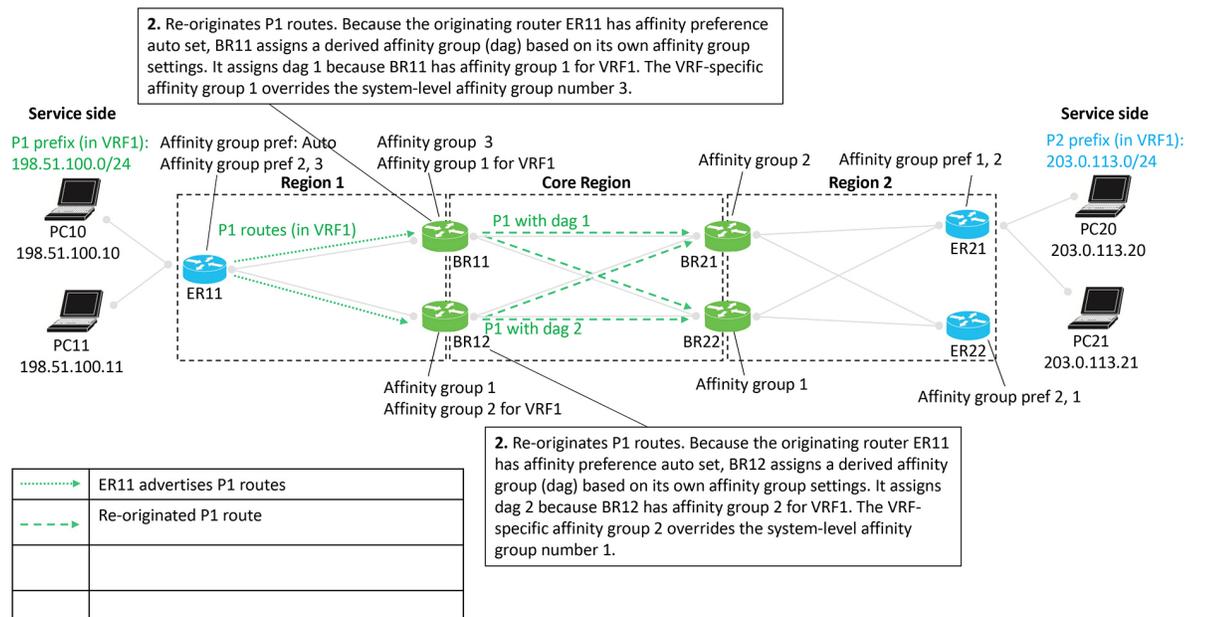


Figure 37: Border Routers BR21 and BR22 Re-Originate the P1 Routes

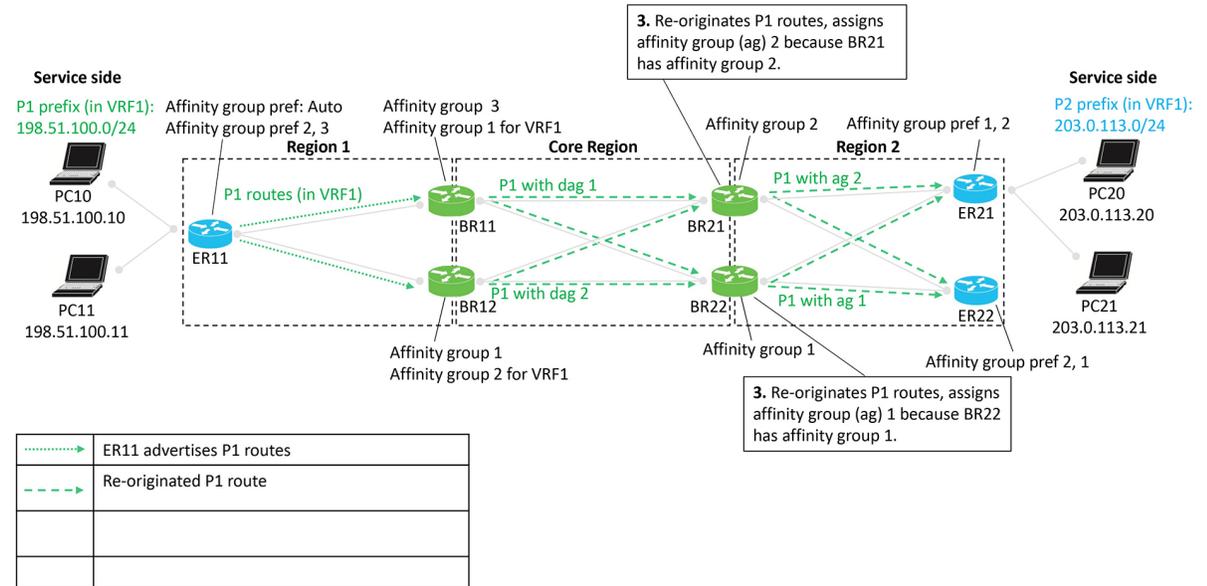


Figure 38: Route Preference According to Affinity Group and Derived Affinity Group

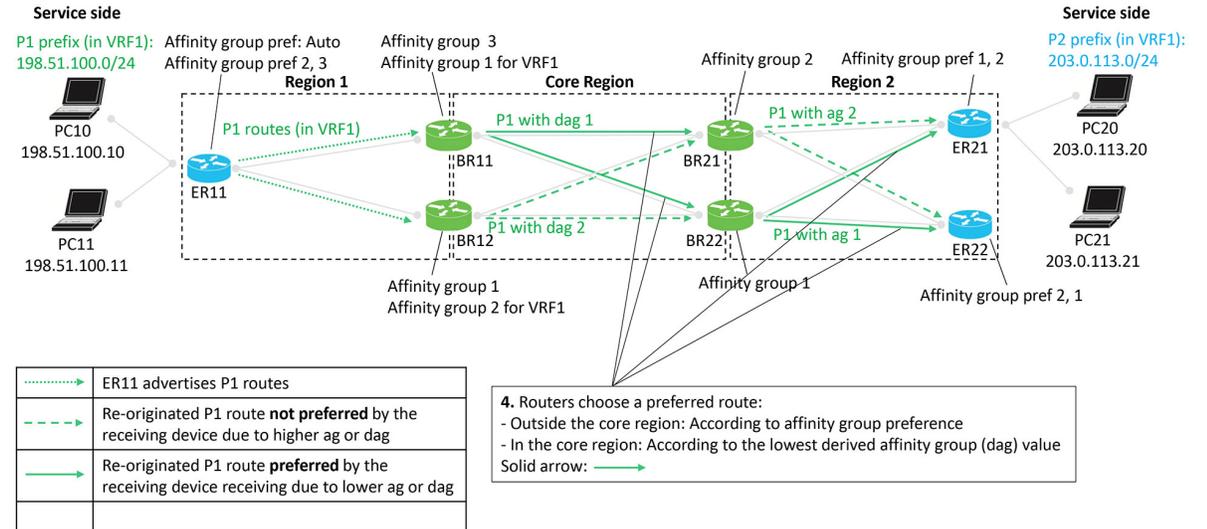
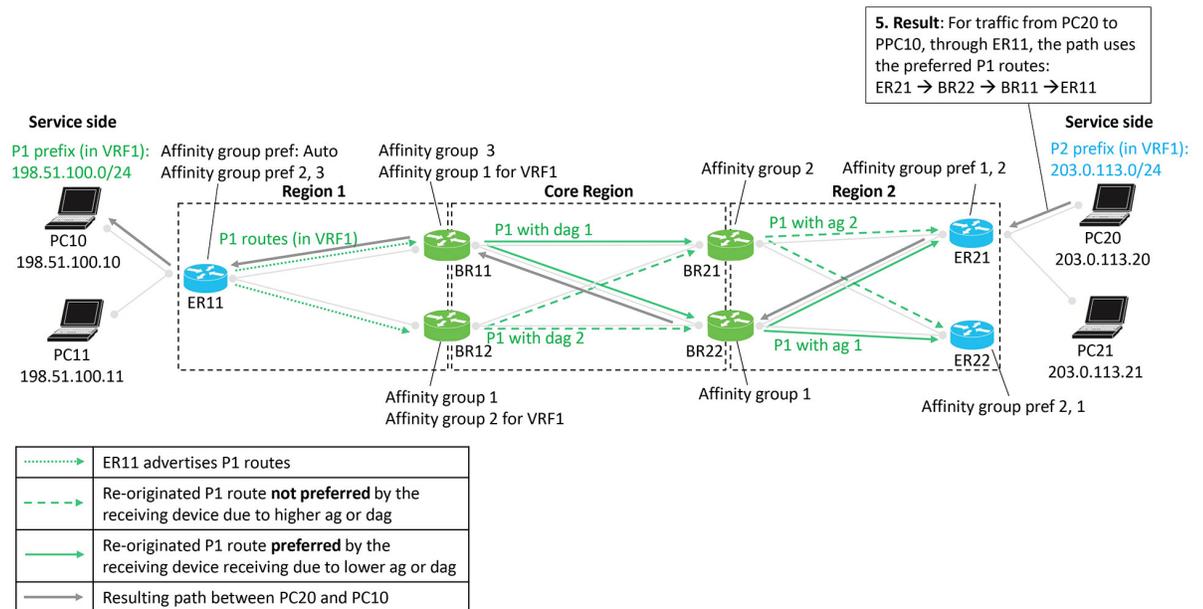


Figure 39: Resulting Path of Traffic to P1



Advertising P2 Routes

Edge routers ER21 and ER22 advertise P2 routes. The process of re-originating these routes to border routers, and eventually to ER11 occurs from right to left in the illustration. In the process, border routers assign affinity groups and derived affinity groups when re-originating routes.

Routers in the network choose preferred routes as follows:

- Outside the core region: According to affinity group preference
- In the core region: According to the lowest derived affinity group (dag) value

Figure 40: Edge Router ER21 Advertises P2 Routes

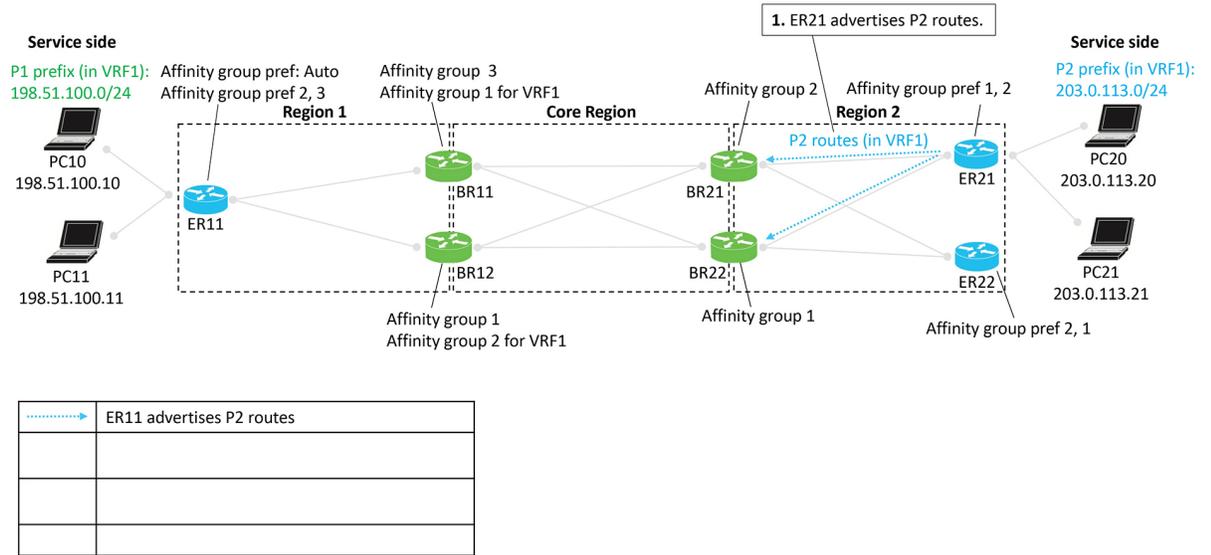


Figure 41: Border Routers BR21 and BR22 Re-Originate the P2 Routes

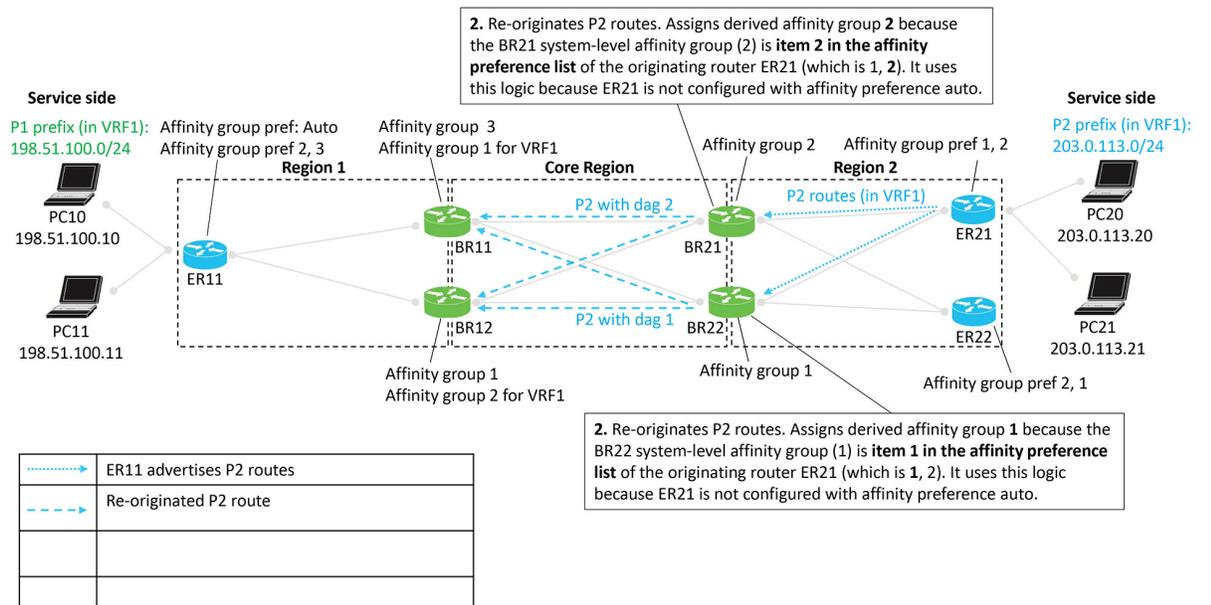


Figure 42: Border Routers BR11 and BR12 Re-Originate the P2 Routes

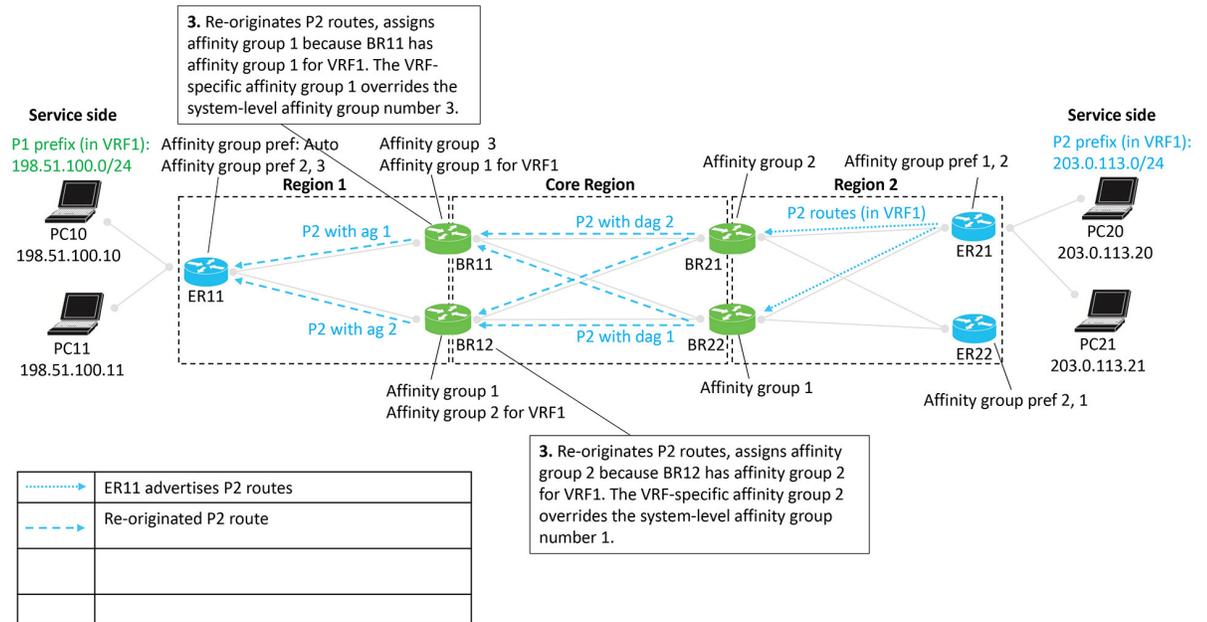


Figure 43: Route Preference According to Affinity Group and Derived Affinity Group

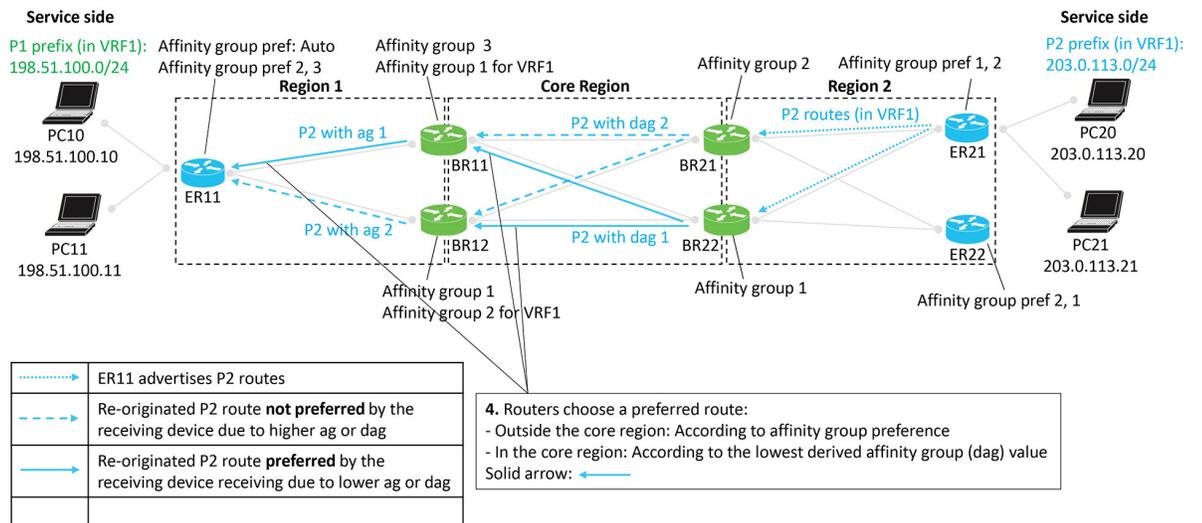
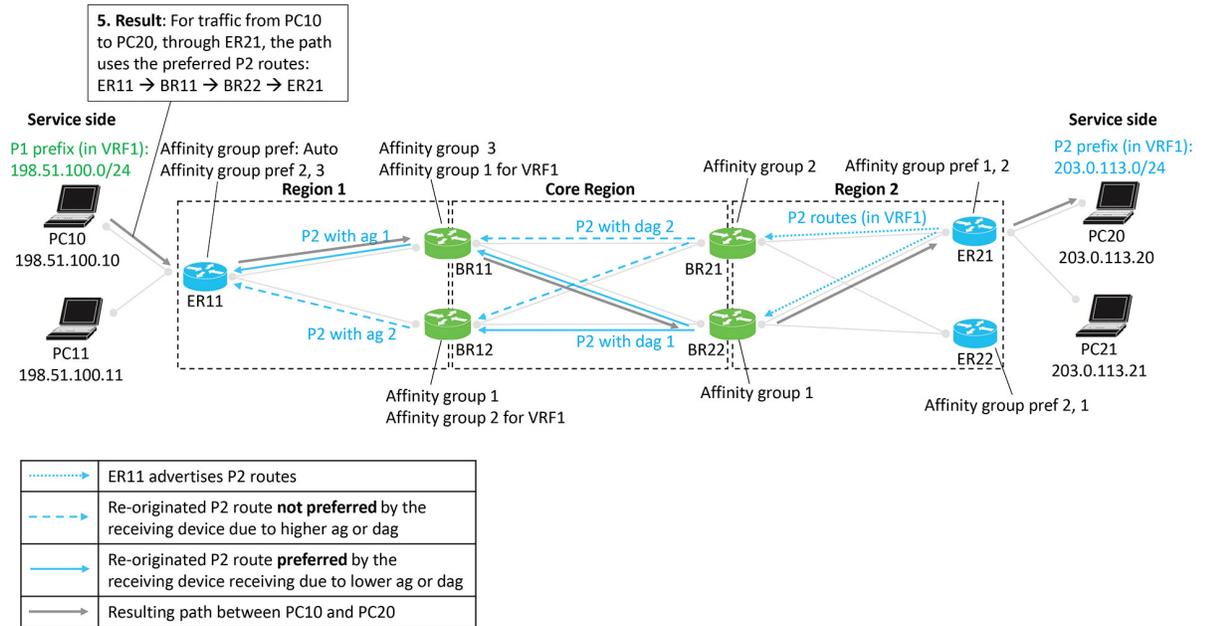


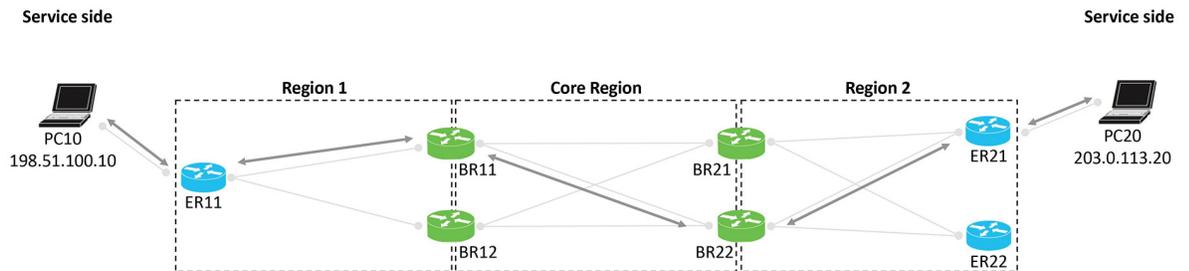
Figure 44: Resulting Path of Traffic to P2



Result

The following figure shows that the result of the configuration is symmetric routing for flows between, in this example, PC10 and PC20:

Figure 45: Result Is Symmetric Routing



Supported Scenarios

The approach to configuring symmetric routing described here applies in the following networking scenarios:

- Hub-and-spoke topology with multiple hub routers
This includes scenarios in which the hub router serves a multi-homed data center.
- Multi-Region Fabric with multiple border routers
This includes scenarios in which a Multi-Region Fabric region includes a multi-homed data center.

- Multi-Region Fabric with transport gateways serving subregions

The following sections briefly describe various specific scenarios, and show example configurations that support symmetric routing in the scenario.

Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active

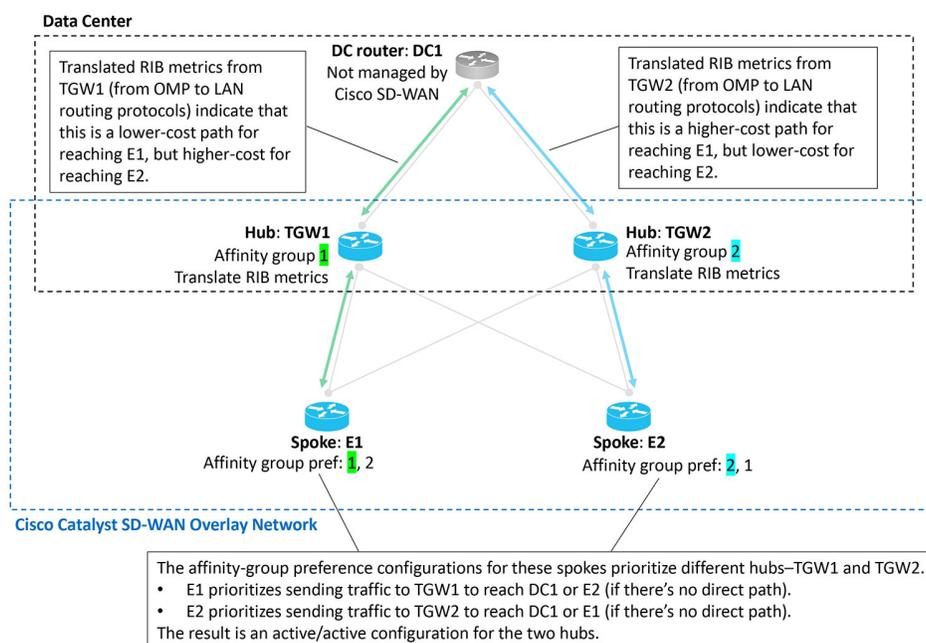
In this scenario, two hubs serve a data center. The two hubs are both active, for an active/active arrangement.

The data center LAN is not part of the Cisco Catalyst SD-WAN overlay network.



Note 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, on page 239](#).

Figure 46: Data Center, Two Hubs, Active/Active

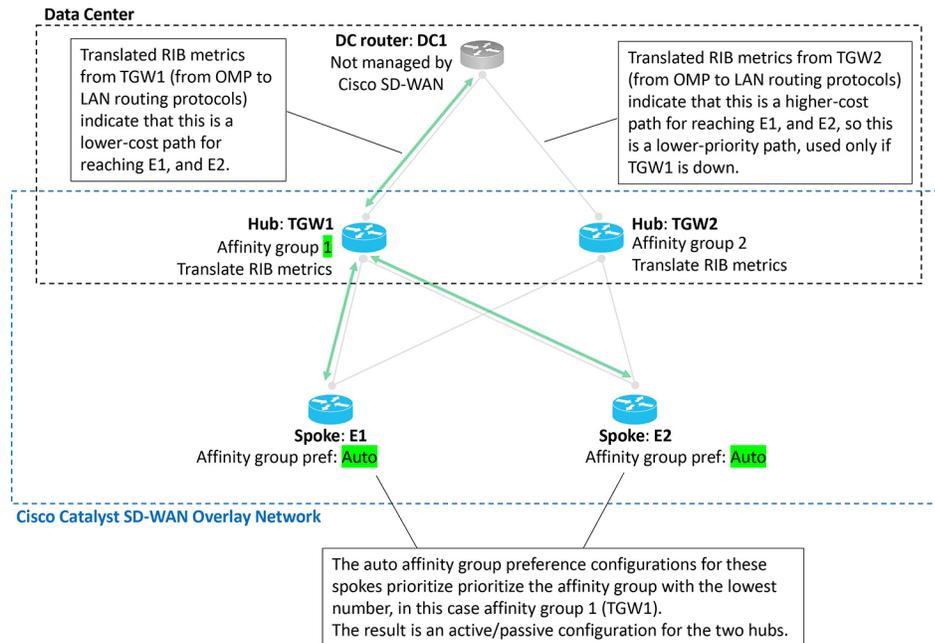


Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Passive

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 47: Data Center, Two Hubs, Active/Passive

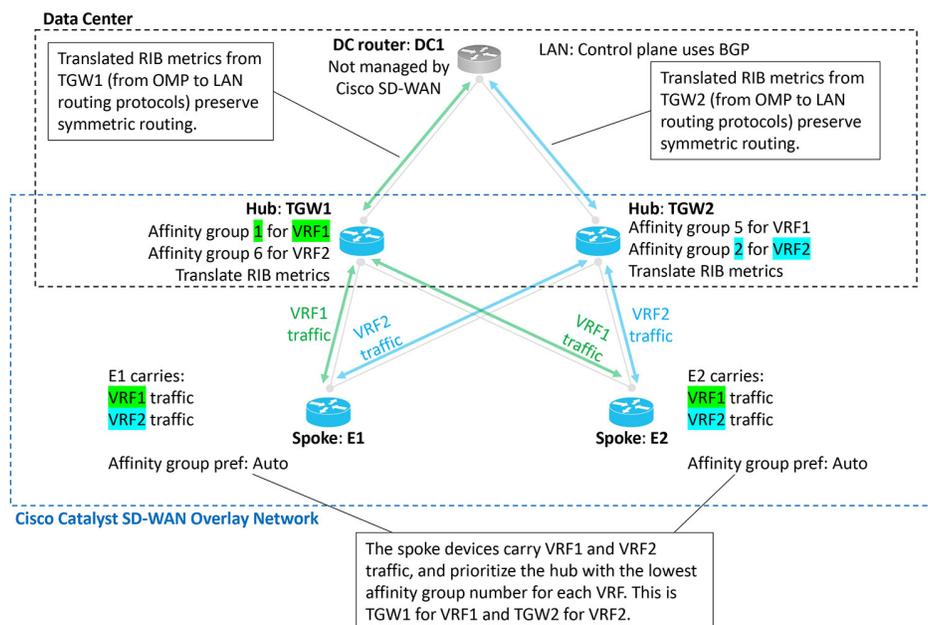


Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active by VRF

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 48: Data Center, Two Hubs, Active/Active, Segregated by VRF



Scenario: Multi-Region Fabric Environment

The [Example of Configuration for Symmetric Routing and the Mechanism](#), on page 223 section describes a Multi-Region Fabric scenario in detail.

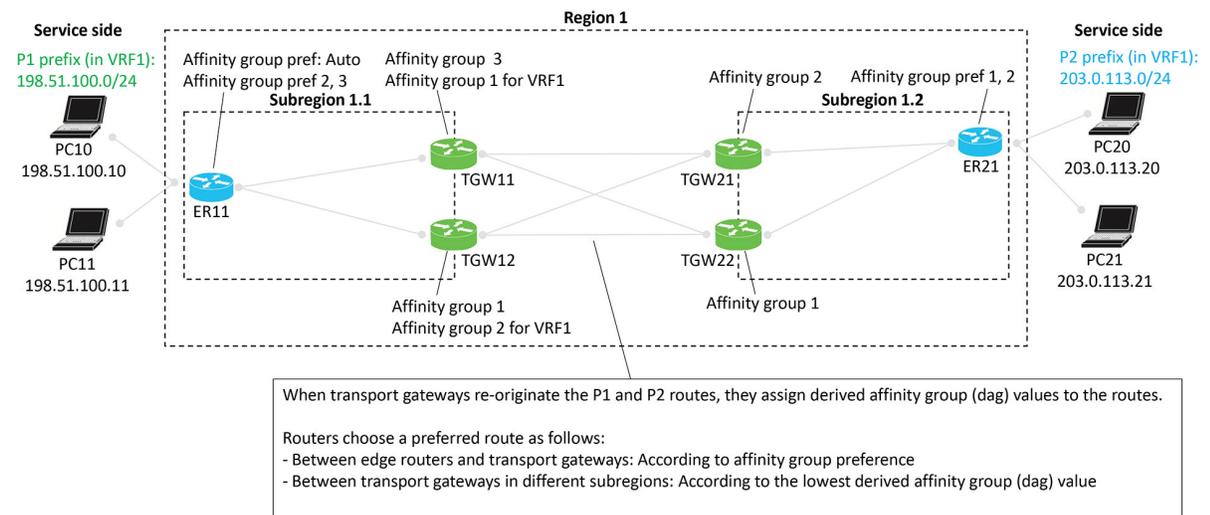
Scenario: Multi-Region Fabric, Transport Gateways Serving Subregions

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](#), on page 223.

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 49: 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, on page 223](#).

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, on page 232](#), but with route leaking. 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
- In this specific scenario, a control policy on the Cisco SD-WAN Controllers provides route leaking from VRF1 to VRF2, and VRF2 to VRF1. Route leaking enables connectivity between endpoints within different VRFs.

This route-leaking scenario illustrates a how transport gateways (or similarly, border routes) assign a derived affinity group (dag) when re-originating routes. The logic is a bit subtle, but this example demonstrates it clearly.

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.

- ER 21: Subscribes only to VRF2, and advertises prefix P2 in VRF2.
- All of the transport gateway routers handle both VRF1 and VRF2 traffic, and therefore re-originate both P1 (in VRF1) and P2 (in VRF2) routes.

By default, the network provides VRF isolation, meaning that when a device advertises routes in various VRFs, Cisco SD-WAN Controllers filter the routes before providing them to other devices. Specifically, a Cisco SD-WAN Controller advertises a VRF *x* route only to devices subscribing to VRF *x*. So in this example, by default, ER11, which subscribes only to VRF1, would not receive P2 routes, which are advertised in VRF2. Similarly, ER21, which subscribes only to VRF2, would not receive P1 routes, which are advertised in VRF1.

Consequently, VRF isolation would prevent traffic flow between ER11 and ER21, which subscribe exclusively to different VRFs.

Route Leaking

Route leaking enables devices to advertise routes across VRFs by exporting ("leaking") a route from one VRF to another.

- Source VRF of a route: Original VRF of the route
- Current VRF of a route: VRF to which the route was exported

When advertising exported routes, routers keep track of the source VRF and the current VRF, so the background of each route is preserved. This point factors in to the logic explained below.

This example has the following route leaking configured:

- An inbound control policy for ER11 configures it to receive VRF1 routes and export them to VRF2. Result: ER11 advertises the P1 prefix in both VRF1 and VRF2 to its associated transport gateways, TGW11 and TGW12.
- An inbound control policy for ER21 configures it to receive VRF2 routes and export them to VRF1. Result: ER21 advertises the P2 prefix in both VRF2 and VRF1 to its associated transport gateways, TGW21 and TGW22.

As mentioned earlier, after the route leaking, devices track, for each route, the source VRF (where the route came from) and the current VRF (VRF to which it was leaked).

Calculating the Derived Affinity Group

A transport gateway device, as in this example, or a border router in a similar example, assigns a derived affinity group (dag) to a route that it is re-originating as follows:

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.

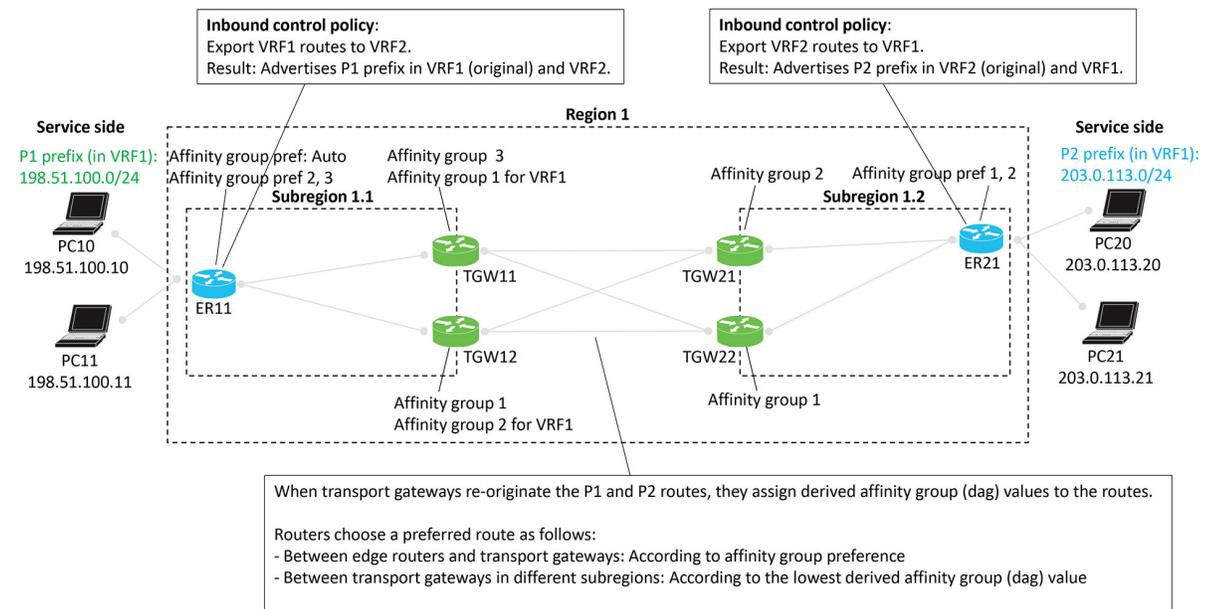
Hypothetically, in the example, if TGW21, in addition to having a system-level affinity group of 2, also had a VRF-specific affinity group of 1 for VRF1, then when TGW21 received from ER21 a P2 route leaked to VRF1, it would consider the preference order of the originating device (ER21). In this hypothetical 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 described in the illustration, a route leaked from VRF2 to VRF1 has source VRF value of 2 and a current VRF value of 1. When a transport gateway re-originates this route, it assigns it a dag according to the number 1, which is the lower of the two VRF numbers. For example, if TGW12 is re-originating a route with a source VRF value of 1 and current VRF value of 2, it chooses 1, which 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 is calculating the dag according to VRF1, it assigns the re-originated route a dag value of 2, taken from the VRF-specific affinity group.

To consider a hypothetical, 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 50: Multi-Region Fabric with Subregions, Route Leaking



Prerequisites for Symmetric Routing

Prerequisite	Description
Transport gateway access to VRFs	For a transport gateway's affinity-group-per-VRF configuration to have effect, the transport gateway requires access to the VRFs for which an affinity group is configured.
Edge routers require affinity group preference	For information, see Configuration Overview, on page 220 .
Transport gateways and border routers require affinity groups	For information, see Configuration Overview, on page 220 .
Transport gateways and border routers conducting traffic with a LAN must redistribute OMP metrics to the LAN	For information, see Configuration Overview, on page 220 .

Restrictions for Symmetric Routing

Restriction	Description
Translating OMP metrics	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, on page 216 .

Configure Symmetric Routing

The following sections describe procedures for the configuration required for symmetric routing.

Configure a Router to Use Automatic Affinity Group Preference Using Cisco SD-WAN Manager

Before You Begin

If you configure a router with an affinity preference order manually and also configure auto preference order, the auto preference order has priority for 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*.

Configure a Router to Use Automatic Affinity Group Preference

1. From the Cisco SD-WAN Manager menu, choose **Configuration > Templates**.
2. Click **Feature Templates**.
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**.
4. In the **Affinity Group Preference Auto** field, choose **On**.
5. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure a Router Affinity Group or Affinity Group Preference

See the following procedures for configuring router affinity groups and affinity group preference:

[Configure an Affinity Group or Affinity Group Preference on a Device, Using Cisco SD-WAN Manager](#)

[Configure an Affinity Group on a Router Using the CLI](#)

[Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager, on page 238](#)

[Configure Router Affinity Groups for Specific VRFs Using a CLI Template, on page 238](#)

[Configure Affinity Group Preference on a Router Using the CLI](#)

[Configure a Router to Use Automatic Affinity Group Preference Using a CLI Template, on page 239](#)

Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager

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



Note 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, on page 238](#).

5. Click **Save** if creating a new template, or **Update** if editing an existing template.

Configure Router Affinity Groups for Specific VRFs Using a CLI Template

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.

1. Enter system configuration mode.

```
system
```

2. Configure an affinity group to apply to a specific VRF or range of VRFs.

```
affinity-per-vrf affinity-group vrf-range vrf-range
```

Example

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
```



Note For information about verifying the VRF-specific affinity group configuration, see [Verify the VRF-Specific Affinity Group Configuration on a Router, on page 242](#).

Configure a Router to Use Automatic Affinity Group Preference Using a CLI Template

Before You Begin

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.

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*.

Configure a Router to Use Automatic Affinity Group Preference

1. Enter system configuration mode.

```
system
```
2. Configure automatic affinity group preference.

```
affinity-group preference-auto
```

Example

```
system
  affinity-group preference-auto
```

Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template

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.



Note This configuration is not available through a feature template.

1. Do one of the following:

- If the underlay network uses the border gateway protocol (BGP), enter router configuration mode and specify a BGP autonomous system. For information about the BGP autonomous system parameter, see the *IP Routing Configuration Guide, Cisco IOS XE 17.x*.

```
router bgp bgp-AS
```

- If the underlay network uses the open shortest path first (OSPF) protocol, enter router configuration mode and specify an OSPF.

```
router ospf process-id [vrf vrf-name]
```

- If the underlay network uses the open shortest path first version 3 (OSPFv3) protocol, enter router configuration mode and specify an OSPFv3.

```
router ospfv3 process-id
```

2. If you specified BGP or OSPFv3 in the previous step, enter address-family mode, specify IPv4 or IPv6, and specify the VRF for which to translate OMP metrics.

```
address-family {ipv4 | ipv6} vrf vrf-name
```

3. Enable translation of OMP route metrics to BGP, OSPF, or OSPFv3 during redistribution of routes to a device outside of the Cisco Catalyst SD-WAN overlay network.



Note 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.

```
redistribute omp translate-rib-metric
```

4. In a scenario in which the underlay network uses BGP, enable propagation of the AS-Path metric. Omitting this causes a router to treat the AS-Path metric as empty.

```
propagate-aspath
```

Example 1

This example applies to a scenario in which the underlay network uses BGP.

```
router bgp 1
  address-family ipv4 vrf 2
    redistribute omp translate-rib-metric
  propagate-aspath
```

Example 2

This example applies to a scenario in which the underlay network uses OSPF.

```
router ospf 1 vrf 1
  redistribute omp translate-rib-metric
```

Example 3

This example applies to a scenario in which the underlay network uses OSPFv3 IPv4.

```
router ospfv3 1
  address-family ipv4 vrf 1
  redistribute omp translate-rib-metric
```

Example 4

This example applies to a scenario in which the underlay network uses OSPFv3 IPv6.

```
router ospfv3 1
  address-family ipv6 vrf 1
  redistribute omp translate-rib-metric
```

Verify Symmetric Routing

The following sections describe procedures for verifying the configurations required for symmetric routing.

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. For information about this command, see [show sdwan omp routes](#) in the *Cisco IOS XE SD-WAN Qualified Command Reference*.

Example

```
Router#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.

Example

```
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 following procedures for configuring VRF-specific affinity groups:

- [Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager, on page 238](#) Configure Router Affinity Groups for Specific VRFs Using Cisco SD-WAN Manager
- [Configure Router Affinity Groups for Specific VRFs Using a CLI Template, on page 238](#) Configure Router Affinity Groups for Specific VRFs Using a CLI Template



Note You can define VRF-specific affinity groups on a router without any requirement that the particular VRF exists.

Example

```
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](#) in the *Cisco SD-WAN Policies Configuration Guide, Cisco IOS XE Release 17.x*.

Use **show running-config apply-policy** on a Cisco SD-WAN Controller to show the sites to which a control policy is applied.

Example 1

The following 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.

Example

In the following 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, on page 216](#).

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 following example 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

```
Router#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.

```
Router#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.

```
Router#show ipv6 route vrf 1 omp
m   2001:DB8::/128 [251/66]
    via 172.16.0.1%default
...
```

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.

The following example shows the derived BGP metrics for the 10.1.1.1 route. This example shows an IPv4 route, but 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.)

```
Router#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
```

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.

```
Router#show ip ospf
Routing Process "ospf 10" with ID 10.100.10.1
...
Redistributing External Routes from,
  omp, includes subnets in redistribution, translate rib metric
  Maximum limit of redistributed prefixes 10240
  Threshold for warning message 75%
```

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

To 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 13

Troubleshoot Cisco Catalyst SD-WAN Routing



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

- [Overview, on page 247](#)
- [Support Articles, on page 247](#)
- [Feedback Request, on page 248](#)
- [Disclaimer and Caution, on page 248](#)

Overview

This chapter provides links to documents authored by Cisco subject matter experts (SMEs). They aim to help you resolve technical issues without requiring a support ticket. If these documents are unable to resolve your issue, we recommend visiting the applicable [Cisco Community](#). There is a wealth of information and advice available from fellow Cisco customers who may have experienced this issue already and provided a solution. If you are not able to find a resolution on the Community, it may be best that you raise a support ticket at [Cisco Support](#). In cases where a support ticket has to be raised, these documents provide guidance about the data that should be collected and added to the support ticket. Specify the support document you referred, and TAC can create an improvement request with the document owner.

Support Articles

The documents in this section were created using specific software and hardware listed in the Components Used section of each article. However, this does not mean that they are limited to what is listed in Components Used, and generally remain relevant for later versions of software and hardware. Note that there could be some changes in the software or hardware that can cause commands to stop working, the syntax to change, or GUIs and CLIs to look different from one release to another.

The following are the support articles associated with this technology:

Document	Description
Exclude Routes from Redistributing into OMP	This document describes how to exclude unwanted routes from being redistributed into Overlay Management Protocol (OMP).
Cisco IOS-XE Catalyst SD-WAN Installs OSPF External Route with DN-Bit	This document describes the expected behavior of Cisco IOS [®] -XE SD-WAN software when Open Shortest Path First (OSPF) external routes are installed into the routing table.
Troubleshoot OMP Route Instability in Failover Scenario	This document describes how to troubleshoot Overlay Management Protocol (OMP) routes and explains Cisco SD-WAN Controller route selection order of operations.
How to Avoid BGP-OMP Routing Loop in SD-WAN Overlay at Dual-Homed Sites with Two Routers	This document describes how to avoid a routing loop in SD-WAN fabric when Border Gateway Protocol (BGP) routing and Site of Origin (SoO) is used.
OMP Best Path Selection Peculiarities and Typical Confusions	This document describes a typical misunderstanding of the Overlay Management Protocol (OMP) best-path selection and order of operation between OMP best-path selection, egress policy, and send-path-limit feature.
Quick Start Guide - Data Collection for Various SD-WAN Issues	This document describes several Cisco Catalyst SD-WAN issues along relevant data that must be collected in advance before you open a TAC case to improve the speed of troubleshooting and/or problem resolution.
Collect an Admin-Tech in SDWAN Environment and Upload to TAC Case	This document describes how to initiate an <code>admin-tech</code> in an Cisco Catalyst SD-WAN environment.

Feedback Request

Your input helps. A key aspect to improving these support documents is customer feedback. Note that these documents are owned and maintained by multiple teams within Cisco. If you find an issue specific to the document (unclear, confusing, information missing, etc):

- Provide feedback using the **Feedback** button located at the right panel of the corresponding article. The document owner will be notified, and will either update the article, or flag it for removal.
- Include information regarding the section, area, or issue you had with the document and what could be improved. Provide as much detail as possible.

Disclaimer and Caution

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.