



Symmetric Routing

- [Symmetric Routing, on page 1](#)
- [Information About Symmetric Routing, on page 1](#)
- [Configuration Overview, on page 9](#)
- [Supported Scenarios, on page 18](#)
- [Prerequisites for Symmetric Routing, on page 25](#)
- [Restrictions for Symmetric Routing, on page 26](#)
- [Configure Symmetric Routing, on page 26](#)
- [Verify Symmetric Routing, on page 30](#)
- [Monitor RIB Metric Translation, on page 33](#)

Symmetric Routing

Table 1: 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 1: 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 2: 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 25, 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>

Mechanism	Description
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 27 • Configure Router Affinity Groups for Specific VRFs Using a CLI Template, on page 27 <p>Uses the affinity-per-vrf <i>affinity-group vrf-range vrf-range</i> command.</p>
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 26.)

Mechanism	Description
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 <p>Uses the affinity-group preference list command.</p> <ul style="list-style-type: none"> • (Automatic configuration) Configure a Router to Use Automatic Affinity Group Preference Using Cisco SD-WAN Manager, on page 26 • (Automatic configuration) Configure a Router to Use Automatic Affinity Group Preference Using a CLI Template, on page 28 <p>Uses the affinity-group preference-auto command.</p>
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 5.</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 28 <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 3: Translating OMP Metrics



Related Topics

[Translating OMP Metrics to BGP Attributes, on page 6](#)

[Translating OMP Metrics to an OSPF Metric, on page 9](#)

[Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template, on page 28](#)

[Monitor RIB Metric Translation, on page 33](#)

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 33](#).

The translation from OMP to BGP is as follows:

Table 2: 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	Two possibilities: <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 3: Translation from OMP to BGP without RIB Metric Translation Enabled

Example	OMP Metrics: Example Combinations		Translation to BGP Attributes: propagate- <i>aspath</i> enabled translate- <i>rib-metric</i> not enabled		
	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)

	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
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 4: 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)
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)

	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
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 33](#).

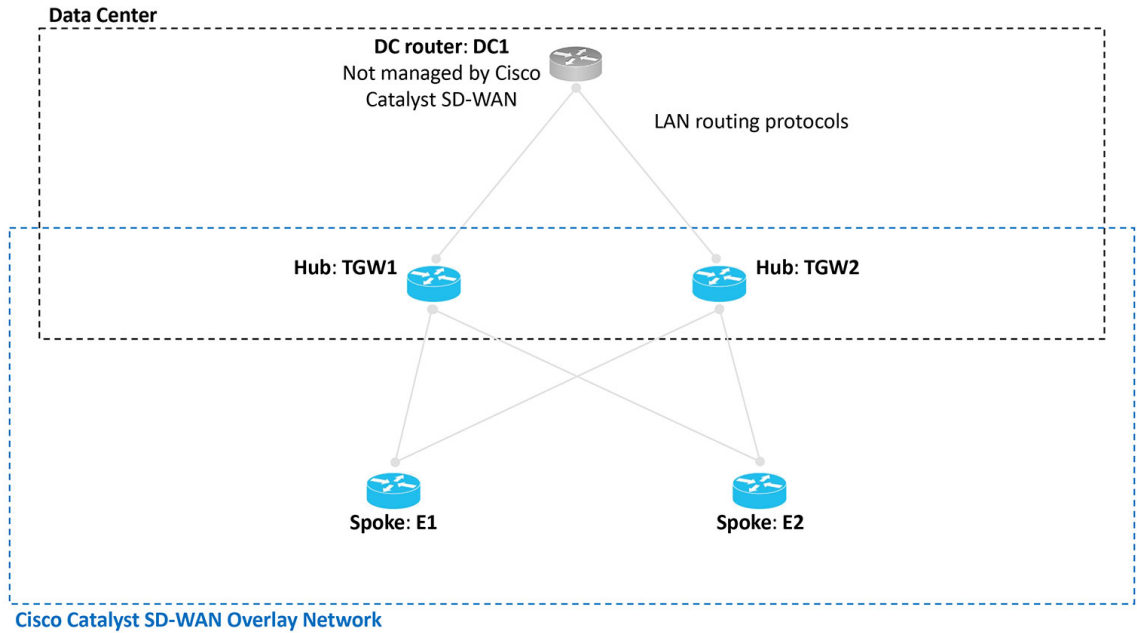
Configuration Overview

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 4: 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 5: 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 27.</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 27.</p>
3. Enable translation of RIB metrics	Transport gateways Border routers in a Multi-Region Fabric scenario	<p>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.</p> <p>For a full explanation, see Translating OMP Metrics for Devices Outside of the Overlay Network, on page 5.</p> <p>For configuration instructions, see Configure a Router to Translate OMP Metrics to BGP or OSPF Using a CLI Template, on page 28.</p>

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 6: Transport Gateway Scenario with a Data Center LAN, Showing a Configuration for Symmetric Routing

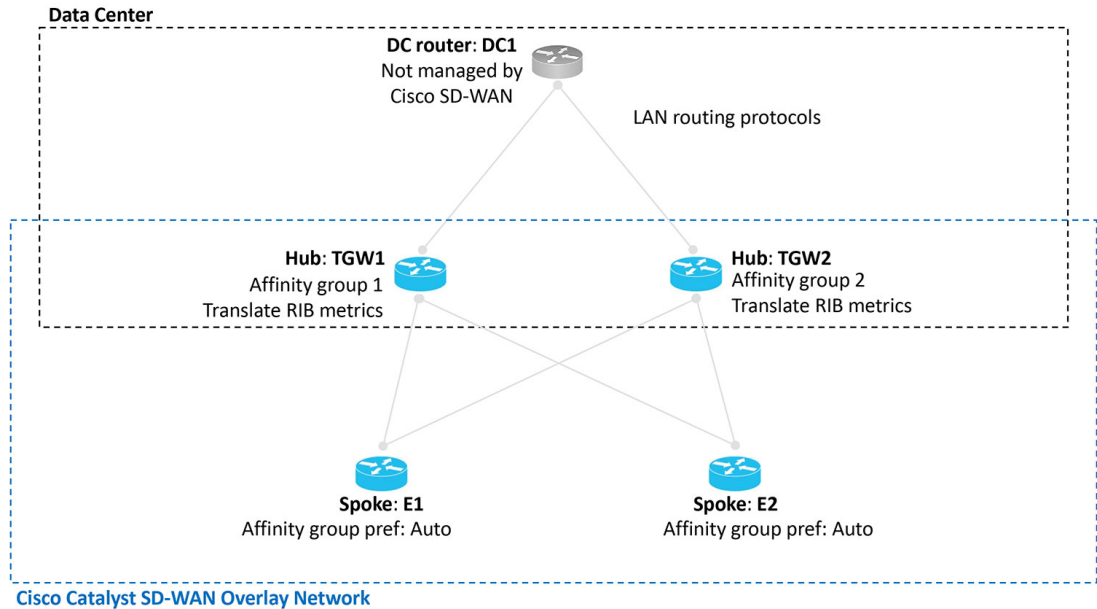



Figure 7: Multi-Region Fabric Scenario, Showing a Configuration for Symmetric Routing

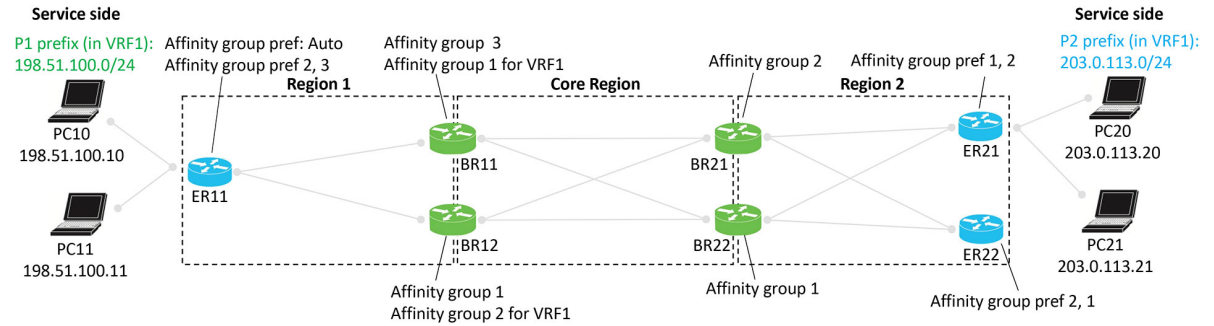


Example of Configuration for Symmetric Routing and the Mechanism

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 8: Multi-Region Fabric Scenario, Configuration for Symmetric Routing



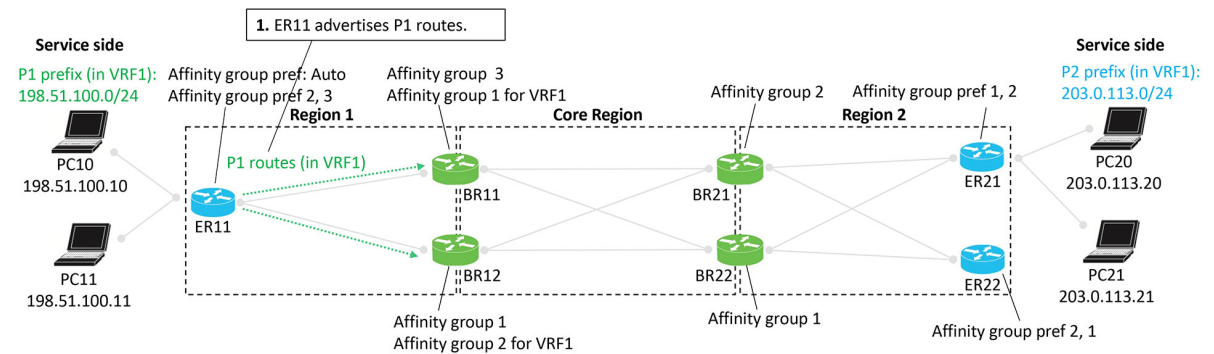
Advertising P1 Routes

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 9: Edge Router ER11 Advertises P1 Routes



	ER11 advertises P1 routes

Figure 10: Border Routers BR11 and BR12 Re-Originate the P1 Routes

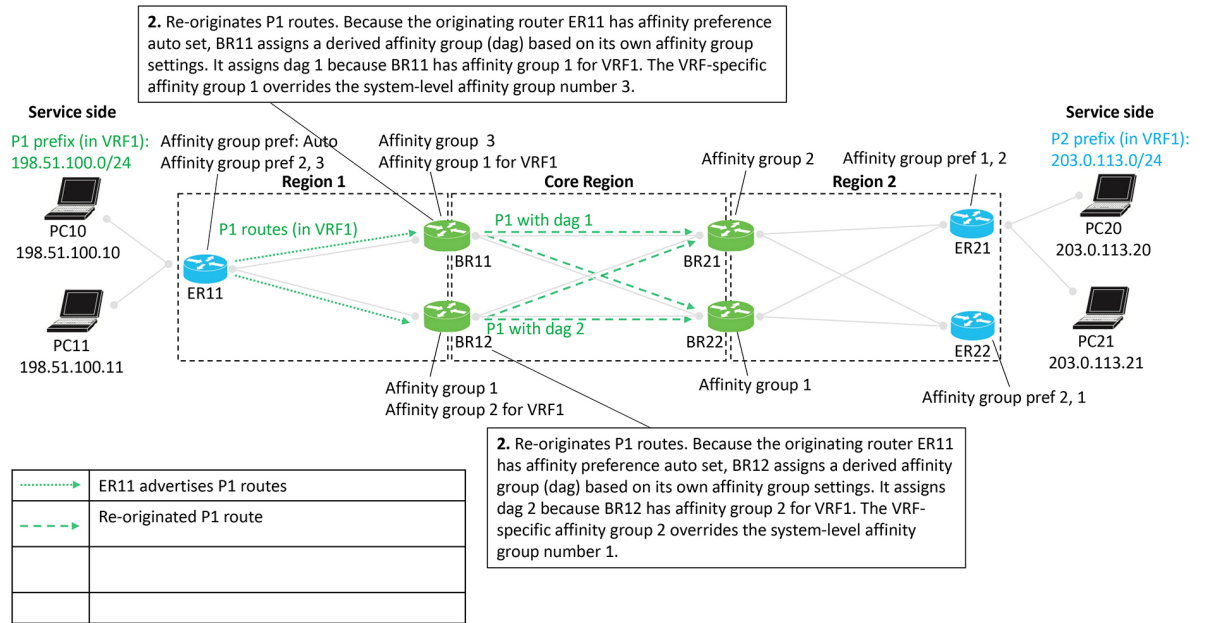


Figure 11: Border Routers BR21 and BR22 Re-Originate the P1 Routes

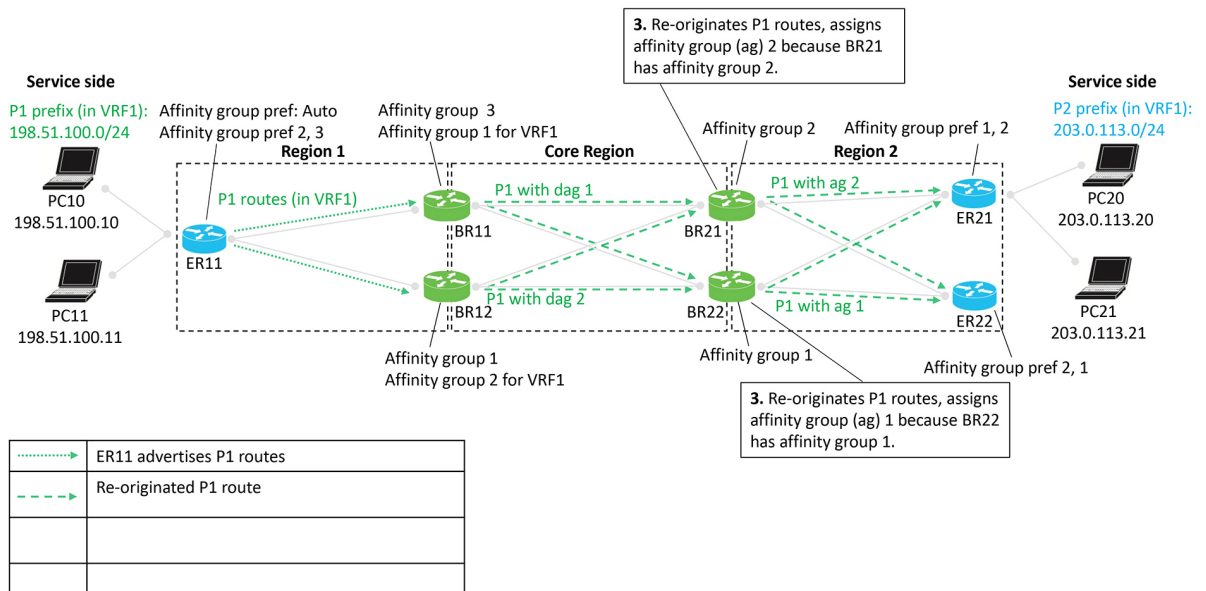


Figure 12: Route Preference According to Affinity Group and Derived Affinity Group

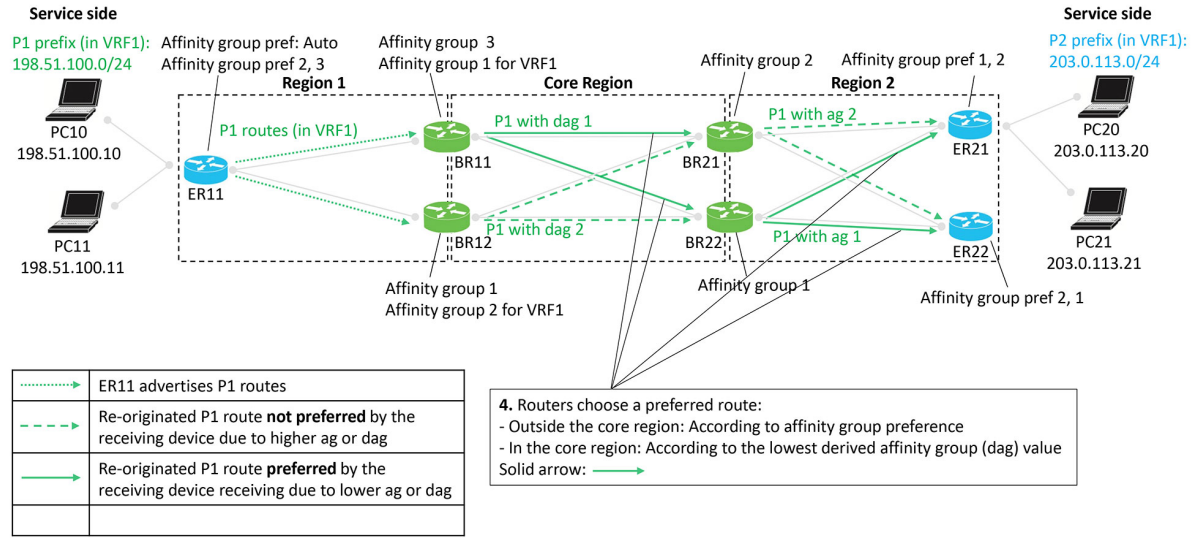
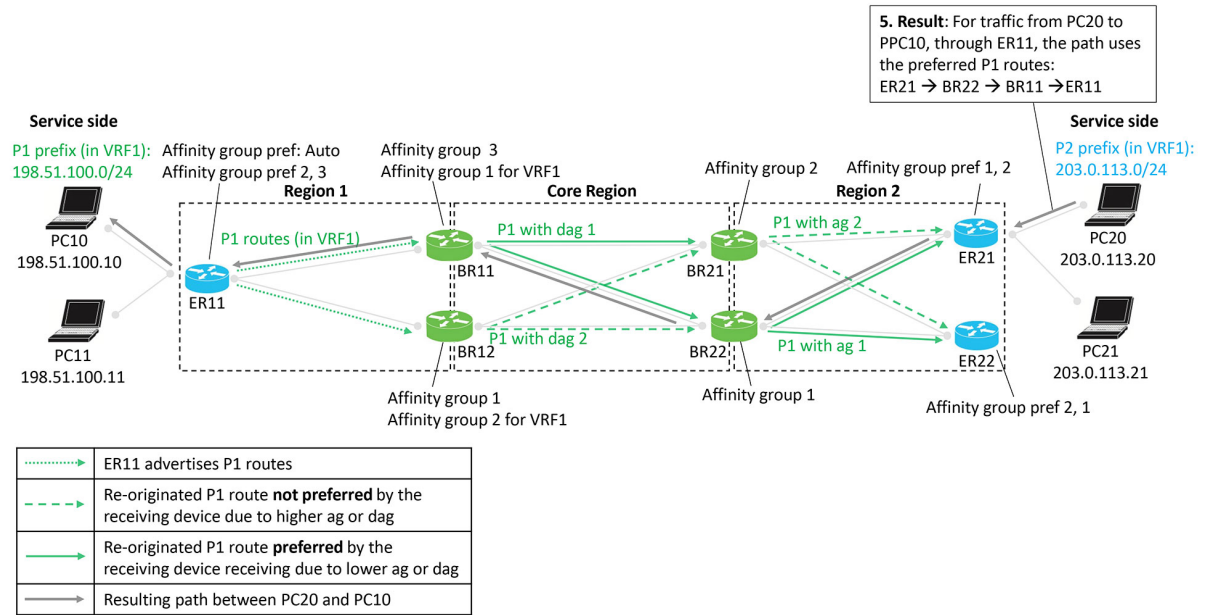


Figure 13: Resulting Path of Traffic to P1



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 14: Edge Router ER21 Advertises P2 Routes

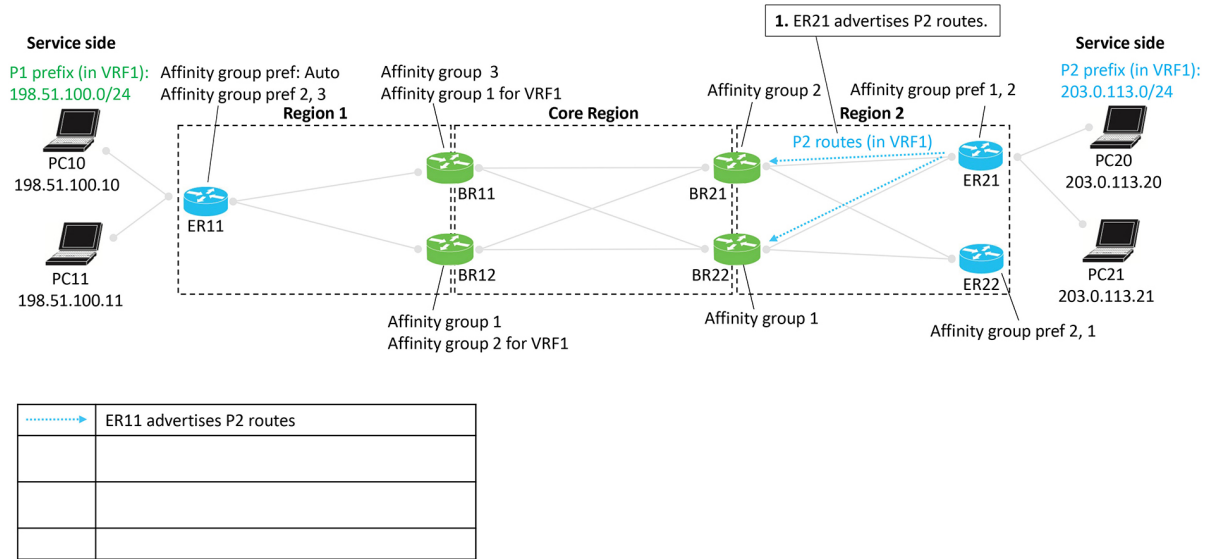


Figure 15: Border Routers BR21 and BR22 Re-Originate the P2 Routes

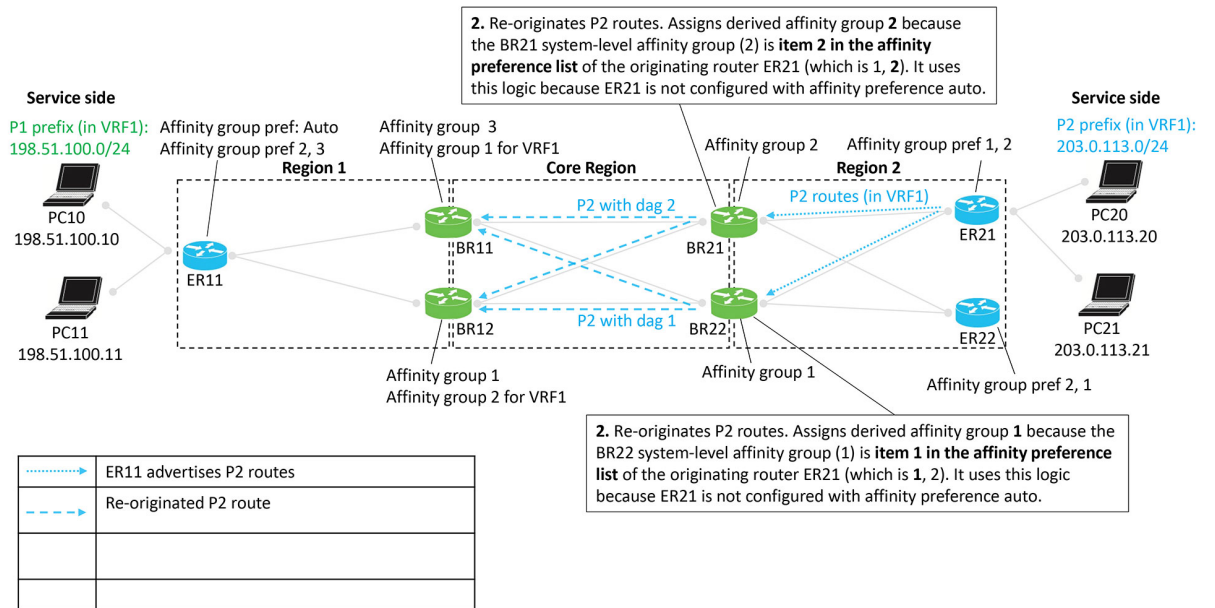


Figure 16: Border Routers BR11 and BR12 Re-Originate the P2 Routes




Figure 17: Route Preference According to Affinity Group and Derived Affinity Group

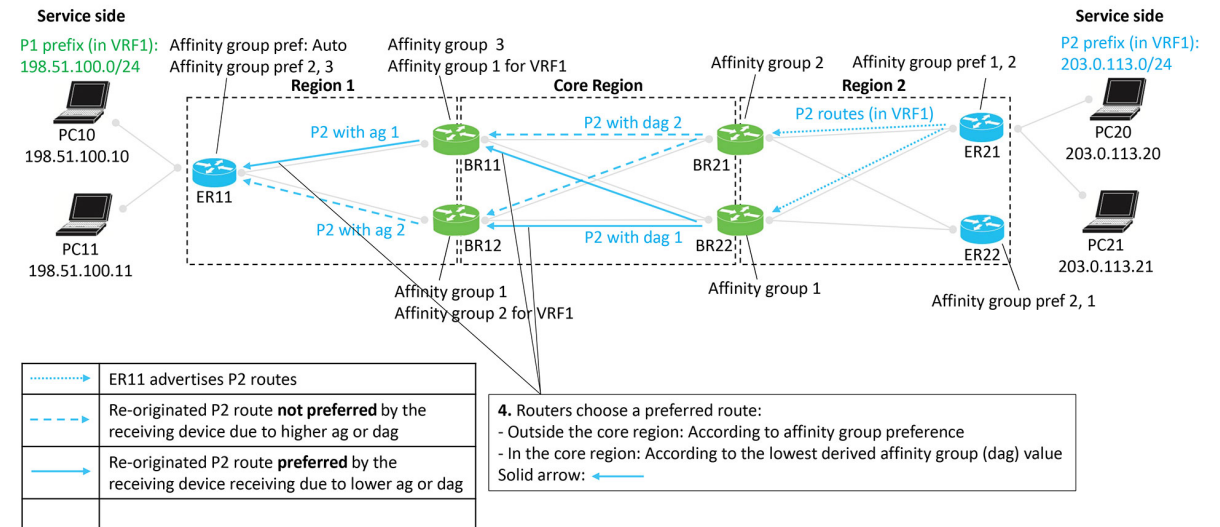
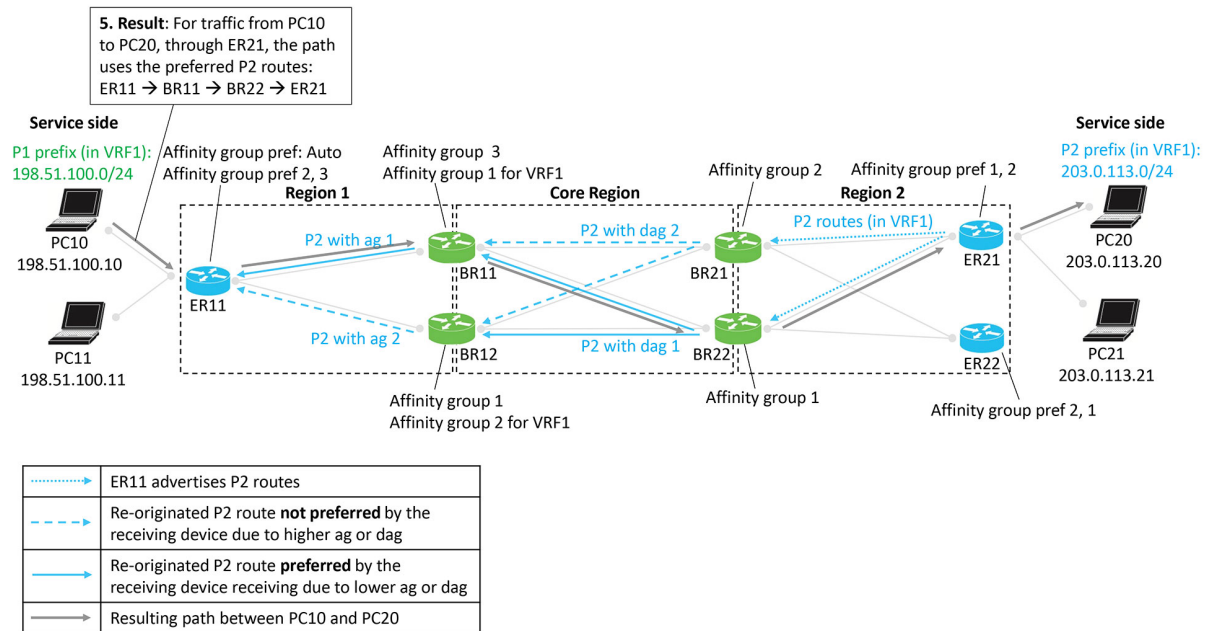


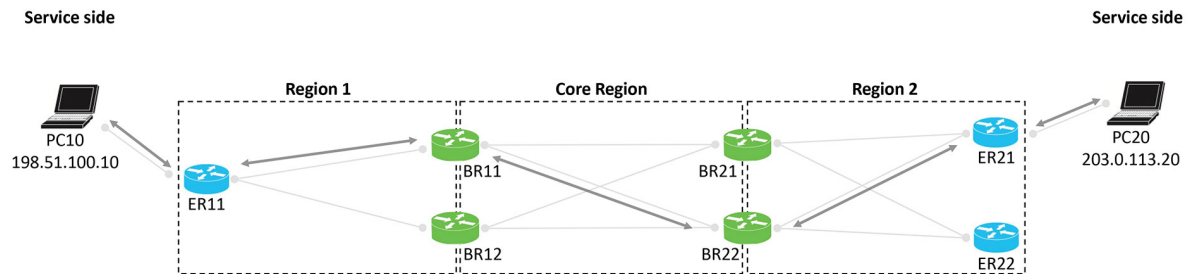
Figure 18: Resulting Path of Traffic to P2



Result

The following figure shows that the result of the configuration is symmetric routing for flows between, in this example, PC10 and PC20:

Figure 19: 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.

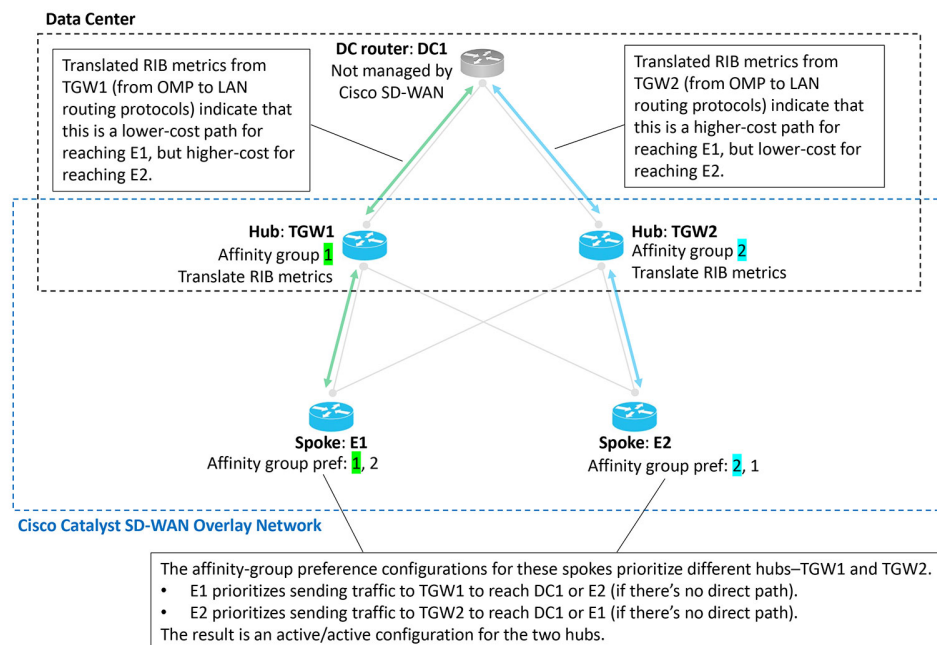
Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active

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 28](#).

Figure 20: Data Center, Two Hubs, Active/Active

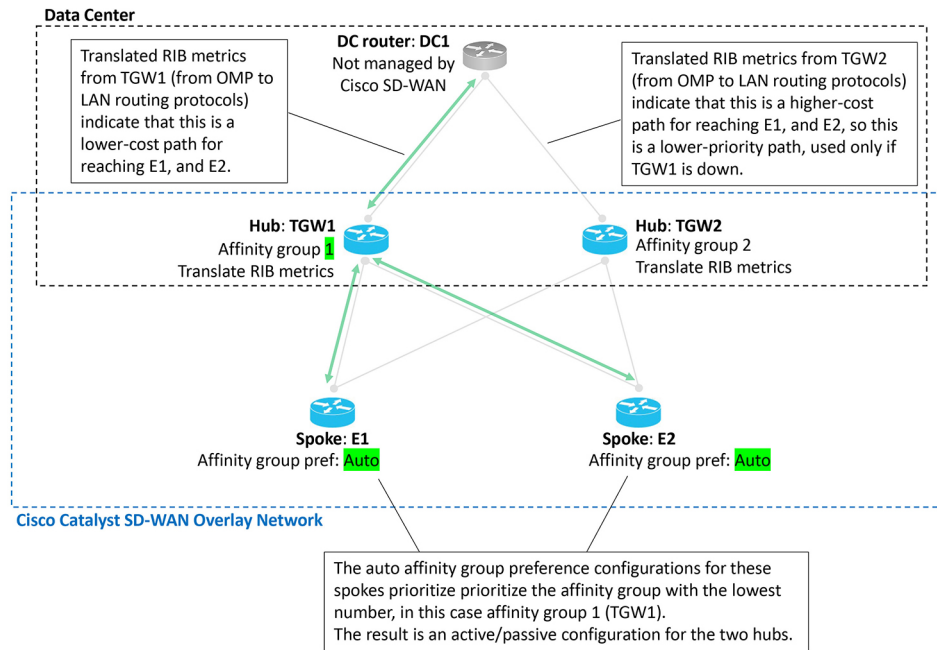


Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Passive

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 21: Data Center, Two Hubs, Active/Passive

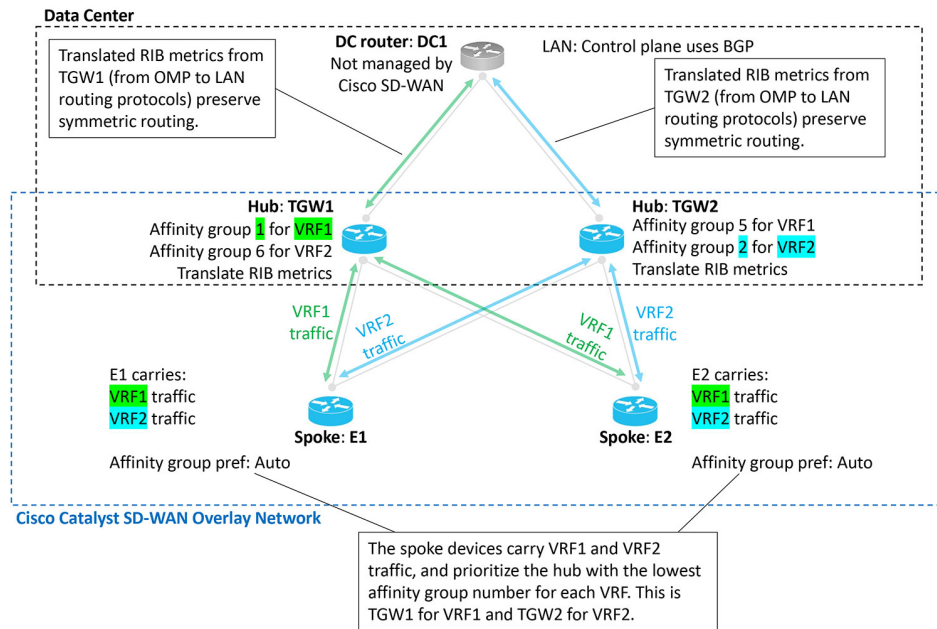


Scenario: Hub-and-Spoke Topology, Multiple Hubs Serving a Data Center, Active/Active by VRF

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 22: 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 12 section describes a Multi-Region Fabric scenario in detail.

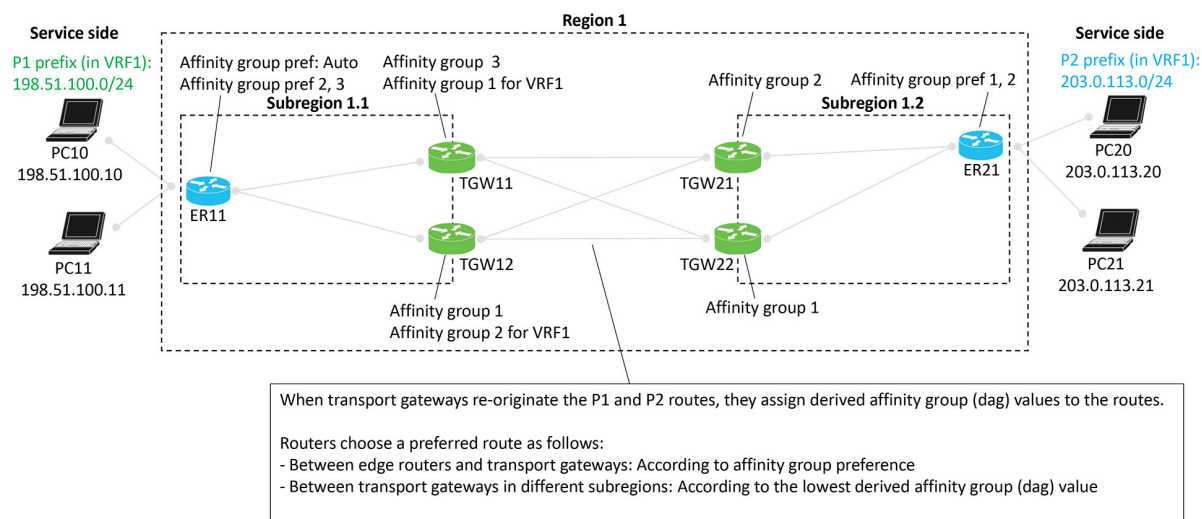
Scenario: Multi-Region Fabric, Transport Gateways Serving Subregions

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 12.

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 23: 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 12](#).

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 21](#), 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.

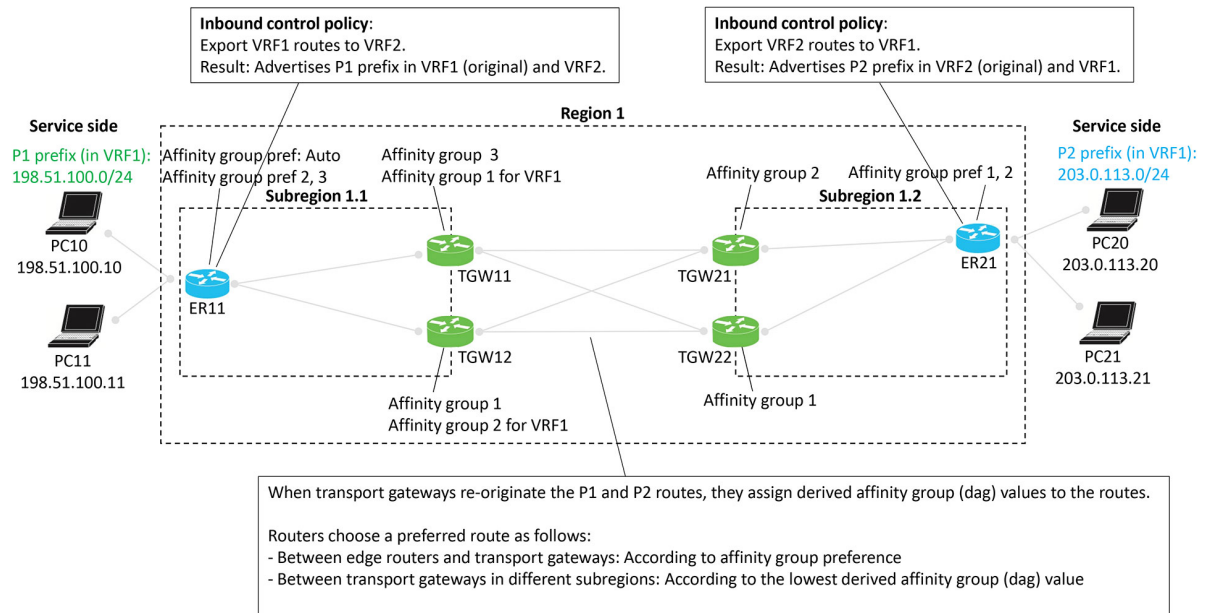
Hypothetically, in the example, if TGW21, in addition to having a system-level affinity group of 2, also had a VRF-specific affinity group of 1 for VRF1, then when TGW21 received from ER21 a P2 route leaked to VRF1, it would consider the preference order of the originating device (ER21). In this hypothetical 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 24: 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 9 .
Transport gateways and border routers require affinity groups	For information, see Configuration Overview, on page 9 .
Transport gateways and border routers conducting traffic with a LAN must redistribute OMP metrics to the LAN	For information, see Configuration Overview, on page 9 .

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 5 .

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 27](#)

[Configure Router Affinity Groups for Specific VRFs Using a CLI Template, on page 27](#)

[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 28](#)

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 27](#).

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 31](#).

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 27](#) 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 27](#) 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 5](#).

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
...
```

