

Sites Connected via SR-MPLS

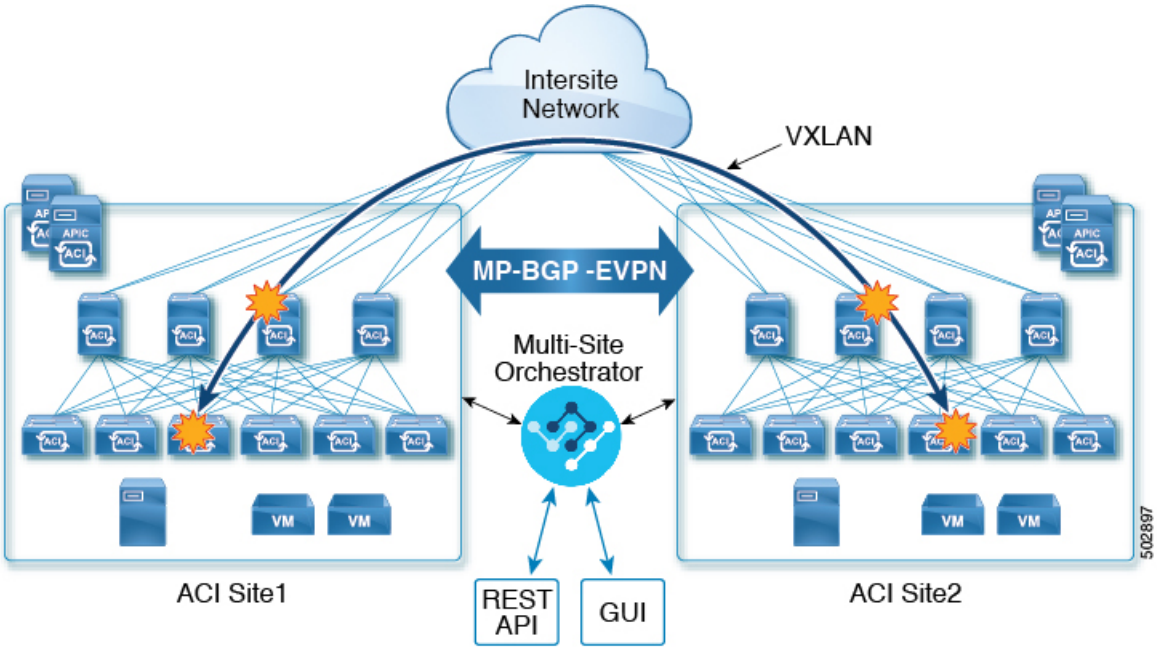
- SR-MPLS and Multi-Site, on page 1
- Infra Configuration, on page 3
- SR-MPLS Tenant Requirements and Guidelines, on page 10
- Creating SR-MPLS Route Map Policy, on page 12
- Enabling Template for SR-MPLS, on page 14
- Creating VRF and SR-MPLS L3Out, on page 14
- Configuring Site-Local VRF Settings, on page 15
- Configuring Site-Local SR-MPLS L3Out Settings, on page 15
- Communication Between EPGs Separated by MPLS Network, on page 16
- Deploying Configuration, on page 17

SR-MPLS and Multi-Site

Starting with Multi-Site Orchestrator, Release 3.0(1) and APIC Release 5.0(1), the Multi-Site architecture supports APIC sites connected via MPLS networks.

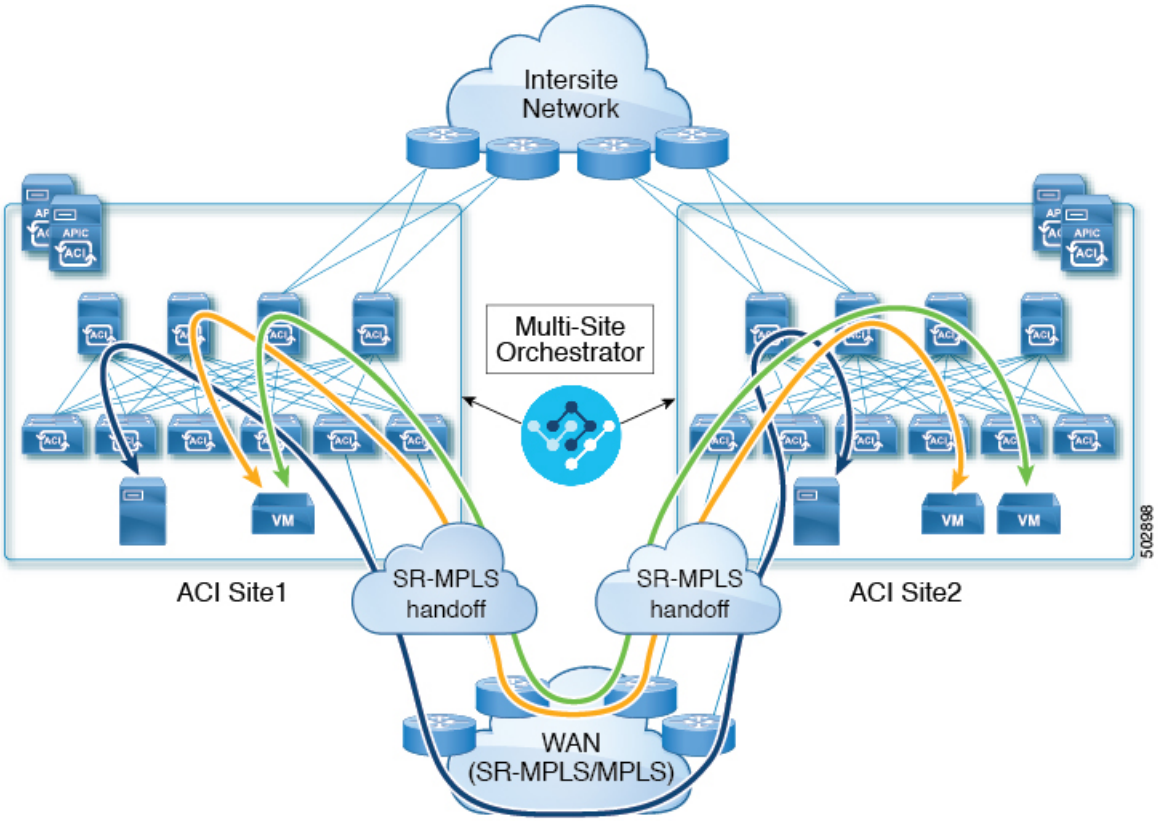
In a typical Multi-Site deployment, traffic between sites is forwarded over an intersite network (ISN) via VXLAN encapsulation:

Figure 1: Multi-Site and ISN



With Release 3.0(1), MPLS network can be used in addition to or instead of the ISN allowing inter-site communication via WAN:

Figure 2: Multi-Site and MPLS



The following sections describe guidelines, limitations, and configurations specific to managing Schemas that are deployed to these sites from the Multi-Site Orchestrator. Detailed information about MPLS hand off, supported individual site topologies (such as remote leaf support), and policy model is available in the *Cisco APIC Layer 3 Networking Configuration Guide*.

Infra Configuration

SR-MPLS Infra Guidelines and Limitations

If you want to add an APIC site that is connected to an SR-MPLS network to be managed by the Multi-Site Orchestrator, keep the following in mind:

- Any changes to the topology, such as node updates, are not reflected in the Orchestrator configuration until site configuration is refreshed, as described in Refreshing Site Connectivity Information.
- Objects and policies deployed to a site that is connected to an SR-MPLS network cannot be stretched to other sites.

  When you create a template and specify a Tenant, you will need to enable the `SR-MPLS` option on the tenant. You will then be able to map that template only to a single ACI site.

- Tenants deployed to a site that is connected via an SR-MPLS network will have a set of unique configuration options specifically for SR-MPLS configuration. For information about configuring tenants, see Tenants.

Supported Hardware

The SR-MPLS connectivity is supported for the following platforms:

- **Border Leaf switches**: The "FX", "FX2", and "GX" switch models.
- **Spine switches**:
  - Modular spine switch models with "LC-EX", "LC-FX", and "GX" at the end of the linecard names.
  - The Cisco Nexus 9000 series N9K-C9332C and N9K-C9364C fixed spine switches.
- **DC-PE routers**:
  - Network Convergence System (NCS) 5500 Series
  - ASR 9000 Series
  - NCS 540 or 560 routers

SR-MPLS Infra L3Out

You will need to create an SR-MPLS Infra L3Out for the fabrics connected to SR-MPLS networks as described in the following sections. When creating an SR-MPLS Infra L3Out, the following restrictions apply:

- Each SR-MPLS Infra L3Out must have a unique name.
- You can have multiple SR-MPLS infra L3Outs connecting to different routing domains, where the same border leaf switch can be in more than one L3Out, and you can have different import and export routing policies for the VRFs toward each routing domain.
- Even though a border leaf switch can be in multiple SR-MPLS infra L3Outs, a border leaf switch/provider edge router combination can only be in one SR-MPLS infra L3Out as there can be only one routing policy for a user VRF/border leaf switch/DC-PE combination.
- If there is a requirement to have SR-MPLS connectivity from multiple pods and remote locations, ensure that you have a different SR-MPLS infra L3Out in each of those pods and remote leaf locations with SR-MPLS connectivity.
- If you have a multi-pod or remote leaf topology where one of the pods is not connected directly to the SR-MPLS network, that pod's traffic destined for the SR-MPLS network will use standard IPN path to another pod, which has an SR-MPLS L3Out. Then the traffic will use the other pod's SR-MPLS L3Out to reach its destination across SR-MPLS network.
- Routes from multiple VRFs can be advertised from one SR-MPLS Infra L3Out to provider edge (PE) routers connected to the nodes in this SR-MPLS Infra L3Out.

  PE routers can be connected to the border leaf directly or through other provider (P) routers.
- The underlay configuration can be different or can be the same across multiple SR-MPLS Infra L3Outs for one location.

For example, assume the same border leaf switch connects to PE-1 in domain 1 and PE-2 in domain 2, with the underlay connected to another provider router for both. In this case, two SR-MPLS Infra L3Outs will be created: one for PE-1 and one for PE-2. But for the underlay, it's the same BGP peer to the provider router. Import/export route-maps will be set for EVPN session to PE-1 and PE-2 based on the corresponding route profile configuration in the user VRF.

Guidelines and Limitations for MPLS Custom QoS Policies

Following is the default MPLS QoS behavior:

- All incoming MPLS traffic on the border leaf switch is classified into QoS Level 3 (the default QoS level).
- The border leaf switch will retain the original DSCP values for traffic coming from SR-MPLS without any remarking.
- The border leaf switch will forward packets with the default MPLS EXP (0) to the SR-MPLS network.

Following are the guidelines and limitations for configuring MPLS Custom QoS policies:

- Data Plane Policers (DPP) are not supported at the SR-MPLS L3Out.
- Layer 2 DPP works in the ingress direction on the MPLS interface.
- Layer 2 DPP works in the egress direction on the MPLS interface in the absence of an egress custom MPLS QoS policy.
- VRF level policing is not supported.

Creating SR-MPLS QoS Policy

This section describes how to configure SR-MPLS QoS policy for a site that is connected via an MPLS network. If you have no such sites, you can skip this section.

SR-MPLS Custom QoS policy defines the priority of the packets coming from an SR-MPLS network while they are inside the ACI fabric based on the incoming MPLS EXP values defined in the MPLS QoS ingress policy. It also marks the CoS and MPLS EXP values of the packets leaving the ACI fabric through an MPLS interface based on IPv4 DSCP values defined in MPLS QoS egress policy.

If no custom ingress policy is defined, the default QoS Level (`Level3`) is assigned to packets inside the fabric. If no custom egress policy is defined, the default EXP value of `0` will be marked on packets leaving the fabric.

Procedure

Step 1 Log in to the Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, select **Application Management** > **Policies**.

Step 3 In the main pane, select **Add Policy** > **Create QoS Policy**.

Step 4 In the **Add QoS Policy** screen, provide the name for the policy.

Step 5 Click **Add Ingress Rule** to add an ingress QoS translation rule.

These rules are applied for traffic that is ingressing the ACI fabric from an MPLS network and are used to map incoming packet's experimental bits (EXP) values to ACI QoS levels, as well as to set differentiated services code point (DSCP) values in the VXLAN header for the packet while it's inside the ACI fabric.

The values are derived at the border leaf using a custom QoS translation policy. The original DSCP values for traffic coming from SR-MPLS without any remarking. If a custom policy is not defined or not matched, default QoS Level (`Level3`) is assigned

a) In the **Match EXP From** and **Match EXP To** fields, specify the EXP range of the ingressing MPLS packet you want to match.
b) From the **Queuing Priority** dropdown, select the ACI QoS Level to map.

   This is the QoS Level you want to assign for the traffic within ACI fabric, which ACI uses to prioritize the traffic within the fabric.. The options range from Level1 to Level6. The default value is `Level3`. If you do not make a selection in this field, the traffic will automatically be assigned a `Level3` priority.

c) From the **Set DSCP** dropdown, select the DSCP value to assign to the packet when it's inside the ACI fabric.

   The DSCP value specified is set in the original traffic received from the external network, so it will be re-exposed only when the traffic is VXLAN decapsulated on the destination ACI leaf node.

   If you set the value to `Unspecified`, the original DSCP value of the packet will be retained.

d) From the **Set CoS** dropdown, select the CoS value to assign to the packet when it's inside the ACI fabric.

   The CoS value specified is set in the original traffic received from the external network, so it will be re-exposed only when the traffic is VXLAN decapsulated on the destination ACI leaf node.

   If you set the value to `Unspecified`, the original CoS value of the packet will be retained, but only if the CoS preservation option is enabled in the fabric. For more information about CoS preservation, see *Cisco APIC and QoS*.

e) Click the checkmark icon to save the rule.
f) Repeat this step for any additional ingress QoS policy rules.

Step 6 Click **Add Egress Rule** to add an egress QoS translation rule.

These rules are applied for the traffic that is leaving the ACI fabric via an MPLS L3Out and are used to map the packet's IPv4 DSCP value to the MPLS packet's EXP value as well as the internal ethernet frame's CoS value.

Classification is done at the non-border leaf switch based on existing policies used for EPG and L3Out traffic. If a custom policy is not defined or not matched, the default EXP value of `0` is marked on all labels. EXP values are marked in both, default and custom policy scenarios, and are done on all MPLS labels in the packet.

Custom MPLS egress policy can override existing EPG, L3out, and Contract QoS policies

a) Using the **Match DSCP From** and **Match DSCP To** dropdowns, specify the DSCP range of the ACI fabric packet you want to match for assigning the egressing MPLS packet's priority.
b) From the **Set MPLS EXP** dropdown, select the EXP value you want to assign to the egressing MPLS packet.
c) From the **Set CoS** dropdown, select the CoS value you want to assign to the egressing MPLS packet.
d) Click the checkmark icon to save the rule.
e) Repeat this step for any additional egress QoS policy rules.

Step 7 Click **Save** to save the QoS policy.

What to do next

After you have created the QoS policy, enable MPLS connectivity and configure MPLS L3Out as described in #unique_179.

Creating SR-MPLS Infra L3Out

This section describes how to configure SR-MPLS L3Out settings for a site that is connected to an SR-MPLS network.

- The SR-MPLS infra L3Out is configured on the border leaf switch, which is used to set up the underlay BGP-LU and overlay MP-BGP EVPN sessions that are needed for the SR-MPLS handoff.
- An SR-MPLS infra L3Out will be scoped to a pod or a remote leaf switch site.
- Border leaf switches or remote leaf switches in one SR-MPLS infra L3Out can connect to one or more provider edge (PE) routers in one or more routing domains.
- A pod or remote leaf switch site can have one or more SR-MPLS infra L3Outs.

Before you begin

You must have:

- Added a site that is connected via SR-MPLS network as described in Adding Cisco ACI Sites.
- If necessary, created SR-MPLS QoS policy as described in Creating SR-MPLS QoS Policy, on page 5.

Procedure

Step 1 Log in to the Cisco Multi-Site Orchestrator GUI.

Step 2 Ensure that SR-MPLS Connectivity is enabled for the site.

a) In the main navigation menu, select **Infrastructure** > **Infra Configuration**.
b) In the **Infra Configuration** view, click **Configure Infra**.
c) In the left pane, under **Sites**, select a specific site.
d) In the right **<*Site*> Settings** pane, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range

The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label.

The Segment Routing Global Block (SRGB) is the range of label values reserved for Segment Routing (SR) in the Label Switching Database (LSD). The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label.

The default range is `16000-23999`.

Step 3 In the main pane, click **+Add SR-MPLS L3Out** within a pod.

Step 4 In the right **Properties** pane, provide a name for the SR-MPLS L3Out.

Step 5 (Optional) From the **QoS Policy** dropdown, select a QoS Policy you created for SR-MPLS traffic.

Select the QoS policy you created in Creating SR-MPLS QoS Policy, on page 5.

Otherwise, if you do not assign a custom QoS policy, the following default values are assigned:

- All incoming MPLS traffic on the border leaf switch is classified into QoS Level 3 (the default QoS level).
- The border leaf switch does the following:
  - Retains the original DSCP values for traffic coming from SR-MPLS without any remarking.
  - Forwards packets to the MPLS network with the original CoS value of the tenant traffic if the CoS preservation is enabled.
  - Forwards packets with the default MPLS EXP value (`0`) to the SR-MPLS network.
- In addition, the border leaf switch does not change the original DSCP values of the tenant traffic coming from the application server while forwarding to the SR network.

Step 6 From the **L3 Domain** dropdown, select the Layer 3 domain.

Step 7 Configure BGP settings.

You must provide BGP connectivity details for the BGP EVPN connection between the site's border leaf (BL) switches and the provider edge (PE) router.

a) Click **+Add BGP Connectivity**.

b) In the **Add BGP Connectivity** window, provide the details.

For the **MPLS BGP-EVPN Peer IPv4 Address** field, provide the loopback IP address of the DC-PE router, which is not necessarily the device connected directly to the border leaf.

For the **Remote AS Number**, enter a number that uniquely identifies the neighbor autonomous system of the DC-PE. the Autonomous System Number can be in 4-byte as plain format from 1 to 4294967295. Keep in mind, ACI supports only `asplain` format and not `asdot` or `asdot+` format AS numbers. For more information on ASN formats, see Explaining 4-Byte Autonomous System (AS) ASPLAIN and ASDOT Notation for Cisco IOS document.

For the **TTL** field, specify a number large enough to account for multiple hops between the border leaf and the DC-PE router, for example `10`. The allowed range `2-255` hops.

(Optional) Choose to enable the additional BGP options based on your deployment.

c) Click **Save** to save BGP settings.

d) Repeat this step to for any additional BGP connections.

Typically, you would be connecting to two DC-PE routers, so provide BGP peer information for both connections.

Step 8 Configure settings for border leaf switches and ports connected to the SR-MPLS network.

You need to provide information about the border leaf switches as well as the interface ports which connect to the SR-MPLS network.

a) Click **+Add Leaf** to add a leaf switch.

b) In the **Add Leaf** window, select the leaf switch from the **Leaf Name** dropdown.

c) Provide a valid segment ID (SID) offset.

When configuring the interface ports later in this section, you will be able to choose whether you want to enable segment routing. The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label. If you plan to enable segment routing, you must specify the segment ID for this border leaf.

- The value must be within the SRGB range you configured earlier.
- The value must be the same for the selected leaf switch across all SR-MPLS L3Outs in the site.
- The same value cannot be used for more than one leaf across all sites.
- If you need to update the value, you must first delete it from all SR-MPLS L3Outs in the leaf and re-deploy the configuration. Then you can update it with the new value, followed by re-deploying the new configuration.

d) Provide the local **Router ID**.

Unique router identifier within the fabric.

e) Provide the **BGP EVPN Loopback** address.

The BGP-EVPN loopback is used for the BGP-EVPN control plane session. Use this field to configure the MP-BGP EVPN session between the EVPN loopbacks of the border leaf switch and the DC-PE to advertise the overlay prefixes. The MP-BGP EVPN sessions are established between the BP-EVPN loopback and the BGP-EVPN remote peer address (configured in the **MPLS BGP-EVPN Peer IPv4 Address** field in the **BGP Connectivity** step before).

While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

f) Provide the **MPLS Transport Loopback** address.

The MPLS transport loopback is used to build the data plane session between the ACI border leaf switch and the DC-PE, where the MPLS transport loopback becomes the next-hop for the prefixes advertised from the border leaf switches to the DC-PE routers.

While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

g) Click **Add Interface** to provide switch interface details.

From the **Interface Type** dropdown, select whether it is a typical interface or a port channel. If you choose to use a port channel interface, it must have been already created on the APIC.

Then provide the interface, its IP address, and MTU size. If you want to use a subinterface, provide the **VLAN ID** for the sub-interface, otherwise leave the VLAN ID field blank.

In the **BGP-Label Unicast Peer IPv4 Address** and **BGP-Label Unicast Remote AS Number**, specify the BGP-LU peer information of the next hop device, which is the device connected directly to the interface. The next hop address must be part of the subnet configured for the interface.

Choose whether you want to enable segment routing (SR) MPLS.

(Optional) Choose to enable the additional BGP options based on your deployment.

Finally, click the checkmark to the right of the **Interface Type** dropdown to save interface port information.

h) Repeat the previous sub-step for all interfaces on the switch that connect to the MPLS network.

i) Click **Save** to save the leaf switch information.

Step 9 Repeat the previous step for all leaf switches connected to MPLS networks.

What to do next

After you have enabled and configured MPLS connectivity, you can create and manage Tenants, route maps, and schemas as described in the *Multi-SiteConfiguration Guide, Release 3.0(x)*.

SR-MPLS Tenant Requirements and Guidelines

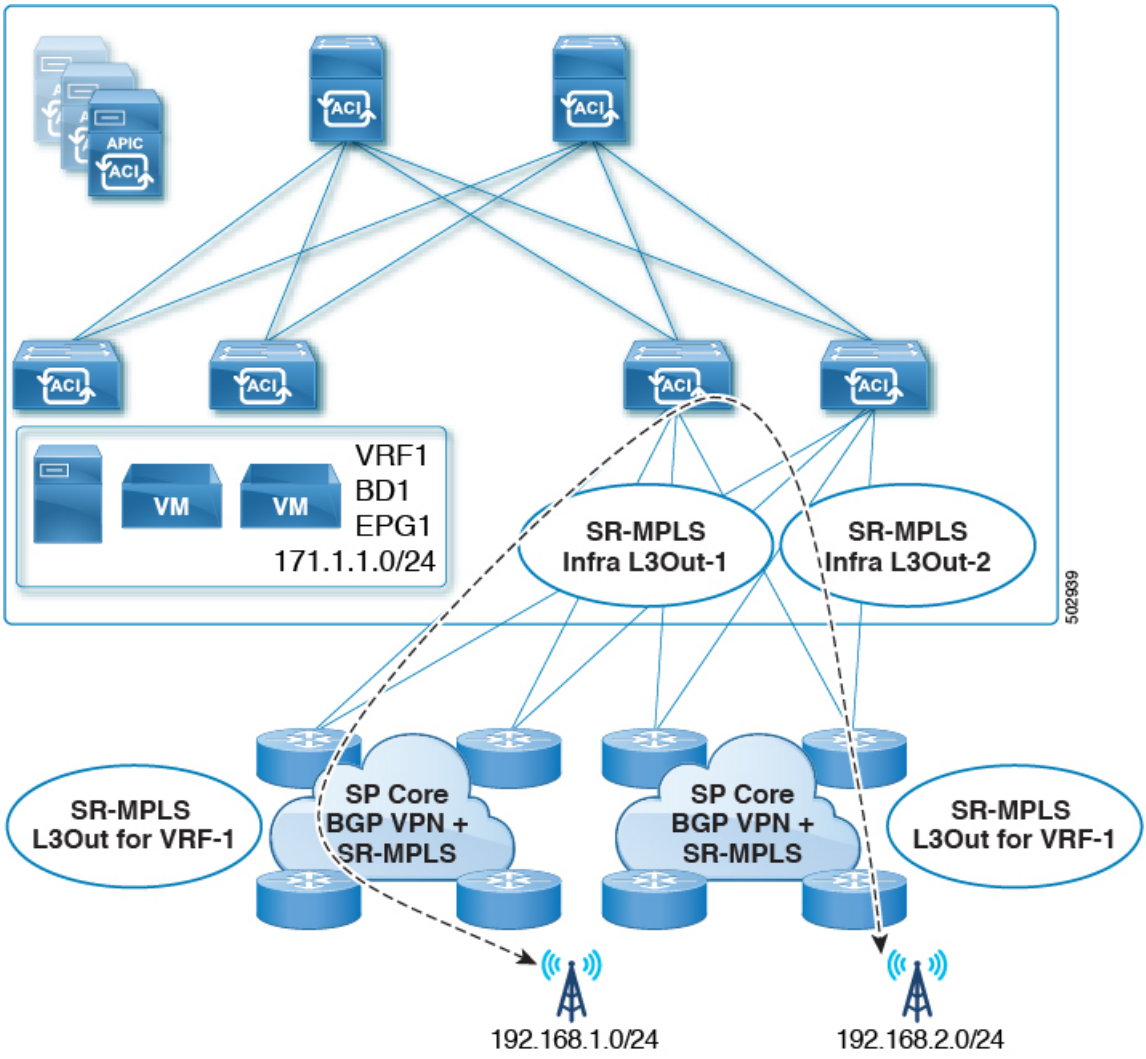
While the Infra MPLS configuration and requirements are described in the Day-0 operations chapter, the following restrictions apply for any user Tenants you will deploy to sites that are connected to SR-MPLS networks.

- You must have created and configured the SR-MPLS Infra L3Outs, including the QoS policies, as described in the Day-0 operations chapter.
- In case when traffic between two EPGs in the fabric needs to go through the SR-MPLS network:
  - Contracts must be assigned between each EPG and the external EPG defined on the local Tenant SR-MPLS L3Out.
  - If both EPGs are part of the same ACI fabric but separated by an SR-MPLS network (for example, in multi-pod or remote leaf cases), the EPGs must belong to different VRFs and not have a contract between them nor route-leaking configured.
  - If EPGs are in different sites, they can be in the same VRF, but there must **not** be a contract configured directly between them.

    Keep in mind, if the EPGs are in different sites, each EPG must be deployed to a single site only. Stretching EPGs between sites is not supported when using SR-MPLS L3Outs.
- When configuring a route map policy for the SR-MPLS L3Out:
  - Each L3Out must have a single export route map. Optionally, it can also have a single import route map.
  - Routing maps associated with any SR-MPLS L3Out must explicitly define all the routes, including bridge domain subnets, which must be advertised out of the SR-MPLS L3Out.
  - If you configure a `0.0.0.0/0` prefixe and choose to not aggregate the routes, it will allow the default route only.

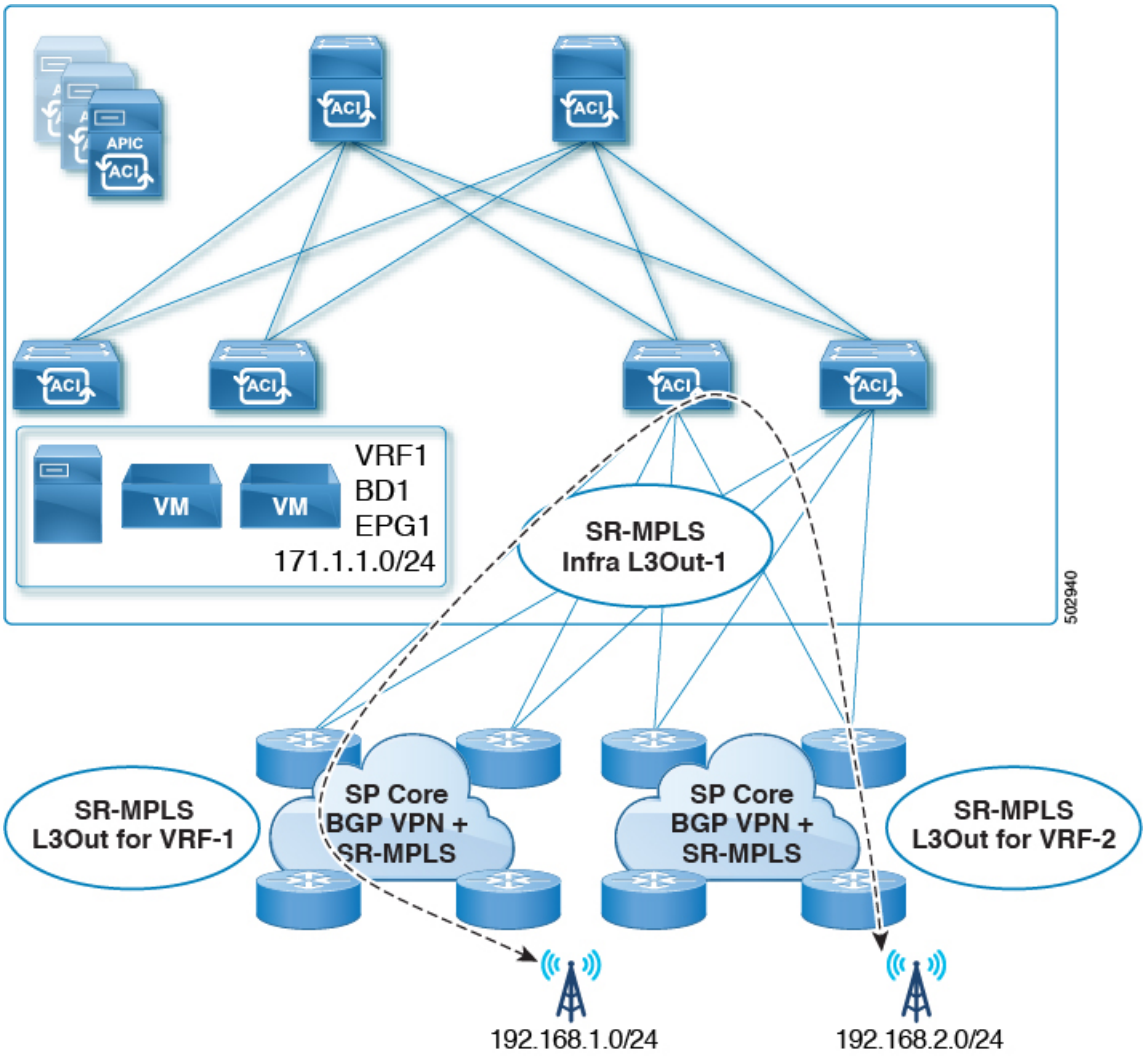
    However, if you choose to aggregate routes `0` through `32` for the `0.0.0.0/0` prefix, it will allow all routes.
  - You can associate any routing policy with any tenant L3Out.
- Transit routing is supported, but with some restrictions:
  - Transit routing between two SR-MPLS networks **using the same VRF** is not supported. The following figure shows an example of this unsupported configuration.

Figure 3: Unsupported Transit Routing Configuration Using Single VRF



- Transit routing between two SR-MPLS networks **using different VRFs** is supported. The following figure shows an example of this supported configuration.

Figure 4: Supported Transit Routing Configuration Using Different VRFs



Creating SR-MPLS Route Map Policy

This section describes how to create a route map policy. Route maps are sets of `if-then` rules that enable you to specify which routes are advertised out of the Tenant SR-MPLS L3Out. Route maps also enable you to specify which routes received from the DC-PE routers will be injected into the BGP VPNv4 ACI control plane.

If you have no sites connected to MPLS networks, you can skip this section.

Procedure

Step 1 Log in to the Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, select **Application Management** > **Policies**.

Step 3 In the main pane, select **Add Policy** > **Create Route Map Policy**.

Step 4 In the **Add Route Map Policy** screen, select a Tenant and provide the name for the policy.

Step 5 Click **Add Entry** under **Route-Map Entry Order** to add a route map entry.

a) Provide the **Context Order** and **Context Action**.

Each context is a rule that defines an action based on one or more matching criteria.

Context order is used to determine the order in which contexts are evaluated. The value must be in the `0-9` range.

Action defines the action to perform (`permit` or `deny`) if a match is found.

b) If you want to match an action based on an IP addres or prefix, click **Add IP Address**.

In the **Prefix** field, provide the IP address prefix. Both IPv4 and IPv6 prefixes are supported, for example `2003:1:1a5:1a5::/64` or `205.205.0.0/16`.

If you want to aggregate IPs in a specific range, check the **Aggregate** checkbox and provide the range. For example, you can specify `0.0.0.0/0` prefix and choose to aggregate routes `0` through `32`.

c) If you want to match an action based on community lists, click **Add Community**.

In the **Community** field, provide the community string. For example, `regular:as2-nn2:200:300`.

Then choose the **Scope**.

d) Click **+Add Action** to specify the action that will be taken should the context match.

You can choose one of the following actions:

- `Set Community`
- `Set Route Tag`
- `Set Weight`
- `Set Next Hop`
- `Set Preference`
- `Set Metric`
- `Set Metric Type`

After you have configured the action, click the checkmark icon to save the action.

e) (Optional) You can repeat the previous substeps to specify multiple match criteria and actions within the same Context entry.

f) Click **Save** to save the Context entry.

Step 6 (Optional) Repeat the previous step if you want to add multiple entries to the same route policy.

Step 7 Click **Save** to save the route map policy.

Enabling Template for SR-MPLS

There is a number of template configuration settings that are unique when deploying them to sites connected via MPLS. Enabling SR-MPLS for a Tenant restricts and filters certain configurations that are not available for MPLS sites while bringing in additional configurations only available for such sites.

Before you can update MPLS-specific settings, you must enable the **SR-MPLS** knob in the template's Tenant properties.

Procedure

Step 1 Log in to the Multi-Site Orchestrator GUI.

Step 2 In the main navigation menu, select **Application Management** > **Schemas**.

Step 3 Create a new or select an existing Schema where you will configure SR-MPLS Tenant.

Step 4 Select the Tenant.

If you created a new Schema, choose a Tenant as you typically would. Otherwise click an existing Template in the left sidebar.

Step 5 In the right sidebar **Template** properties, enable **SR-MPLS** knob.

Creating VRF and SR-MPLS L3Out

This section describes how to create the VRF, tenant SR-MPLS L3Out, and External EPG you will use to configure communication between application EPGs separated by an MPLS network.

Before you begin

You must have:

- Created a template and enabled SR-MPLS for its tenant, as described in Enabling Template for SR-MPLS, on page 14.

Procedure

Step 1 Select the template.

Step 2 Create a VRF.

a) In the main pane, scroll down to the **VRF** area and click the + sign to add a VRF.
b) In the right properties sidebar, provide the name for the VRF.

Step 3 Create an SR-MPLS L3Out.

a) In the main pane, scroll down to the **SR-MPLS L3Out** area and click the + sign to add an L3Out.
b) In the right properties sidebar, provide the name for the L3Out.

c) From the **Virtual Routing & Forwarding** dropdown, select the same VRF you selected for the external EPG in the previous step.

Step 4 Create an external EPG.

a) In the main pane, scroll down to the **External EPG** area and click the + sign to add an external EPG.
b) In the right properties sidebar, provide the name for the external EPG.
c) From the **Virtual Routing & Forwarding** dropdown, select the VRF you created in the previous step.

Configuring Site-Local VRF Settings

You must provide BGP route information for the VRF used by the SR-MPLS L3Out.

Before you begin

You must have:

- Created a template and enabled SR-MPLS for its tenant, as described in Enabling Template for SR-MPLS, on page 14.
- Created a VRF and SR-MPLS L3Out, as described in Creating VRF and SR-MPLS L3Out, on page 14.
- Added the template to an MPLS site.

Procedure

Step 1 Select the schema that contains your template.
Step 2 In the left sidebar of the schema view under **Sites**, select the template to edit its site-local properties.
Step 3 In the main pane, scroll down to **VRF** area and select the VRF.
Step 4 In the right properties sidebar, click **+Add BGP Route Target Address**.
Step 5 Configure the BGP settings.

a) From the **Address Family** dropdown, select whether it is IPv4 or IPv6 address.
b) In the **Route Target** field, provide the route string.

For example, `route-target:ipv4-nn2:1.1.1.1:1901`.

c) From the **Type** dropdown, select whether to import or export the route.
d) Click **Save** to save the route information.

Step 6 (Optional) Repeat the previous step to add any additional BGP route targets.

Configuring Site-Local SR-MPLS L3Out Settings

Similar to how you configure site-local L3Out properties for typical external EPGs, you need to provide SR-MPLS L3Out details for external EPGs deployed to sites connected via MPLS.

Before you begin

You must have:

- Created a template and enabled SR-MPLS for its tenant, as described in Enabling Template for SR-MPLS, on page 14.
- Created a VRF and SR-MPLS L3Out, as described in Creating VRF and SR-MPLS L3Out, on page 14.
- Configured the VRF's site-local properties, as described in Configuring Site-Local VRF Settings, on page 15.
- Added the template to an MPLS site.

Procedure

Step 1 Select the schema that contains your template.

Step 2 In the left sidebar of the schema view under **Sites**, select the template to edit its site-local properties.

Step 3 In the main pane, scroll down to **SR-MPLS L3Out** area and select the MPLS L3Out.

Step 4 In the right properties sidebar, click **+Add SR-MPLS Location**.

Step 5 Configure the SR-MPLS Location settings.

a) From the **SR-MPLS Location** dropdown, select the Infra SR-MPLS L3Out you created when configuring Infra for that site.

b) Under **External EPGs** section, select an external EPG from the dropdown and click the checkmark icon to add it.

You can add multiple external EPGs.

c) Under **Route Map Policy** section, select a route map policy you created in previous section from the dropdown, specify whether you want to import or export the routes, then click the checkmark icon to add it.

You must configure a single export route map policy. Optionally, you can configure an additional import route map policy.

d) Click **Save** to add the location to the MPLS L3Out.

Step 6 (Optional) Repeat the previous step to add any additional SR-MPLS Locations for your SR-MPLS L3Out.

Communication Between EPGs Separated by MPLS Network

Typically, if you wanted to establish communication between two EPGs, you would simply assign the same contract to both EPGs with one EPG being the provider and the other one a consumer.

However, if the two EPGs are separated by an MPLS network, the traffic has to go through each EPG's MPLS L3Out and you establish the contracts between each EPG and its MPLS L3Out instead. This behavior is the same whether the EPGs are deployed to different sites or within the same fabric but separated by an SR-MPLS network, such as in Multi-Pod or Remote Leaf cases.

Before you begin

You must have:

- Added one or more sites connected to MPLS network(s) to the Orchestrator.
- Configured Infra MPLS settings, as described in "Day-0 Operations" chapter.
- Created a schema, added a Tenant, and enabled the Tenant for SR-MPLS, as described in Enabling Template for SR-MPLS, on page 14.

Procedure

Step 1 Log in to the Multi-Site Orchestrator GUI.

Step 2 Create two application EPGs as you typically would.

For example, `epg1` and `epg2`.

Step 3 Create two separate external EPGs

These EPGs can be part of the same template or different templates depending on the specific deployment scenario.

For example, `mpls-extepg-1` and `mpls-extepg-2`

Step 4 Configure two separate Tenant SR-MPLS L3Outs.

For example, `mpls-l3out-1` and `mpls-l3out-2`

For each Tenant SR-MPLS, configure the VRF, route map policies, and external EPGs as described in Configuring Site-Local VRF Settings, on page 15 and Configuring Site-Local SR-MPLS L3Out Settings, on page 15.

Step 5 Create a contract you will use to allow traffic between the two application EPGs you created in Step 2.

You will need to create and define a filter for the contract just as you typically would.

Step 6 Assign the contracts to the appropriate EPGs.

In order to allow traffic between the two application EPGs you created, you will actually need to assign the contract twice: once between `epg1` and its `mpls-l3out-1` and then again between `epg2` and its `mpls-l3out-2`.

As an example, if you want `epg1` to provide a service to `epg2`, you would:

a) Assign the contract to `epg1` with type `provider`.
b) Assign the contract to `mpls-l3out-1` with type `consumer`.
c) Assign the contract to `epg2` with type `consumer`.
d) Assign the contract to `mpls-l3out-2` with type `provider`.

Deploying Configuration

You can deploy the configuration Template to an MPLS site as you typically would, with one exception: because you cannot stretch objects and policies between MPLS site and another site, you can only select a single site when deploying the template.

Procedure

Step 1 Add the site to which you want to deploy the template.

a) In the left sidebar of the **Schema** view under **Sites**, click the **+** icon.
b) In the **Add Sites** window, select the site where you want to deploy the Template.

You can only select a single site if your template is MPLS-enabled.

c) From the **Assign to Template** dropdown, select one or more Template you have created in this Schema.
d) Click **Save** to add the site.

Step 2 Deploy the configuration

a) In the main pain of the **Schemas** view, click **Deploy to Sites**.
b) In the **Deploy to Sites** window, verify the changes that will be pushed to the site and click **Deploy**.