Node and Interface for L3Out

Modifying Interfaces for L3Out

Modifying Interfaces for L3Out Using the GUI

This procedure modifies an L3Out interface.

Note The steps for filling out the fields are not necessarily listed in the same order that you see them in the GUI.

Before you begin

- The Cisco ACI fabric is installed, the Cisco APICs are online, and the Cisco APIC cluster is formed and healthy.
- A Cisco APIC fabric administrator account is available that enables creating the necessary fabric infrastructure configurations.
- The target leaf switches are registered in the Cisco ACI fabric and available.
- Port channels are configured when port channels are used for L3Out interfaces.

Procedure

Step 1 On the menu bar, choose **Tenants** > **All Tenants**.

Step 2 In the Work pane, double click the tenant's name.

Step 3 In the Navigation pane, expand **tenant_name** > **Networking** > **L3Outs** > ***L3Out*** > **Logical Node Profiles** > ***node_profile*** > **Logical Interface Profiles** and choose the profile that you want to modify.

Step 4 Choose an interface type tab: **Routed Sub-Interfaces**, **Routed Interfaces**, **SVI**, or **Floating SVI**.

Step 5 Double click an existing interface to modify it, or click the **Create** (**+**) button to add a new interface to the logical interface profile.

Step 6 For interface types other than floating SVI, perform the following substeps:

a) To add a new interface in the **Path Type** field, choose the appropriate path type.

For the routed sub-interface and routed interface interface types, choose **Port** or **Direct Port Channel**. For the SVI interface type, choose **Port**, **Direct Port Channel**, or **Virtual Port Channel**.

b) In the **Node** drop-down list, choose a node.

Note
This is applicable only for the non-port channel path types. If you selected **Path Type** as **Port**, then perform this step. Otherwise, proceed to the next step.

c) In the **Path** drop-down list, choose the interface ID or the port channel name.

An example of an interface ID is eth 1/1. The port channel name is the interface policy group name for each direct or virtual port channel.

Step 7 For the floating SVI interface type, in the **Anchor Node** drop-down list, choose a node.

Step 8 (Optional) In the **Description** field, enter a description of the L3Out interface.

Step 9 For the routed sub-interfaces, SVI, and floating SVI interface types, in the **Encap** drop-down list, choose **VLAN** and enter an integer value for this entry.

Step 10 For the SVI and floating SVI interface types, perform the following substeps:

a) For the **Encap Scope** buttons, choose the scope of the encapsulation used for the Layer 3 Outside profile.

- **VRF**: Use the same transit VLAN in all Layer 3 Outsides in the same VRF instance for a given VLAN encapsulation. This is a global value.
- **Local**: Use a unique transit VLAN per Layer 3 Outside.

b) For the **Auto State** buttons, choose whether to enable or disable this feature.

- **disabled**: The SVI or floating SVI remains active even if no interfaces are operational in the corresponding VLANs.
- **enabled**: When a VLAN interface has multiple ports in the VLAN, the SVI or floating SVI goes to the down state when all the ports in the VLAN go down.

c) For the **Mode** buttons, choose the VLAN tagging mode.

Step 11 In the **IPv4 Primary / IPv6 Preferred Address** field, enter the primary IP addresses of the path attached to the Layer 3 outside profile.

Step 12 In the **IPv4 Secondary / IPv6 Additional Addresses** table, click the **+** to enter the secondary IP addresses of the path attached to the Layer 3 outside profile.

Step 13 (Optional) In the **Link-local Address** field, enter an IPv6 link-local address. This is the override of the system-generated IPv6 link-local address.

Step 14 In the **MAC Address** field, enter the MAC address of the path attached to the Layer 3 outside profile.

Step 15 In the **MTU (bytes)** field, set the maximum transmit unit of the external network. The range is 576 to 9216. To inherit the value, enter *inherit* in the field.

Step 16 In the **Target DSCP** drop-down list, choose the target differentiated services code point (DSCP) of the path attached to the Layer 3 outside profile.

Step 17 Click **Submit**.

Create OSPF Interface Profile

The OSPF interface profile enables OSPF on the interface. Optionally, the OSPF interface profile can have a relation to an OSPF interface policy for more granular control over interface proprieties.

Faults

Faults are raised in the following scenarios, which will bring down the OSPF session:

- No pre-shared key provided under Key (key-string) provided under "key" in the KeyChain policy
- No Key (key-string) configured in the KeyChain policy
- If you specify unsupported encryption algorithm such as 3DES and AES. These algorithms are supported for Authentication.

No fault will be raised if the OSPF session goes down because the Send/ Accept lifetime of the key was expired, with no active key. The KeyChain state under the OSPF interface will be in “not-ready” state.

Before you begin

- The Cisco ACI fabric is installed, the Cisco APICs are online, and the Cisco APIC cluster is formed and healthy.
- A Cisco APIC fabric administrator account is available that enables creating the necessary fabric infrastructure configurations.
- The target leaf switches are registered in the Cisco ACI fabric and available.
- Port channels are configured when port channels are used for L3Out interfaces.

Procedure

Step 1 On the menu bar, choose **Tenants > All Tenants**.

Step 2 In the Work pane, double click the tenant's name.

Step 3 In the Navigation pane, expand **tenant_name > Networking > L3Outs > L3Out > Logical Node Profiles > node_profile > Logical Interface Profiles > OSPF Interface Profile**.

Step 4 In the **Name** field, enter a name for the OSPF interface. This name can be between 1 and 64 alphanumeric characters.

Note
You cannot change this name after the object has been saved.

Step 5 [Optional] In the **Description** field, enter a description for the OSPF interface profile. The description can be 0 to 128 alphanumeric characters.

Step 6 Enter a value for the target interface policy name. This name can be between 1 and 64 alphanumeric characters. You cannot change this name after the object has been saved.

Step 7 To configure the OSPF interface profile by using the MD5 or the simple authentication, complete the following steps:

a) In the **OSPFv2 Authentication Key** field, enter the authentication key. The authentication key is a password (up to 8 characters) that can be assigned on an interface basis. The authentication key must match for each router on the interface

Note
To use authentication, the OSPF authentication type for this interface's area should be set to Simple (the default is None).

b) In the **Confirm OSPFv2 Authentication Key** field, renter the authentication key.

c) In the **OSPFv2 Authentication Key ID** field, enter the authentication key identifier.

d) In the OSPFv2 Authentication Type field, select the appropriate option.

The OSPF authentication type. Authentication enables the flexibility to authenticate OSPF neighbors. You can enable authentication in OSPF to exchange routing update information in a secure manner.

Note
When you configure authentication, you must configure an entire area with the same type of authentication.

The authentication types are:

- **None**—No authentication is used.
- **Simple**—You need to specify the authentication key, **OSPFv2 Authentication Key** that you specified earlier. The authentication key is a password (up to 8 characters) that can be assigned on an interface basis. The authentication key must match for each router on the interface
- **Md5**—The password does not pass over the network. MD5 is a message-digest algorithm specified in RFC 1321. MD5 is considered the most secure OSPF authentication mode. When you configure authentication, you must configure an entire area with the same type of authentication.

The default is **None**.

Step 8 To configure the OSPF interface profile by using the KeyChain authentication, complete the following steps:

a) In the **OSPFv2 KeyChain Policy** field, select OSPFv2 KeyChain policy.

The OSPFv2 KeyChain policy supports HMAC-SHA authentication along with Simple and MD5 authentication. When you select this option, you can have multiple keys under the same key chain.

For enhanced security, you can use the rotating keys by specifying a life time for each key. When the lifetime expires for a key it automatically rotates to next key. If you do not specify any algorithm, OSPF will use MD5, which is the default cryptographic authentication algorithm

Note
The new key is the preferred key and will take precedence against the existing keys.

Note
You can configure the authentication by specifying the legacy way, which is the *OSPFv2 authentication type - MD5 authentication /Simple authentication* or by specifying the *OSPFv2 keychain policy*.

Configuring the Keychain policy will override the selected Authentication Type.

Step 9 (applicable only for OSPFv3) **OSPFv3 IPsec Policy**: to associate an OSPFv3 IPsec policy to an L3Out interface, select an IPsec policy from the drop-down list. For creating an OSPFv3 IPsec policy, see the Create an OSPF IPsec Policy procedure

What to do next

To specify the rotating keys for the OSPFv2 KeyChain, refer to the Create Key Policy, on page 5.

Create Key Policy

Before you begin

Ensure that you have create the OSPFv2 interface profile. Refer to Create OSPF Interface Profile, on page 3 for more information.

Procedure

Step 1 On the menu bar, click **Tenants> All Tenants**.

Step 2 In the Work pane, double click the tenant's name.

Step 3 In the Navigation pane, navigate to **Policies > Protocol > KeyChains**.

Step 4 Right-click **KeyChains**, select **Create Key Policy**, and perform the following steps:

a) Type a name for the policy and optionally add a description.
b) In the **Key ID** field, enter a valid key ID.
c) In the **Pre-Shared Key** field, enter the pre-shared key.
d) In the **Cryptographic Algorithm** field, enter the algorithm.
e) In the **Start Time** field, specify the start time in **YYYY-MM--DD- HH-MM-SS** format.
f) In the **End Time** field, specify the end time in **YYYY-MM--DD- HH-MM-SS** format.
g) In the **Key accept lifetime start time** field, specify the start time in the **YYYY-MM--DD- HH-MM-SS** format.

This is a rotating key. you will be specifying a life time for each key. When the lifetime expires for a key it automatically rotates to next key. If you do not specify any algorithm, OSPF will use MD5, which is the default cryptographic authentication algorithm.

This field is not applicable for OSPFv3 IPSec policy.

Note
The new key is the preferred key and will take precedence against the existing keys.

h) In the **Key accept lifetime end time** field, specify the end time in the **YYYY-MM--DD- HH-MM-SS** format.

This field is not applicable for OSPFv3 IPSec policy.

Step 5 Click **Submit**.

Create an OSPF IPsec Policy

Beginning with Cisco APIC release 6.1(2), encryption and authentication for OSPFv3 sessions are supported. Use this procedure

Note OSPFv3 is not supported on infra tenant; the OSPF IPSec policy support is for user tenant only.

Before you begin

Create a keychain policy.

Procedure

Step 1 On the menu bar, click **Tenants**> *tenant name*.

Step 2 On the left Navigation pane, navigate to **Policies** > **Protocol** > **OSPF** > **OSPF IPSec**.

Step 3 In the **Create IPSec Authentication Policy** window, enter these details:

- Name: for the IPSec policy.
- Description: description for the IPSec policy.
- IP Security Protocol: select either Authentication Header (AH) or Encapsulating Security Payload (ESP).

  If you select Authentication Header, only authentication is supported. If you select ESP, the available options are: authentication, encryption or both (authentication and encryption).

  Supported keychain algorithms:

  - Authentication Header: MD5 (default), HMAC-SHA1.
  - Encapsulating Security Protocol: for authentication: HMAC-SHA1; for encryption: 3DES (default), AES.

  Note
  If you do not select any algorithm or choose an unsupported algorithm, the default is automatically selected.

- Security Parameter Index: unique value for creating the IPSec protocol. Select a value from the drop-down list. The supported range is from 256 to 4294967295.
- OSPFv3 Authentication Keychain: select a keychain value from the drop-down list. If you have selected the AH option for the IP Security Protocol field, this field is mandatory. If you leave the field blank, a fault is generated.

  To check for faults, navigate to the OSPF Interface Profile screen and click the **Faults** tab.
- OSPFv3 Encryption Keychain: select a keychain value from the drop-down list. This field is not applicable if you have selected the AH option for the **IP Security Protocol** field. If you have selected the ESP option for the IP SecurityProtocol field, it is mandatory to enter a value for the **Authentication Keychain** field or the **Encryption Keychain** field.

Step 4 Click **Submit**.

You can use the **show ipv6 ospfv3 interface** *interface-id* command on the switch over a SSH or console session to check the created IPSec policy.

Step 5 To associate the created OSPFv3 IPSec policy to an L3Out interface, see Step-9 of the Create OSPF Interface Profile procedure.

Create OSPF Timers Policy

OSPF timers control the behavior of protocol messages and shortest path first (SPF) calculations. Use this procedure to configure this policy that will be used in the VRF.

Before you begin

Ensure that you have create the OSPFv2 interface profile. Refer to Create OSPF Interface Profile, on page 3 for more information.

Procedure

Step 1 On the menu bar, choose **Tenants > Policies>OSPF>OSPF Timers**.

Step 2 In the Work pane, double click the tenant's name.

Step 3 In the **Name** field, context-level OSPF policy name. This name can be between 1 and 64 alphanumeric characters.

Note
You cannot change this name after the object has been saved.

Step 4 [Optional] In the **Description** field, enter a description for the OSPF interface profile. The description can be 0 to 128 alphanumeric characters.

Step 5 In the **Bandwidth Reference** field, enter the OSPF bandwidth reference. This is used to calculate the default metrics for an interface. The range is 1 to 40000. The default value is 40000.

Step 6 In the **Admin Distance Preference** field, enter the preferred administrative distance. The administrative distance is the feature that the routers use order to select the best path when there are two or more different routes to the same destination from two different routing protocols. Administrative distance defines the reliability of a routing protocol. Each routing protocol is prioritized in order of most to least reliable (believable) with the help of an administrative distance value. The range is 1 to 255. The default value is 110.

Step 7 In the **Maximum ECMP** field, enter the maximum ECMP for the OSPF protocol. The maximum range is 1 to 16. The default value is 8.

Step 8 In the **Control Knobs** field, enter any of the following OSPF policy control types:

- **Enable name lookup fo router IDs**
- **Prefix suppression**—The prefix is not advertised.

Step 9 Put a check in the **Graceful Restart Controls** checkbox for a graceful restart, or nonstop forwarding (NSF), which allows OSPF to remain in the data forwarding path through a process restart.

Step 10 In the **Initial Spf Schedule Delay Interval (Ms)** field, enter the initial delay interval for the SPF schedule. The interval indicates the amount of time to wait until the first SPF calculation occurs. Each interval after the initial calculation is

twice as long as the previous one until the wait interval reaches the maximum wait time specified. The range is 1 to 600000. The default value is 200.

Step 11 In the **Minimum Hold Time Between Spf Calculations (Ms)** field, enter the minimum hold time between SPF calculations. The interval indicates the minimum amount of time to wait until the SPF calculation after the initial interval occurs. Each interval after the initial calculation is twice as long as the previous one until the wait interval reaches the maximum wait time specified. The range is 1 to 600000. The default value is 1000.

Step 12 In the **Maximum Wait Time Between Spf Calculations (Ms)** field, enter the maximum interval between SPF calculations. Each interval after the initial calculation is twice as long as the previous one until the wait interval reaches the maximum wait time specified. The range is 1 to 600000. The default value is 5000.

Step 13 In the **LSA Group Pacing Interval (Secs)** field, enter the interval in which LSAs are grouped and refreshed, checksummed, or aged. The duration of the LSA group pacing is inversely proportional to the number of LSAs that the router is handling. For example, if you have about 10,000 LSAs, you should decrease the pacing interval. If you have a very small database (40 to 100 LSAs), you should increase the pacing interval to 10 to 20 minutes. The range is from 1 to 1800 seconds. The default value is 10 seconds.

Step 14 In the **LSA Generation Throttle Start Wait Interval (Ms)** field, enter the generation throttle start-wait interval between LSAs. This is the minimum interval for accepting the same LSA. If an instance of the same LSA arrives sooner than the interval that is set, the LSA is dropped. The range is 0 to 5000. The default value is 0.

Step 15 In the **LSA Generation Throttle Hold Interval (Ms)** field, enter the incremental time (in milliseconds) used to calculate the subsequent rate limiting times for LSA generation. The throttle interval range is from 50 to 30000. The default value is 5000.

Step 16 In the **LSA Generation Throttle Maximum Interval (Ms)** field, enter the maximum time (in milliseconds) that is used to calculate the subsequent rate limiting times for LSA generation. The range is 50 to 30000. The default value is 5000.

Step 17 In the **Minimum Interval Between Arrival of a LSA (Ms)** field, enter the minimum interval in which the software accepts the same link-state advertisement (LSA) from Open Shortest Path First (OSPF) neighbors. The same LSA is an LSA instance that contains the same LSA ID number, LSA type, and advertising router ID. If an instance of the same LSA arrives sooner than the interval that is set, the software drops the LSA. The range is 10 to 600000. The default value is 1000.

Step 18 In the **Maximum Number of Not Self-generated LSAs** field, enter the maximum number of Not Self-generated LSAs. The default value is 20000.

Step 19 In the **LSA Maximum Sleep Ignore Count Interval (ignore-time)** specify the time in minutes the OSPF process stays in the ignore state after LSA limit is exceeded.

Step 20 In the **LSA Maximum Sleep Ignore Count (ignore-count)** field, set the maximum number of times the OSPF process can enter the ignore state before manual recovery is required.

Step 21 In the **LSA Sleep Count Reset Interval (reset-time)** field, specify the time in minutes for the OSPF process to operate normally to reset the ignore state counter to zero.

Step 22 In the **LSA Threshold (percentage)** field, enter the number of LSAs expressed in a percentage of the total threshold maximum. The default value is 75 percent.

Step 23 In the **LSA Maximum Action**field, select either the **Log** or **Reject** option.

Step 24 Click **Submit**.

What to do next

To deploy the newly created OSPF Timer Policy, navigate to **Tenants> Networking> VRFs> Policy> OSPF Timers** and associate it to the VRF.

OSPF Link-State Database Overload Protection

The OSPF Link-State Database OverloadProtection feature allows you to limit the number of nonself-generated link-state advertisements (LSAs) for a given Open Shortest Path First (OSPF) process. Excessive LSAs generated by other routers in the OSPF domain can substantially drain the CPU and memory resources of the router.

Prerequisites forOSPF Link-State Database Overload Protection

It is presumed that you have OSPF running on your network.

Benefits of Using OSPF Link-State Database Overload Protection

The OSPF Link-State Database Overload Protection feature provides a mechanism at the OSPF level to limit the number of nonself-generated LSAs for a given OSPF process. When other routers in the network have been misconfigured, they may generate a high volume of LSAs, for instance, to redistribute large numbers of prefixes. This protection mechanism prevents routers from receiving a large number of LSAs and therefore experiencing CPU and memory shortages.

How OSPF Link-State Database Overload Protection Works

When the OSPF Link-State Database Overload Protection feature is enabled, the router keeps a count of the number of received (non-self-generated) LSAs. When the configured threshold number of LSAs is reached, a fault with the description **Number of non-self-originated LSAs <> has exceeded allowed limit <>** is raised in the APIC. If the configured maximum number of LSAs is exceeded, the router will send a notification. If the count of received LSAs remains higher than the configured maximum after one minute, the OSPF process takes down all adjacencies and clears the OSPF database. In this ignore state, all OSPF packets received on any interface belonging to this OSPF process are ignored, and no OSPF packets are generated on any of these interfaces. The OSPF process remains in the ignore state for the duration configured in the **LSA Maximum Sleep Ignore Count Interval** (`ignore-time`) field. Each time the OSPF process enters the ignore state, a counter is incremented. If this counter exceeds the value configured in the **LSA Maximum Sleep Ignore Count** (`ignore-count`) field, the OSPF process stays permanently in the ignore state, and manual intervention is required to bring the OSPF process out of this state. The ignore state counter is reset to 0 when the OSPF process operates normally for the amount of time specified by the **LSA Sleep Count Reset Interval**(`reset-time`) field.

If the LSA Maximum Action is set to **Log** (warning-only), the OSPF process logs the fault when the LSA maximum limit is exceeded, but no adjacencies are affected.

Note To recover from a permanent-ignore state, one of the following manual interventions is required:

1. Disable and re-enable the OSPF process in L3Out (if it is in the user tenant).
2. Restart the process (For example: `kill -9 $(pidof <ospf>)`.
3. Reload the device.

OSPF max-metric

OSPF max-metric feature controls the flow of routing information within a network. This feature lets a router advertise its locally generated link-state advertisements (LSAs) with the maximum metric. This makes the router less preferable as a transit path for data traffic. This approach is especially useful during switch reloads, as it prevents the device from being selected for transit traffic until it is operational.

Guidelines of OSPF max-metric policy

Follow these guidelines when configuring the OSPF max-metric policy:

- Create the OSPF max-metric policy under either the common or user tenant.
- The OSPF max-metric policy allows you to select controls such as External LSAs, Stub Links, or Summary LSAs, and specify a max-metric value from 1 to 16,777,215. This max-metric value only applies to External and Summary LSAs. Router LSAs and router LSAs for stub networks are always advertised with a fixed max-metric value of 65,535.
- If the OSPF max-metric policy is enabled without selecting the External LSAs, Stub Links, or Summary LSAs controls, only Router LSAs for non-stub networks are advertised with the max-metric value.
- When the On Startup control is enabled, the OSPF max-metric policy is advertised only during the configured Startup Interval Time after the switch boots. If On Startup is not enabled, the configured max-metric will be advertised from all OSPF-enabled border leafs in the VRF.
- Associate the OSPF max-metric policy with the VRF where max-metric must be enabled.
- OSPF max-metric policy is not supported under the infra tenant.
- Deploy the OSPF max-metric policy only on border leaf switches.
- This feature is not supported for the overlay-1 VRF under the infra tenant because overload mode is implicitly enabled by default on the PE.

Limitations of OSPF max-metric policy

Be aware of these limitations when configuring the OSPF max-metric policy:

- The **wait-for-bgp** option in the OSPF max-metric policy is not supported due to a limitation in BGP functionality on ACI.
- The OSPF max-metric policy is not supported for OSPFv3 in Release 6.1.4

Create OSPF max-metric policy using the GUI

Use this task when you want to limit the OSPF traffic through the router for a short time.

Follow these steps to create an OSPF max-metric policy using the GUI.

Procedure

Step 1 On the menu bar, choose **Tenants**.

Step 2 In the Navigation pane, expand *Tenant_name* > **Policies** > **Protocol** > **OSPF** , right-click **OSPF Max-Metric** and choose **Create OSPF Max-Metric Policy**.

The **Create Max-Metric Link State Advertisement per Domain/VRF** dialog box appears.

Step 3 Enter the name of the policy.

Step 4 Configure the required maximum metric controls.

a) Check the **External LSAs** check box to configure external LSAs to maximum metric.
b) Check the **On Startup** check box to configure the router to advertise a maximum metric at startup.
c) Check the **Stub Links** check box to configure stub links to maximum metric.
d) Check the **Summary LSAs** check box to configure summary LSAs to maximum metric.

If you do not select External LSAs, Stub Links, or Summary LSAs, the system advertises the max-metric only for router LSAs. If any of these controls are enabled, the corresponding LSAs and the router LSAs will be advertised with the max-metric.

Step 5 Enter a value in the **Maximum metric value for external LSAs** field to specify the maximum metric values for external LSAs.

The default value is 65,535. The range is 1 to 16,777,215. The router LSAs are always advertised with a metric value of 65,535 when max-metric is enabled.

Step 6 Enter a value in the **Maximum metric value for summary LSAs** field to specify the maximum metric values for summary LSAs.

The default value is 65,535. The range is 1 to 16,777,215.

Step 7 Enter a value in the **Startup Interval Time (in secs)** field to specify the time interval to advertise the maximum metric.

The default value is 600 seconds. The range is 5 to 86,400 seconds.

Step 8 Click **Submit**.

The created OSPF max-metric policy appears under **Protocol** > **OSPF** > **OSPF Max-Metric**.

You can also create the OSPF max-metric policy during VRF creation. In step 1 of the **Create VRF** wizard, check the **Configure OSPF Policies** check box. In step 2 of the wizard, choose **Create OSPF Max-Metric Policy** from the **OSPF Max-Metric** drop-down list.

Associate OSPF max-metric policy to VRF using the GUI

Follow these steps to associate an OSPF max-metric policy to a VRF using the GUI.

Procedure

Step 1 On the menu bar, choose **Tenants**.

Step 2 In the Navigation pane, expand *Tenant_name* > **Networking** > **VRFs** and choose a VRF.

You can view the details of the selected VRF in the right pane.

Step 3 Click the **Policy** tab.

Step 4 From the **OSPF Max-Metric** drop-down list, choose an OSPF max-metric policy.

Step 5 Click **Submit** to associate an OSPF max-metric policy to a VRF.

When the associated OSPF max-metric policy is deleted under **Policies** > **Protocol** > **OSPF**, a fault is raised.

Verify OSPF max-metric policy configurations using the CLI

Follow these steps to verify the OSPF max-metric policy configurations using the CLI.

Procedure

Run the **show ip ospf vrf** command to display OSPF information within a specific VRF instance.

Example:

```
nextaci1-leaf1# show ip ospf vrf ospf_max-metric:ospf_vrf_backbone
 Routing Process default with ID 101.101.101.101 VRF ospf_max-metric:ospf_vrf_ba
ckbone
 Routing Process Instance Number 1
 Stateful High Availability enabled
 Graceful-restart helper mode is enabled
 Supports only single TOS(TOS0) routes
 Supports opaque LSA
 This router is an autonomous system boundary
 Maximum number of non self-generated LSA allowed 20000
    (feature configured but inactive)
    Current number of non self-generated LSA 2
    Threshold for warning message 75%
    Ignore-time 5 minutes, reset-time 10 minutes
    Ignore-count allowed 5, current ignore-count 0
 Table-map using route-map exp-ctx-2818053-deny-external-tag
 Redistributing External Routes from
   static route-map exp-ctx-st-2818053
   direct route-map exp-ctx-st-2818053
   bgp-100 route-map exp-ctx-proto-2818053
   eigrp-default route-map exp-ctx-proto-2818053
   coop route-map exp-ctx-st-2818053
 Administrative distance 110
 Reference Bandwidth is 40000 Mbps
 SPF throttling delay time of 200.000 msecs,
   SPF throttling hold time of 1000.000 msecs,
   SPF throttling maximum wait time of 5000.000 msecs
 LSA throttling start time of 0.000 msecs,
```

```
  LSA throttling hold interval of 5000.000 msecs,
  LSA throttling maximum wait time of 5000.000 msecs
Minimum LSA arrival 1000.000 msec
LSA group pacing timer 10 secs
Maximum paths to destination 8
Originating router LSA with maximum metric
  Condition: Always
Number of external LSAs 1, checksum sum 0x5c0e
Number of opaque AS LSAs 0, checksum sum 0
```

Customizing SVI for L3Out

SVI External Encapsulation Scope

About SVI External Encapsulation Scope

In the context of a Layer 3 Out configuration, a switch virtual interfaces (SVI), is configured to provide connectivity between the ACI leaf switch and a router.

By default, when a single Layer 3 Out is configured with SVI interfaces, the VLAN encapsulation spans multiple nodes within the fabric. This happens because the ACI fabric configures the same bridge domain (VXLAN VNI) across all the nodes in the fabric where the Layer 3 Out SVI is deployed as long as all SVI interfaces use the same external encapsulation (SVI) as shown in the figure.

However, when different Layer 3 Outs are deployed, the ACI fabric uses different bridge domains even if they use the same external encapsulation (SVI) as shown in the figure:

Figure 1: Local Scope Encapsulation and One Layer 3 Out

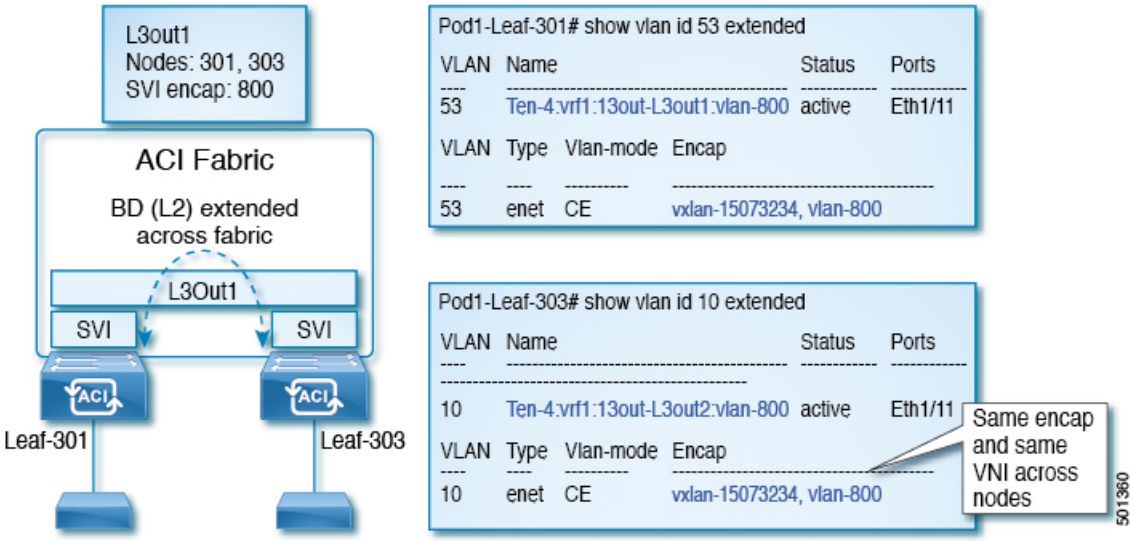

Figure 2: Local Scope Encapsulation and Two Layer 3 Outs

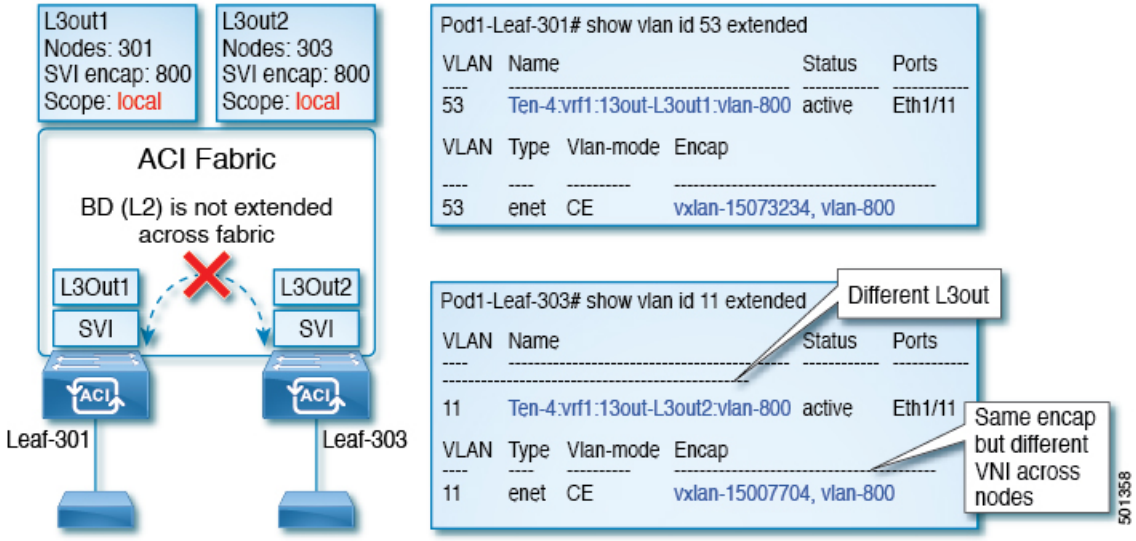

Starting with Cisco APIC release 2.3, it is now possible to choose the behavior when deploying two (or more) Layer 3 Outs using the same external encapsulation (SVI).

The encapsulation scope can now be configured as Local or VRF:

- Local scope (default): The example behavior is displayed in the figure titled *Local Scope Encapsulation and Two Layer 3 Outs*.
- VRF scope: The ACI fabric configures the same bridge domain (VXLAN VNI) across all the nodes and Layer 3 Out where the same external encapsulation (SVI) is deployed. See the example in the figure titled *VRF Scope Encapsulation and Two Layer 3 Outs*.

Figure 3: VRF Scope Encapsulation and Two Layer 3 Outs

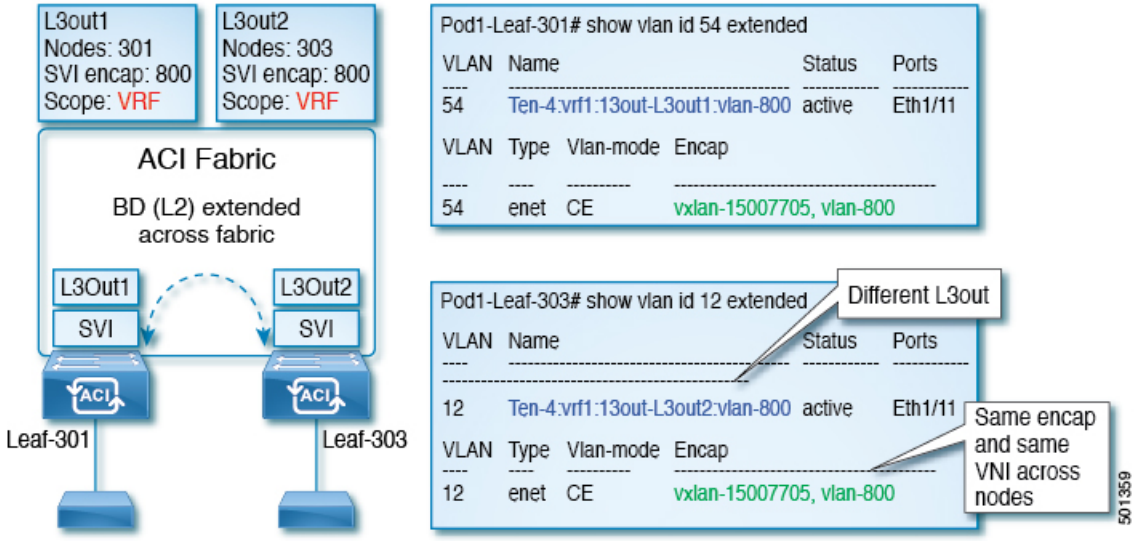

Encapsulation Scope Syntax

The options for configuring the scope of the encapsulation used for the Layer 3 Out profile are as follows:

- **Ctx**—The same external SVI in all Layer 3 Outs in the same VRF for a given VLAN encapsulation. This is a global value.
- **Local** —A unique external SVI per Layer 3 Out. This is the default value.

The mapping among the CLI, API, and GUI syntax is as follows:

Table 1: Encapsulation Scope Syntax

CLI	API	GUI
l3out	local	Local
vrf	ctx	VRF

Note The CLI commands to configure encapsulation scope are only supported when the VRF is configured through a named Layer 3 Out configuration.

Guidelines for SVI External Encapsulation Scope

To use SVI external encapsulation scope, follow these guidelines:

- If deploying the Layer 3 Outs on the same node, the OSPF areas in both the Layer 3 Outs must be different.
- If deploying the Layer 3 Outs on the same node, the BGP peer configured on both the Layer 3 Outs must be different.

Configuring SVI External Encapsulation Scope Using the GUI

Before you begin

- The tenant and VRF configured.
- An L3Out is configured and a logical node profile under the L3Out is configured.

Procedure

Step 1 On the menu bar, click > **Tenants** > ***Tenant_name***.

Step 2 In the **Navigation** pane, click **Networking** > **L3Outs** > ***L3Out_name*** > **Logical Node Profiles** > ***LogicalNodeProfile_name*** > **Logical Interface Profiles**.

Step 3 In the **Navigation** pane, right-click **Logical Interface Profiles**, and click **Create Interface Profile**.

Step 4 In the **Create Interface Profile** dialog box, perform the following actions:

a) In the **Step 1 Identity** screen, in the **Name** field, enter a name for the interface profile.

b) In the remaining fields, choose the desired options, and click **Next**.
c) In the **Step 2 Protocol Profiles** screen, choose the desired protocol profile details, and click **Next**.
d) In the **Step 3 Interfaces** screen, click the **SVI** tab, and click the **+** icon to open the **Select SVI** dialog box.
e) In the **Specify Interface** area, choose the desired values for the various fields.
f) In the **Encap Scope** field, choose the desired encapsulation scope value. Click **OK**.

The default value is **Local**.

The SVI External encapsulation scope is configured in the specified interface.

Support for Multiple Encapsulation for L3Outs With SVI

When an L3Out is configured with SVI interfaces on different leaf switches using the same encapsulation VLAN, the SVI VLAN will be mapped to the same VXLAN network identifier (VNID). This forms a single bridge domain (external bridge domain) and broadcast domain across the fabric. An SVI interface configured with a different VLAN will form a separate external bridge domain as illustrated in the diagram below. Prior to release 5.2(3) it was not possible to create a single external bridge domain with different encapsulation VLANs on different switches.

Figure 4: Separate VNID Associated to External Bridge Domains with Different Encapsulation (pre-ACI 5.2(3) Releases).

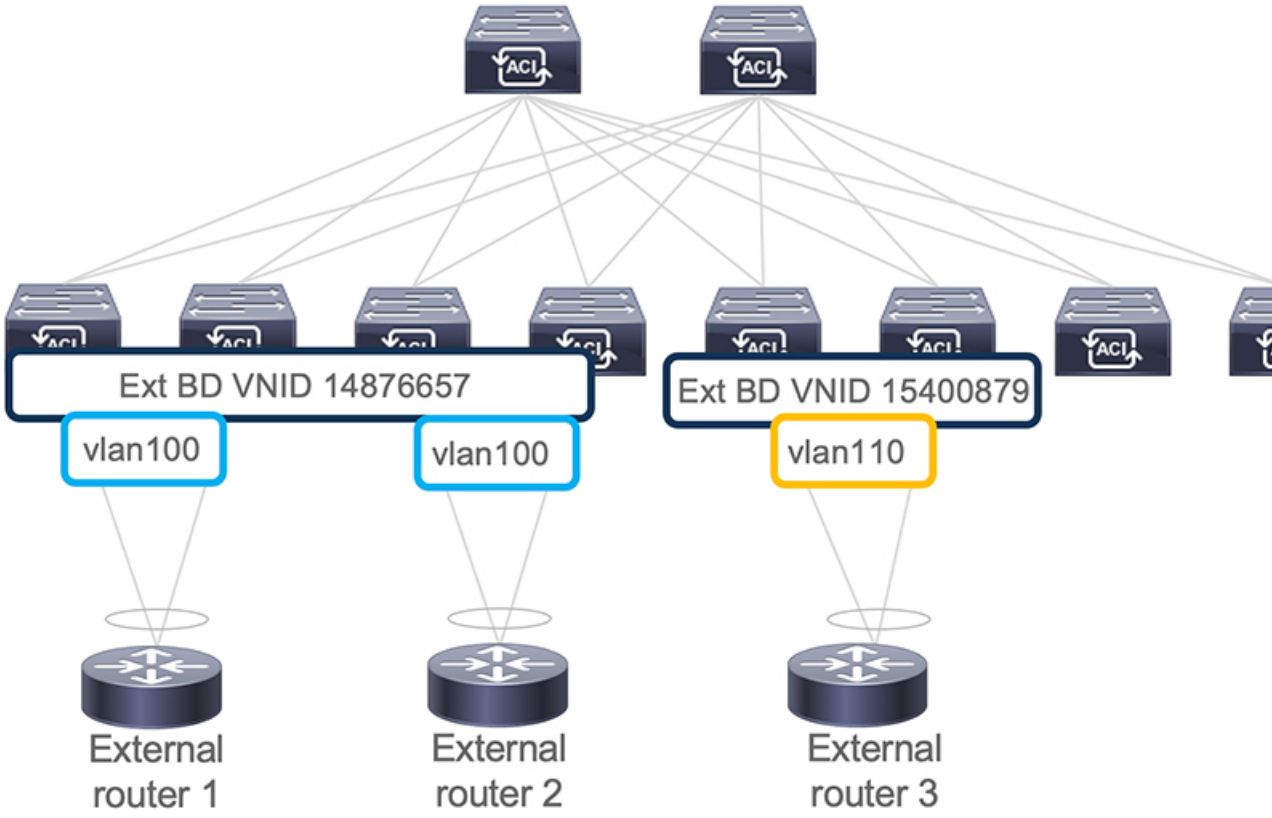

Release 5.2(3) added support for configuring a single external bridge that can be configured with different encapsulation VLANs on different leaf switches. The multiple encapsulation support feature uses the floating

SVI object to define the external bridge domain for floating L3Outs or an external bridge group profile for defining the external bridge domain for regular L3Outs. The use case for this feature may be where the same VLAN cannot be used on different leaf switches because it may already be in use.

Figure 5: Single VNID Associated to External Bridge Domains with Different Encapsulation (post-ACI 5.2(3) Releases).

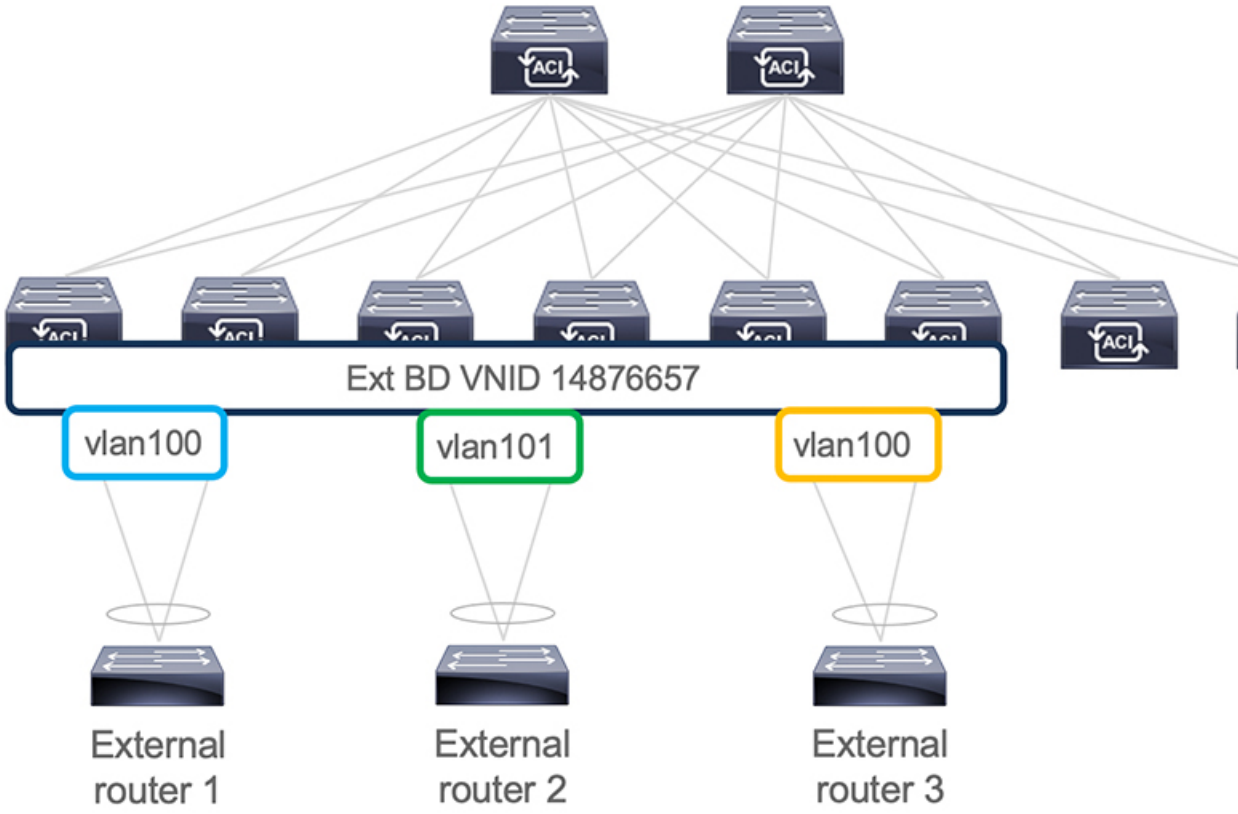

As of ACI release 6.0(1), this feature is supported for physical domain L3Outs only, not for VMM domain L3Outs.

Grouping Multiple SVIs With Different Access Encapsulation

The following figure shows a configuration where multiple SVIs are grouped together with different access encapsulation.

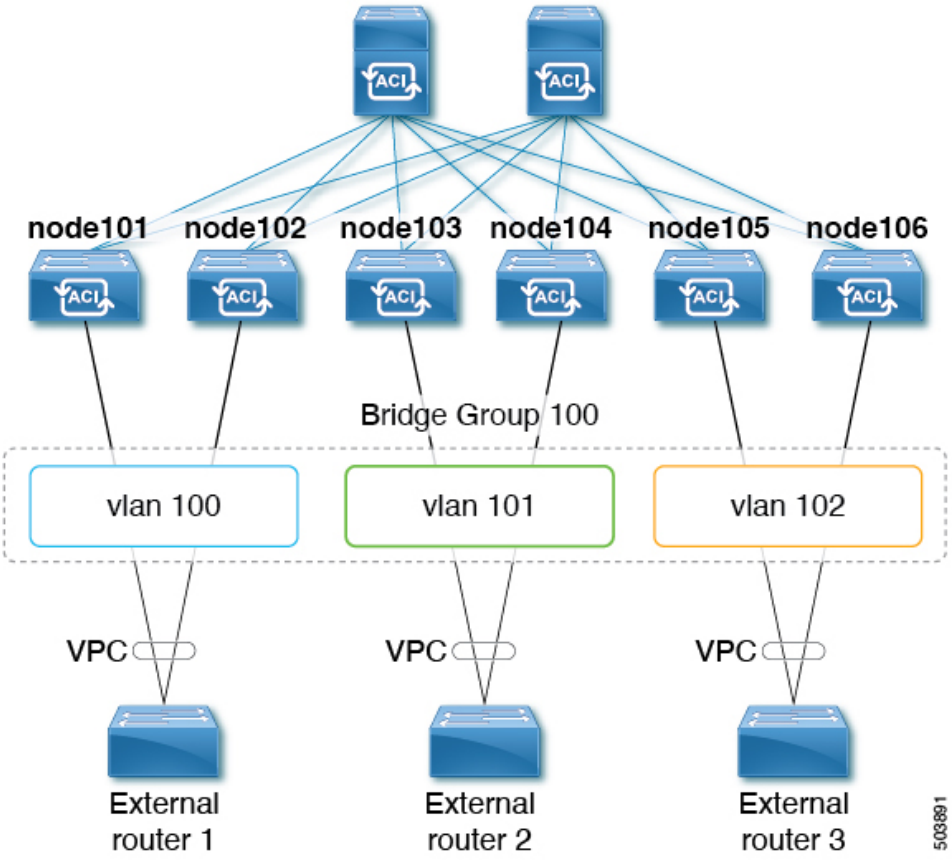

For this use case:

- The following leaf switches are VPC pairs:
  - node101 and node102
  - node103 and node104
  - node105 and node106

To configure the use case shown above, where you are grouping multiple SVIs into a Layer 2 bridge group:

1. Create three regular SVIs for each VPC pair:
   - Create the regular SVI **svi-100** on leaf switches node101 and node102
   - Create the regular SVI **svi-101** on leaf switches node103 and node104
   - Create the regular SVI **svi-102** on leaf switches node105 and node106
2. Configure the leaf switches with access encapsulations:
   - Configure leaf switches node101 and node102 with access encapsulation **vlan100**
   - Configure leaf switches node103 and node104 with access encapsulation **vlan101**
   - Configure leaf switches node105 and node106 with access encapsulation **vlan102**

3. Group the regular SVIs **svi-100**, **svi-101**, and **svi-102** together to behave as part of a single Layer 2 broadcast domain:
   a. Create a bridge domain profile.

      The bridge domain profile is represented by the new MO *l3extBdProfile*
   b. Provide a unique name string for the bridge domain profile.
   c. Associate each of the regular SVIs that need to be grouped together to the same bridge domain profile.

      Two new MOs are available for this association: *l3extBdProfileCont* and *l3extRsBdProfile*.

Guidelines and Limitations

- Layer 2 loops are blocked by the external device/hypervisor. Loops may occur if this feature is used with external switches that rely on spanning tree protocol to prevent loops.
- The SVI will be deleted and re-added after configuring the external bridge domain profile on them.
- The external bridge domain profile is L3Out-scoped. On a node, you cannot have two different access encapsulation mappings to the same external bridge domain profile.
- Bridge domain grouping is not supported with encapsulation scope **ctx** (the **VRF** option in the APIC GUI).
- Grouped SVIs with different line encapsulation can not share any common nodes.
- If you downgrade from release 5.2(3) to a previous release where multiple encapsulation for L3Outs with SVI is not supported, the following actions will be performed on the L3Out that was configured with multiple encapsulations and/or the external bridge domain profile:
  - The new allocator used for the multiple encapsulation support (*l3extBdProfileEncapAllocator*) will be deleted
  - All external bridge domain profiles (new *l3extBdProfile* MOs) will be deleted
  - All new *l3extBdProfileCont* MOs will be deleted
  - All new *l3extRsBdProfile* MOs will be deleted

Configuring Multiple Encapsulation for L3Outs With SVI Using the GUI

Procedure

Step 1 Create the regular SVIs and configure the leaf switches with access encapsulations.

See Configuring SVI External Encapsulation Scope Using the GUI, on page 15 for those procedures.

Step 2 Create an external bridge group profile that will be used for SVI grouping.

a) Navigate to **Tenants** > ***tenant-name*** > **Policies** > **Protocol** > **External Bridge Group Profiles**.
A page showing the already-configured external bridge group profiles is displayed.

b) Right-click on **External Bridge Group Profiles** and choose **Create External Bridge Group Profile**.

The **Create External Bridge Group Profile** page is displayed.

c) Enter a name for the external bridge group profile, then click **Submit**.
The page showing the already-configured external bridge group profiles is updated with the new external bridge group profile.

Step 3 Associate a regular SVI with the bridge domain profile.

a) Navigate to **Tenants** > ***tenant-name*** > **Networking** > **L3Outs** > ***L3Out-name*** > **Logical Node Profile** > ***log-node-profile-name*** > **Logical Interface Profile** > ***log-int-profile-name***.
The **General** page for this logical interface profile is displayed.

b) Click on the **SVI** tab.
A page showing the already-configured switch virtual interfaces is displayed.

c) Double-click on the switch virtual interface that you want to associate with the external bridge domain profile.
General information for this switch virtual interface is displayed.

d) In the **External Bridge Group Profile** field, select the external bridge domain profile that you want to associate with this switch virtual interface.

e) Click **Submit**.

Configuring Multiple Encapsulation for L3Outs With SVI Using the CLI

Procedure

Step 1 Create the regular SVIs and configure the leaf switches with access encapsulations.

See Configuring SVI Interface Encapsulation Scope Using NX-OS Style CLI for those procedures.

Step 2 Log into your APIC through the CLI, then go into configuration mode and tenant configuration mode.

```
apic1#
apic1# configuration
apic1(config)# tenant <tenant-name>
apic1(config-tenant)#
```

Step 3 Enter the following commands to create an external bridge profile that will be used for SVI grouping.

```
apic1(config-tenant)# external-bridge-profile <bridge-profile-name>
apic1(config-tenant-external-bridge-profile)# ?
```

Step 4 Enter the following commands to associate a regular SVI with the bridge domain profile.

```
apic1(config)# leaf <leaf-ID>
apic1(config-leaf)# interface vlan <vlan-num>
apic1(config-leaf-if)# vrf member tenant <tenant-name> vrf <VRF-name>
apic1(config-leaf-if)# ip address <IP-address>
apic1(config-leaf-if)# external-bridge-profile <bridge-profile-name>
```

Configuring Multiple Encapsulation for L3Outs With SVI Using the REST API

Procedure

Step 1 Create the regular SVIs and configure the leaf switches with access encapsulations.

See Configuring SVI Interface Encapsulation Scope Using the REST API for those procedures.

Step 2 Enter a post such as the following example to create an external bridge profile that will be used for SVI grouping.

```
<fvTenant name="t1" dn="uni/tn-t1" >
    <l3extBdProfile name="bd100" status=""/>
</fvTenant>
```

Step 3 Enter a post such as the following example to associate a regular SVI with the bridge domain profile.

```
<fvTenant name="t1">
    <l3extOut name="l1">
        <l3extLNodeP name="n1">
            <l3extLIfP name="i1">
                <l3extRsPathL3OutAtt encap="vlan-108"
                    tDn="topology/pod-1/paths-108/pathep-[eth1/10]"
                    ifInstT="ext-svi">
                    <l3extBdProfileCont>
                        <l3extRsBdProfile tDn="uni/tn-t1/bdprofile-bd100" status=""/
                    </l3extBdProfileCont>
                </l3extRsPathL3OutAtt>
            </l3extLIfP>
        </l3extLNodeP>
    </l3extOut>
</fvTenant>
```

Step 4 Enter a post such as the following example to specify the separate encapsulation for floating nodes.

```
<fvTenant name="t1">
    <l3extOut name="l1">
        <l3extLNodeP name="n1">
            <l3extLIfP name="i1">
                <l3extVirtualLIfP addr="10.1.0.1/24"
                    encap="vlan-100"
                    nodeDn="topology/pod-1/node-101"
                    ifInstT="ext-svi">
                    <l3extRsDynPathAtt floatingAddr="10.1.0.100/24"
                        encap="vlan-104"
                        tDn="uni/phys-phyDom"/>
                </l3extVirtualLIfP>
            </l3extLIfP>
    </l3extOut>
</fvTenant>
```

SVI Auto State

About SVI Auto State

Note This feature is available in the APIC Release 2.2(3x) release and going forward with APIC Release 3.1(1). It is not supported in APIC Release 3.0(x).

The Switch Virtual Interface (SVI) represents a logical interface between the bridging function and the routing function of a VLAN in the device. SVI can have members that are physical ports, direct port channels, or virtual port channels. The SVI logical interface is associated with VLANs, and the VLANs have port membership.

The SVI state does not depend on the members. The default auto state behavior for SVI in Cisco APIC is that it remains in the up state when the auto state value is disabled. This means that the SVI remains active even if no interfaces are operational in the corresponding VLAN/s.

If the SVI auto state value is changed to enabled, then it depends on the port members in the associated VLANs. When a VLAN interface has multiple ports in the VLAN, the SVI goes to the down state when all the ports in the VLAN go down.

Table 2: SVI Auto State

SVI Auto State	Description of SVI State
Disabled	SVI remains in the up state even if no interfaces are operational in the corresponding VLAN/s. Disabled is the default SVI auto state value.
Enabled	SVI depends on the port members in the associated VLANs. When a VLAN interface contains multiple ports, the SVI goes into the down state when all the ports in the VLAN go down.

Guidelines and Limitations for SVI Auto State Behavior

Read the following guidelines:

- When you enable or disable the auto state behavior for SVI, you configure the auto state behavior per SVI. There is no global command.

Configuring SVI Auto State Using the GUI

Before you begin

- The tenant and VRF configured.
- An L3Out is configured and a logical node profile and a logical interface profile under the L3Out is configured.

Procedure

Step 1 On the menu bar, click > **Tenants** > ***Tenant_name***.

Step 2 In the **Navigation** pane, click **Networking** > **L3Outs** > ***L3Out_name*** > **Logical Node Profiles** > ***LogicalNodeProfile_name*** > **Logical Interface Profiles**.

Step 3 In the **Navigation** pane, expand **Logical Interface Profile**, and click the appropriate logical interface profile.

Step 4 In the Work pane, click the **SVI** tab, then click the + sign to display the **SVI** dialog box.

Step 5 To add an additional SVI, in the **SVI** dialog box, perform the following actions:

a) In the **Path Type** field, choose the appropriate path type.
b) In the **Path** field, from the drop-down list, choose the appropriate physical interface.
c) In the **Encap** field, choose the appropriate values.
d) In the **Auto State** field, choose the SVI in the **Work** pane, to view/change the Auto State value.

The default value is **Disabled**.

Note
To verify or change the Auto State value for an existing SVI, choose the appropriate SVI and verify or change the value.

About Cisco Floating L3Outs

Beginning with the Cisco Application Policy Infrastructure Controller (APIC) release 4.2(1), you no longer need to specify multiple Layer 3 outside network connection (L3Out) logical interface paths to connect external network devices.

The floating L3Out feature enables you to configure a L3Out without specifying logical interfaces. The feature saves you from having to configure multiple L3Out logical interfaces to maintain routing when virtual machines (performing a specific virtual network function) move from one host to another. Floating L3Out is supported for VMM domains with VMware vSphere Distributed Switch (VDS).

Beginning with the Cisco APIC release 5.0(1), physical domains are supported. This means that the same simplified configuration can be used for physical routers deployments as well.

For more information, see the *Using Floating L3Out to Simplify Outside Network Connections* knowledge base article:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/Cisco-ACI-Floating-L3Out.html