



Shared L3Out

This chapter contains the following sections:

- [Shared Layer 3 Out, on page 1](#)
- [Layer 3 Out to Layer 3 Out Inter-VRF Leaking, on page 4](#)

Shared Layer 3 Out

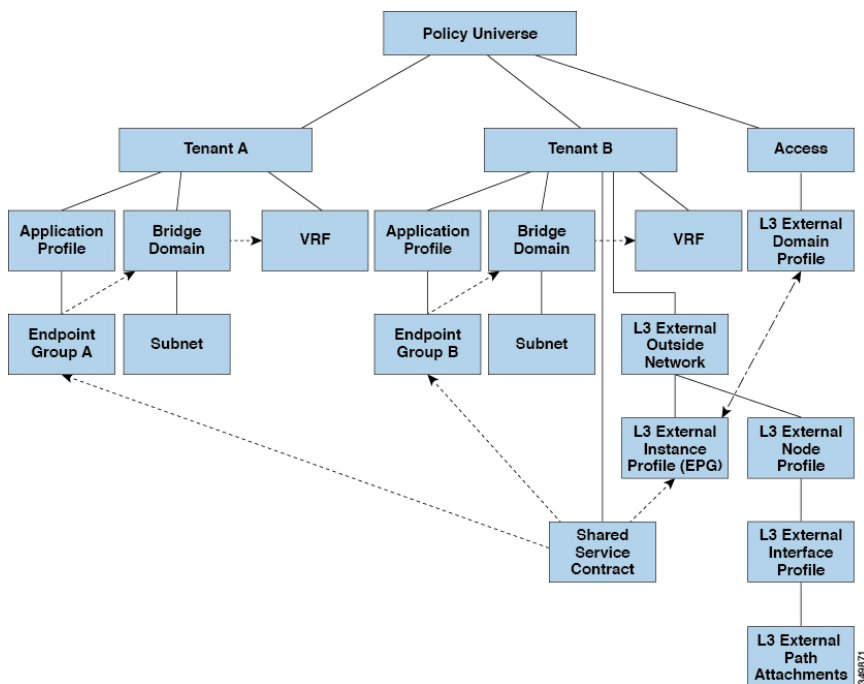
A shared Layer 3 outside network (`L3extOut`) configuration provides routed connectivity to external networks as a shared service. An `L3extOut` profile (`l3extInstP`) EPG provides routed connectivity to external networks. It can be provisioned as a shared service in any tenant (`user`, `common`, `infra`, or `mgmt`). Prior to release 1.2(1x), this configuration was only supported in the `user` and `common` tenants. An EPG in any tenant can use a shared services contract to connect with an `l3extInstP` EPG regardless of where in the fabric that `l3extInstP` EPG is provisioned. This simplifies the provisioning of routed connectivity to external networks; multiple tenants can share a single `l3extInstP` EPG for routed connectivity to external networks. Sharing an `l3extInstP` EPG is more efficient because it consumes only one session on the switch regardless of how many EPGs use the single shared `l3extInstP` EPG.



Note All switches that will use `l3extInstP` EPG shared service contracts require the hardware and software support available starting with the APIC 1.2(1x) and switch 11.2(1x) releases. Refer to the *Cisco APIC Management, Installation, Upgrade, and Downgrade Guide* and Release Notes documentation for more details.

The figure below illustrates the major policy model objects that are configured for a shared `l3extInstP` EPG.

Figure 1: Shared Layer 3 Out Policy Model



Take note of the following guidelines and limitations for shared Layer 3 outside network configurations:

- No tenant limitations: Tenants A and B can be any kind of tenant (*user*, *common*, *infra*, *mgmt*.). The shared `L3extInstP` EPG does not have to be in the *common* tenant.
- Flexible placement of EPGs: EPG A and EPG B in the illustration above are in different tenants. EPG A and EPG B could use the same bridge domain and VRF, but they are not required to do so. EPG A and EPG B are in different bridge domains and different VRFs but still share the same `L3extInstP` EPG.
- A subnet can be *private*, *public*, or *shared*. A subnet that is to be advertised into a consumer or provider EPG of an `L3extOut` must be set to *shared*. A subnet that is to be exported to an `L3extOut` must be set to *public*.
- The shared service contract is exported from the tenant that contains the `L3extInstP` EPG that provides shared Layer 3 outside network service. The shared service contract is imported into the tenants that contain the EPGs that consume the shared service.
- Do not use taboo contracts with a shared L3 out; this configuration is not supported.
- The `L3extInstP` as a shared service provider is supported, but only with non `L3extInstP` consumers (where the `L3extOut` EPG is the same as the `L3extInstP`).
- Traffic Disruption (Flap): When an `L3instP` EPG is configured with an external subnet of 0.0.0.0/0 with the scope property of the `L3instP` subnet set to shared route control (*shared-rctrl*), or shared security (*shared-security*), the VRF is redeployed with a global `pcTag`. This will disrupt all the external traffic in that VRF (because the VRF is redeployed with a global `pcTag`).
- Prefixes for a shared `L3extOut` must be unique. Multiple shared `L3extOut` configurations with the same prefix in the same VRF will not work. Be sure that the external subnets (external prefixes) that are advertised into a VRF are unique (the same external subnet cannot belong to multiple `L3instPs`). An `L3extOut` configuration (for example, named `L3Out1`) with `prefix1` and a second Layer 3 outside

configuration (for example, named L3Out2) also with prefix1 belonging to the same VRF will not work (because only 1 pcTag is deployed). Different behaviors of L3extOut are possible when configured on the same leaf switch under the same VRF. The two possible scenarios are as follows:

- Scenario 1 has an L3extOut with an SVI interface and two subnets (10.10.10.0/24 and 0.0.0.0/0) defined. If ingress traffic on the Layer 3 outside network has the matching prefix 10.10.10.0/24, then the ingress traffic uses the External EPG pcTag. If ingress traffic on the Layer 3 Outside network has the matching default prefix 0.0.0.0/0, then the ingress traffic uses the External Bridge pcTag.
- Scenario 2 has an L3extOut using a routed or routed-sub-interface with two subnets (10.10.10.0/24 and 0.0.0.0/0) defined. If ingress traffic on the Layer 3 outside network has the matching prefix 10.10.10.0/24, then the ingress traffic uses the External EPG pcTag. If ingress traffic on the Layer 3 outside network has the matching default prefix 0.0.0.0/0, then the ingress traffic uses the VRF pcTag.
- As a result of these described behaviors, the following use cases are possible if the same VRF and same leaf switch are configured with L3extOut-A and L3extOut-B using an SVI interface:

Case 1 is for L3extOut-A: This External Network EPG has two subnets defined: 10.10.10.0/24 & 0.0.0.0/1. If ingress traffic on L3extOut-A has the matching prefix 10.10.10.0/24, it uses the external EPG pcTag & contract-A which is associated with L3extOut-A. When egress traffic on L3extOut-A has no specific match found, but there is a maximum prefix match with 0.0.0.0/1, it uses the External Bridge Domain (BD) pcTag & contract-A.

Case 2 is for L3extOut-B: This External Network EPG has one subnet defined: 0.0.0.0/0. When ingress traffic on L3extOut-B has the matching prefix 10.10.10.0/24 (which is defined under L3extOut-A), it uses the External EPG pcTag of L3extOut-A and the contract-A which is tied with L3extOut-A. It does not use contract-B which is tied with L3extOut-B.

- Traffic not permitted: Traffic is not permitted when an invalid configuration sets the scope of the external subnet to shared route control (shared-rtctrl) as a subset of a subnet that is set to shared security (shared-security). For example, the following configuration is invalid:
 - *shared rtctrl*: 10.1.1.0/24, 10.1.2.0/24
 - *shared security*: 10.1.0.0/16

In this case, ingress traffic on a non-border leaf with a destination IP of 10.1.1.1 is dropped, since prefixes 10.1.1.0/24 and 10.1.2.0/24 are installed with a drop rule. Traffic is not permitted. Such traffic can be enabled by revising the configuration to use the *shared-rtctrl* prefixes as *shared-security* prefixes as well.

- Inadvertent traffic flow: Prevent inadvertent traffic flow by avoiding the following configuration scenarios:
 - **Case 1** configuration details:
 - A Layer 3 outside network configuration (for example, named L3extOut-1) with VRF1 is called provider1.
 - A second Layer 3 outside network configuration (for example, named L3extOut-2) with VRF2 is called provider2.
 - L3extOut-1 VRF1 advertises a default route to the Internet, 0.0.0.0/0 which enables both *shared-rtctrl* and *shared-security*.
 - L3extOut-2 VRF2 advertises specific subnets to DNS and NTP, 192.0.0.0/8 which enables *shared-rtctrl*.

- L3extOut-2 VRF2 has specific subnet 192.1.0.0/16, which enables *shared-security*.
- **Variation A:** EPG Traffic Goes to Multiple VRFs.
 - Communications between EPG1 and L3extOut-1 is regulated by an *allow_all* contract.
 - Communications between EPG1 and L3extOut-2 is regulated by an *allow_all* contract.
 - Result:** Traffic from EPG1 to L3extOut-2 also goes to 192.2.x.x.
- **Variation B:** An EPG conforms to the *allow_all* contract of a second shared Layer 3 outside network.
 - Communications between EPG1 and L3extOut-1 is regulated by an *allow_all* contract.
 - Communications between EPG1 and L3extOut-2 is regulated by an *allow_icmp* contract.
 - Result:** Traffic from EPG1 to L3extOut-2 to 192.2.x.x conforms to the *allow_all* contract.
- **Case 2** configuration details:
 - A L3extOut profile (l3instP) has one shared prefix and other non-shared prefixes.
 - Traffic coming in with `src = non-shared` is allowed to go to the EPG
 - **Variation A:** Unintended traffic goes through an EPG.

L3extOut (l3instP) EPG traffic goes through a L3extOut that has these prefixes:

 - 192.0.0.0/8 = `import-security, shared-rtctrl`
 - 192.1.0.0/16 = `shared-security`
 - The EPG has 1.1.0.0/16 = `shared`

Result: Traffic going from 192.2.x.x also goes through to the EPG.
 - **Variation B:** Unintended traffic goes through an EPG. Traffic coming in a shared L3extOut can go through the EPG.
 - The shared L3extOut VRF has an EPG with `pcTag = prov vrf` and a contract set to *allow_all*
 - The EPG <subnet> = `shared`.

Result: The traffic coming in on the Layer 3 out can go through the EPG.

Layer 3 Out to Layer 3 Out Inter-VRF Leaking

Starting with Cisco APIC release 2.2(2e), when there are two Layer 3 Outs in two different VRFs, inter-VRF leaking is supported.

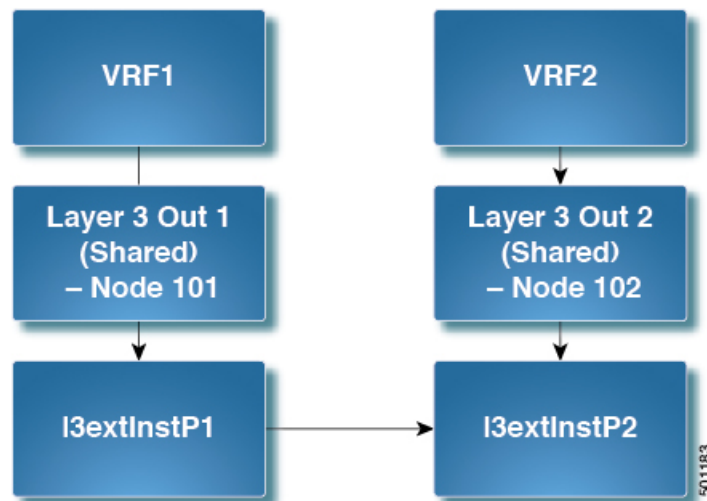
For this feature to work, the following conditions must be satisfied:

- A contract between the two Layer 3 Outs is required.

- Routes of connected and transit subnets for a Layer 3 Out are leaked by enforcing contracts (L3Out-L3Out as well as L3Out-EPG) and without leaking the dynamic or static routes between VRFs.
- Dynamic or static routes are leaked for a Layer 3 Out by enforcing contracts (L3Out-L3Out as well as L3Out-EPG) and without advertising directly connected or transit routes between VRFs.
- Shared Layer 3 Outs in different VRFs can communicate with each other.
- There is no associated L3Out required for the bridge domain. When an Inter-VRF shared L3Out is used, it is not necessary to associate the user tenant bridge domains with the L3Out in tenant `common`. If you had a tenant-specific L3Out, it would still be associated to your bridge domains in your respective tenants.
- Two Layer 3 Outs can be in two different VRFs, and they can successfully exchange routes.
- This enhancement is similar to the Application EPG to Layer 3 Out inter-VRF communications. The only difference is that instead of an Application EPG there is another Layer 3 Out. Therefore, in this case, the contract is between two Layer 3 Outs.

In the following figure, there are two Layer 3 Outs with a shared subnet. There is a contract between the Layer 3 external instance profile (l3extInstP) in both the VRFs. In this case, the Shared Layer 3 Out for VRF1 can communicate with the Shared Layer 3 Out for VRF2.

Figure 2: Shared Layer 3 Outs Communicating Between Two VRFs



Configuring Shared Layer 3 Out Inter-VRF Leaking Using the Advanced GUI

Before you begin

The contract label to be used by the consumer and provider is already created.

Procedure

-
- Step 1** On the menu bar, choose **Tenants > Add Tenant**.
- Step 2** In the **Create Tenant** dialog box, enter a tenant name for the provider.

- Step 3** In the **VRF Name** field, enter a VRF name for the provider, then click **Submit** to create the tenant.
- Step 4** In the **Navigation** pane, under the new tenant name, navigate to **L3Outs**.
- Step 5** Right-click on **L3Outs** and select **Create L3Out**.
The **Create L3Out** wizard appears.
- Step 6** In the **Create L3Out** dialog box, perform the following actions:
- In the **Name** field, enter a name for the L3Out.
 - In the **VRF** field, select the VRF that you created earlier.
 - In the **L3 Domain** field, select an L3 domain.
 - Make the appropriate selections for the protocols, then click **Next**.
- Step 7** Make the necessary selections in the next windows, until you get to the **External EPG** window.
You might see the **Nodes and Interfaces** window and the **Protocols** window, depending on the protocol that you selected in the **Identity** window. The last window in the **Create L3Out** wizard is the **External EPG** window.
- Step 8** In the **External EPG** window, perform the following actions:
- In the **Name** field, enter the external network name.
 - Uncheck the **Default EPG for all external networks** checkbox.
The **Subnets** fields appears.
 - Click + to access the **Create Subnet** window.
 - In the **Create Subnet** dialog box, in the **IP Address** field, enter the match IP address. Click **OK**.
 - Click **Finish** in the **Create L3Out** wizard.
- Step 9** In the **Navigation** pane, navigate to the **L3Out_name > External EPGs > ExternalEPG_name** that you created.
- Step 10** In the **Work** pane, under **Properties** for the external network, verify that the resolved VRF is displayed in the **Resolved VRF** field.
- Step 11** Double-click the IP address for external subnets to open the **Subnet** dialog box.
- Step 12** In the **Scope** field, check the desired check boxes, and then click **Submit**.
In this scenario, check the following check boxes:
- **External Subnets for the External EPG**
 - **Shared Route Control Subnet**
 - **Shared Security Import Subnet**
- Step 13** Navigate to the **L3 Outside** you created earlier.
- Step 14** In the **Provider Label** field, enter the provider name that was created as a prerequisite to starting this task. Click **Submit**.
- Step 15** On the menu bar, click **Tenants > Add Tenant**.
- Step 16** In the **Create Tenant** dialog box, enter a tenant name for the L3Out consumer.
- Step 17** In the **VRF name** field, enter a VRF name for the consumer.
- Step 18** In the **Navigation** pane, under the new tenant name, navigate to **L3Outs** for the consumer.
- Step 19** Right-click on **L3Outs** and select **Create L3Out**.

The **Create L3Out** wizard appears.

Step 20

In the **Create L3Out** dialog box, perform the following actions:

- a) In the **Name** field, enter a name for the L3Out.
- b) In the **VRF** field, from the drop-down menu, choose the VRF that was created for the consumer.
- c) In the **Consumer Label** field, enter the name for the consumer label.
- d) In the **L3 Domain** field, select an L3 domain.
- e) Make the appropriate selections for the protocols, then click **Next**.

Step 21

Make the necessary selections in the next windows, until you get to the **External EPG** window.

You might see the **Nodes and Interfaces** window and the **Protocols** window, depending on the protocol that you selected in the **Identity** window. The last window in the **Create L3Out** wizard is the **External EPG** window.

Step 22

In the **External EPG** window, perform the following actions:

- a) In the **Name** field, enter the external network name.
- b) Uncheck the **Default EPG for all external networks** checkbox.

The **Subnets** fields appears.

- c) Click + to access the **Create Subnet** window.
- d) In the **Create Subnet** dialog box, in the **IP Address** field, enter the match IP address. Click **OK**.
- e) In the **Scope** field, check the desired check boxes, and then click **OK**.

In this scenario, check the check boxes for **Shared Route Control Subnet** and **Shared Security Import Subnet**.

- f) Click **Finish** in the **Create L3Out** wizard.

This completes the configuration of shared Layer 3 Outside Inter-VRF leaking.

