



Endpoint Security Groups


This chapter contains the following topic:

- [Endpoint Security Groups, on page 1](#)
- [Contracts, on page 5](#)
- [ESG Shared Service \(ESG VRF route leaking\), on page 7](#)
- [Layer 4 to Layer 7 Services, on page 9](#)
- [Operational Tools, on page 10](#)
- [Limitations , on page 10](#)
- [ESG Migration Strategy, on page 12](#)
- [Configuring Endpoint Security Groups, on page 15](#)

Endpoint Security Groups

Endpoint Security Groups (ESGs) are the new network security component in Cisco Application Centric Infrastructure (Cisco ACI). Although the endpoint groups (EPGs) have been providing the network security in Cisco ACI, EPGs have to be associated to a single bridge domain (BD) and used to define security zones within a BD. This is because the EPGs define both forwarding and security segmentation at the same time. The direct relationship between the BD and an EPG limits the possibility of an EPG to spanning more than one BD. This limitation of EPGs is resolved by using the new ESG constructs.

Figure 1: Cisco ACI offers multiple segmentation options



The Application endpoint group (fvAEPg) object that represents an EPG has a direct relationship with the bridge domain object (fvBD) that represents the Layer 2 broadcast domain. This is illustrated in the above figure, in the first three columns.

An ESG is a logical entity that contains a collection of physical or virtual network endpoints. In addition, an ESG is associated to a single VRF (Virtual Routing and Forwarding) instead of BD. This allows the definition of a security zone that is independent of the BDs (the fourth column of *Figure 1*, illustrates this point). Just as the EPGs divide a BD into security zones, the ESGs divide the VRF into security zones.

The EPG policy embeds both forwarding and security logic. For example, an EPG provides not only a security zone based on VLAN, but also a VLAN binding on leaf node interfaces. Also, a contract on the EPG is used to enforce the security and determine which leaf nodes the BD subnet should be deployed on, and which subnets to be leaked to which VRF in the case of VRF route leaking (i.e. shared service). On the contrary, an ESG is used only to enforce security using the contracts while the forwarding logics are handled by other components. With an ESG, the routing logic such as BD subnets deployment and VRF route leaking are moved to VRF level. The VLAN binding on leaf node interfaces are still handled at EPG level.

An ESG is a security construct that has certain match criteria to define which endpoint belongs to the ESG, and uses contracts or policies to define the security stance. The match criteria are called the ESG selectors that are based on attributes, such as an IPv4 or IPv6 address spanning across BDs in the associated VRF. For Cisco APIC, release 5.0(1), the available matching criteria is the IP address or the subnet of an endpoint.


The contract usage in the ESGs is the same as the EPGs. Endpoints that belong to the same ESG can communicate without the need for a contract. To enable communication between endpoints that belong to different ESGs, you need to configure contracts between the ESGs. For the communication with devices

outside of the Cisco ACI fabric, you need to configure a contract between the L3Out external EPG (`l3extInstP`) and the ESG. You can also use a Layer 4 to Layer 7 service graph in conjunction with a contract between the ESGs. However, contracts between an EPG and an ESG are not supported.

Traffic Filtering from ESG to ESG

In the figure below, there are four bridge domains associated with one EPG each. The administrator uses the EPG configuration to ensure that traffic from virtual machines or from physical servers is associated with the appropriate bridge domain connected to the appropriate VLAN. For instance EPG1-1 defines the mapping of the traffic from VLAN 10 with BD1, the EPG2-1 maps VLAN 20 to BD2, and so on.

Figure 2: ESGs can be used to aggregate endpoints of different subnets



- 192.168.1.11 on VLAN 10 and 192.168.2.11 on VLAN 20 belong to different subnets and different bridge domains.
- The administrator defines 192.168.1.11 and 192.168.2.11 as belonging to the same ESG.
- Similarly, 192.168.3.11 and 192.168.4.11 are associated to BD3 and BD4 (via EPG3-1 and EPG4-1) respectively, and they both belong to the same ESG.
- With the above configuration, 192.168.1.11 can freely communicate with 192.168.2.11.
- Similarly, 192.168.3.11 can communicate with 192.168.4.11. However, 192.168.1.11 (or 192.168.2.11) cannot communicate with either 192.168.3.11 or 192.168.4.11 without a contract.




Note The contracts that are used by the EPGs cannot be re-used by the ESGs, and vice versa.

Traffic Filtering from Outside to ESG

The configuration to allow outside to ESG communication is performed by a contract between an L3Out external EPG (L3extInstP) and the ESG as illustrated in the figure below. From the L3Out perspective, there is nothing different between contracts with the ESGs and contracts with the EPGs.

Figure 3: ESG to outside connectivity is implemented using the L3 External EPG



ESG Implementation

This section summarizes how the Cisco Application Policy Infrastructure Controller (APIC) programs leaf nodes, when an administrator configures the ESGs.

- Each ESG is associated with a VRF instance, and the ESG selectors define which endpoints within the VRF instance belongs to the ESG.
- The VRF instance (where an ESG is configured) can be configured either in ingress or egress policy enforcement mode.
- Cisco ACI instantiates the ESG configuration on all of the leaf nodes where the associated VRF instance is deployed.
- When an ESG is configured, all of the BD subnets in the associated VRF instance are present as static routes to the spine proxy on all of the leaf nodes where that VRF instance is present.
- ESGs are always deployed with the deployment immediacy of on-demand, and the associated contract rules are programmed only after an endpoint that matches the ESG selectors are learned on the given leaf node.

- The contracts between ESGs are programmed as policy-cam rules in the leaf node TCAM just as with the EPGs.
- The class-id used by the ESG is a global pcTag.
- Unlike the EPGs, contracts between ESGs create only security rules. ESGs are not used for network deployment such as subnet deployment, or route leaking.
- Even when ESGs are used for security enforcement instead of EPGs, EPGs are still required to configure VLAN bindings on leaf node interfaces.



Note Cisco APIC generates a unique number to identify each ESG, just as it does for EPGs. This number is called a pcTag or class-id.

Local pcTags are numbers that are unique within a VRF scope, which means that Cisco APIC can generate the same number to identify another EPG in a different VRF instance.

Global pcTags are numbers that are unique in the entire fabric regardless of which VRF instance the ESG (or EPG) belongs to. ESGs are always assigned a global pcTag.

Contracts

Contracts are the Cisco ACI equivalent of access control lists (ACLs). ESGs can only communicate with other ESGs according to the contract rules. The administrator uses a contract to select the types of traffic that can pass between ESGs, including the protocols and ports allowed. An ESG can be a provider, consumer, or both provider and consumer of a contract, and can consume multiple contracts simultaneously. ESGs can also be part of a preferred group so that multiple ESGs can talk freely with other ESGs that are part of the preferred group.

Supported Contracts relationship:

1. ESG ↔ ESG
2. ESG ↔ L3Out EPG
3. ESG ↔ inband-EPG
4. ESG ↔ vzAny

Contracts between the ESGs and the EPGs (or uSeg EPGs) are not supported. When an endpoint in an ESG needs to communicate with other endpoints in the EPG, the other endpoints need to be migrated to the ESGs first. vzAny or preferred group can be used as an alternative during the migration. Other contract-related features that are supported in a uSeg EPG, such as contract inheritance, an intra ESG contract, or intra ESG isolation, are also supported in an ESG. The exception is the Taboo Contract, which is not supported in an ESG.

vzAny

In alternative to using specific contracts between ESGs, you can also allow traffic between ESGs using a construct called vzAny.

vzAny represents all of the ESGs and EPGs in the given VRF instance. This also includes the L3Out external EPG (`l3extInstP`) within a VRF instance. The vzAny construct provides a shorthand way to refer to all the EPGs and ESGs within that VRF instance. The vzAny referral eases management by allowing for a single point of contract configuration for all EPGs and ESGs within a VRF instance, and optimizes hardware resource consumption by applying the contract to this one group rather than to each EPG or ESG individually.

Figure 4: vzAny is a shorthand to represent all EPGs and ESGs in the same VRF instance




Figure 4 provides an example. If the administrator configures a contract between vzAny and EPG 4-1, in the topology from Figure 2, endpoints 192.168.1.11, 192.168.2.11 (ESG1) and 192.168.3.11 (EPG3-1) can communicate with 192.168.4.11 (EPG4-1).

This does not mean that ESG1 and EPG3-1 belong to the same security zone and 192.168.11 (or 192.168.2.11) can communicate with 192.168.3.11 without a contract. If the desired configuration is to allow any-to-any communication within the VRF instance regardless of an ESG, an EPG, L3Out EPG etc., the user can configure vzAny to provide and consume a contract to allow all traffic instead of disabling **Policy Enforcement (Unenforced)** in the VRF instance.


In summary, the vzAny construct can be used for providing and (or) consuming a contract in order to enable an ESG to communicate with anybody in the VRF instance using the contract just as it does for an EPG. Even though the contracts between ESGs and the EPGs are not allowed, vzAny contracts can be used to allow traffic between the ESGs and EPGs.

Preferred Groups

A preferred group is an alternative to using explicit contracts between ESGs or using vzAny contracts. The user can also configure the preferred group to enable the communication between ESGs in a VRF instance. Any endpoints in the preferred group can communicate with each other freely.

The user can also use preferred groups to enable ESGs to EPGs communication which can be useful in a migration between an EPG-based security configuration to an ESG-based security configuration.

Figure 5: Example with ESG1 and EPG3-1 part of the same preferred group.



In the example of the figure above, ESG1 and EPG3-1 are configured to be part of the preferred group of VRF A and the following communications are allowed:

1. ESG 1 and EPG 3-1 can communicate each other since both are included in the preferred group.
2. ESG 1 and EPG 4-1 cannot communicate each other because:
 - EPG 4-1 is not included in the preferred group.
 - Contracts between EPGs and ESGs are not supported.

Refer to the [Cisco APIC Basic Configuration Guide](#) for information on configuring preferred groups.

ESG Shared Service (ESG VRF route leaking)

When an endpoint needs a service that is shared by another VRF, there are two things required for the communication to happen. The first thing is the routing reachability. The second thing is security permission. In an EPG, these two are coupled closely in one set of configurations, such as the EPG subnet and contracts. In ESG, these two are decoupled in two different configurations:

1. The configuration of route leaking at the VRF level, which is independent of the ESG contract configuration.
2. The configuration of contracts between the ESGs.

With these two configurations completely decoupled, you do not need to configure a subnet or a subset of the subnet under the ESG as you must do for an EPG.

The following sections explain how to configure route leaking for the bridge domain subnets and external prefixes learned from external routers. After you finish configuring route leaking, you can configure a contract between two ESGs, or an ESG and L3Out EPG, to allow the communication. You must use a contract with a scope larger than VRF, such as global.



Note The route leaking configuration at the VRF level is supported only for ESGs.


Route Leaking for Internal Bridge Domain Subnets

This section explains how to configure route leaking between VRF instances for a bridge domain subnet to which the ESG endpoints belong to. This is performed simply by specifying a subnet to leak and the target VRF instance in the source VRF instance at the VRF level (instead of at the EPG level like it is done if you do not use ESGs). The subnet that you enter in the route leaking configuration needs to match the bridge domain subnet or be a subset of a configured bridge domain subnet. The route leaked by this configuration is only the subnet with the specified subnet mask. You cannot specify a range of subnets to leak multiple bridge domain subnets in one configuration.



Note The subnet that you configure under the VRF route leaking configuration can also match subnets used under the EPGs. This can be useful for migration purposes.

Figure 6: Route Leaking with ESGs



The figure above provides an example of VRF leaking between two VRF instances: VRF A and VRF B, where the administrator has configured two ESGs: ESG1 and ESG2.

In addition to having a contract between ESG1 and ESG2 (to allow the traffic), the administrator needs to configure route leaking in the VRF instance as described in the section, [Configuring Route Leaking of Internal Bridge Domain Subnets using the GUI](#).

The configuration of the bridge domain subnet scopes, **Advertised Externally** and **Shared between VRFs**, is not required with VRF level route leaking for an ESG. When a leaked bridge domain subnet needs to be advertised by L3Outs in the target VRF instance, you can set **Allow L3Out Advertisement** to **True** in the VRF level route leaking configuration. Note that the subnet scopes under a bridge domain are ignored when leaking the subnet to the target VRF instance specified in the VRF level route leaking, and the configuration in the VRF level route leaking takes precedence. Those scopes under a bridge domain are still honored at the same time for other configurations like advertising the subnet from an L3Out in the same VRF instance, route leaking to another VRF instance through a traditional configuration that is through EPG contracts, or both.

Route Leaking for External Prefixes

The configuration of route leaking for the purpose of allowing traffic from a L3Out of a VRF to ESGs of another VRF is referred to as **ESG shared L3Out** to differentiate from the shared L3Out for EPGs.

In order to leak routes that are learned from a L3Out for an ESG communication, the administrator must configure the route leaking for external prefixes in VRF level. This is done by using IP prefix-list style configuration. The user can configure a specific prefix or can specify a range of prefixes by using the “le” (less than or equal to) or “ge” (greater than or equal to) as you can with an IP prefix-list in a normal router. Unlike bridge domain subnets, there is no restriction that the leaked prefix must be equal to or smaller than an actual route, because external routes are dynamically learned and are not often predictable. Because of the lack of the restriction, a leaked external prefix can specify a range to leak multiple prefixes with one configuration. In the configuration, you must also specify the target VRF.

Please refer to [Configuring Route Leaking of External Prefixes Using the GUI](#) for the configuration details.

For an ESG shared L3Out configuration, along with configuring route leaking in the VRF and applying a contract with L3Out EPG, you need to define which prefix belongs to which L3Out EPG. To specify which prefix belongs to which L3Out EPG, you must configure an L3Out subnet with the **External Subnets for the External EPG** and **Shared Security Import Subnet** scopes.

Layer 4 to Layer 7 Services

All the Layer 4 to Layer 7 service graph features that are available for the EPGs are supported for the ESGs. The use of service graph between ESGs of different VRFs is not supported in Cisco APIC 5.0.



Note This note is an implementation detail for advanced user information. If a service graph is attached to a contract between ESGs, the Cisco APIC automatically creates hidden service EPGs where the Layer 4 to Layer 7 service device attaches, just as Cisco APIC does for a service graph between EPGs. Unlike a service graph between EPGs, in the case of ESGs, the hidden service EPGs get a global pcTag.

Beginning with Cisco APIC release 5.0(1), all new service EPGs that are created for Layer 4 to Layer 7 Service deployments with vzAny-to-vzAny contracts will get global PcTag.

For more information on Layer 4 to Layer 7 services deployment, refer to [Cisco APIC Layer 4 to Layer 7 Services Deployment Guide](#).

Operational Tools

Capacity Dashboard

The **Capacity Dashboard** tab can be used to get a summary of critical fabric resource thresholds. This allows you to see quickly how close you are to reaching the approved scalability limits. Per leaf node usage is also shown, allowing you to see quickly which leaf node may be hitting resource constraints.

1. In the menu bar, choose **Operations > Capacity Dashboard** to launch the Capacity Dashboard troubleshooting tool.
2. In the **Capacity Dashboard** page, choose **Fabric Capacity** for the fabric resources. Scroll down for the **Endpoint Security Groups** tile and the **Global pcTag** tile to determine the available resources.
3. In the **Capacity Dashboard** page, choose **Leaf Capacity** for the leaf usage. Check the **ESG** tab for details on the resource usage for Endpoint Security Groups.

Endpoint Tracker

The **Endpoint Tracker** tab allows you to enter a fabric-attached endpoint IP or MAC address and quickly see the location of this endpoint, the endpoint group to which the endpoint belongs, the VLAN encapsulation used, and if any transitions (flaps) have occurred for this endpoint.

1. In the menu bar, click **Operations > EP Tracker** to launch the Endpoint Tracker troubleshooting tool.
2. In the **End Point Search** field, enter the IP address or MAC address of the endpoint and click **Search**.
3. Click on the endpoint after it is displayed.

The Endpoint Tracker tool displays the date and time of each state transition along with the IP address, MAC address, owning endpoint group, action (attached or detached), physical node, interface, and VLAN encapsulation during the event.

The Endpoint Tracker tool uses an object called the fvCEp to find the endpoints that are learned in the fabric, for an ESG and as well as an EPG. An endpoint that belongs to an ESG is represented by two fvCEp objects, one for the EPG that provides VLAN binding, another for the ESG that provides security. Therefore, the Endpoint Tracker tool shows two entries (one for an EPG, another for an ESG) when used for the ESG endpoints.

Limitations

As of the Cisco APIC release 5.0(1), the following limitations apply:

- Contracts between ESGs and EPGs are not supported.
- The ESG feature is not integrated with Cisco ACI Multi-Site. Other topologies such as Multi-Pod, Multi-Tier, and Remote Leaf are supported.

- The supported ESG selector is the IP address. MAC addresses, VM tags, or other criteria are not yet supported.
- An ESG contract can be applied only for routed traffic with IP as the selector.
- Taboo contracts are not supported with ESGs.
- Inter-VRF service graphs between ESGs are not supported.
- ESGs are not supported as a source or destination of the following features:
 - On Demand Atomic Counter
 - On Demand Latency Measurement
 - SPAN
- The following features configured at the BD or EPG level are not supported when endpoints in the BD/EPG are classified to an ESG:
 - Endpoint Reachability (Static routes on BD/EPG)
 - Anycast Service
 - Microsoft NLB
 - First Hop Security (FHS)
 - Host Based Routing / Host Route Advertisement
- Only the EX and newer generation of leaf nodes are supported for ESG deployment.
- To prevent Layer 2 traffic (that is, non-routed traffic) from bypassing ESG security when IP is used as the selector, enable an intra EPG contract with a permit-all rule, such as the common default contract, on all of the EPGs that provide VLAN-to-interface binding for the ESG endpoints. If all the endpoints in the EPGs are classified to ESGs, you can alternatively enable intra EPG isolation with proxy ARP on the EPGs instead of the intra EPG contract. If the EPG is used only for VMM DVS integration, you can alternately enable the **Allow Micro-Segmentation** option instead of the other two options mentioned above. Either feature forces all communication between ESG endpoints to go through Layer 3 routing.




Note This note explains the differences between an intra EPG contract with a permit-all rule and intra EPG isolation with proxy-ARP. The main purpose of both features is the same, which is to enforce all traffic to be routed on ACI leaf switches by using proxy-ARP. Note that proxy-ARP is enabled implicitly for the EPG when an intra EPG contract is used. The difference is when there are two or more endpoints that don't belong to ESGs but learned in an EPG. With an intra EPG contract with a permit-all rule, such endpoints can still communicate freely within the same EPG due to the permit-all rule. However, with intra EPG isolation with proxy-ARP, such endpoints can no longer communicate even though they are in the same EPG.

- Label configurations are not supported when you add contracts to an ESG.

ESG Migration Strategy

The contracts between an EPG to an ESG are not supported in the Cisco APIC, release 5.0(1). Therefore, to migrate the existing EPG deployments to an ESG with minimal impact, vzAny can be used to allow communication between an EPG and an ESG endpoints during the migration.

Figure 7: Migration to ESG within VRF



For migration to ESG within VRF


The following procedure explains how to migrate EPGs to ESGs within the same VRF, starting with the provider EPG, EPG4.

1. As shown in the example, firstly create an ESG Service.
2. Configure a new contract, with ESG Service as provider and vzAny as consumer. This will allow the communication between the existing consumer EPG endpoints in EPG1, EPG2 and EPG3 with the newly created ESG service.
3. Now migrate the provider EPG EPG4's endpoints into ESG Service by configuring the appropriate endpoint selectors.
4. This completes the migration of the provider EPG EPG4 to ESG.
5. Prepare the migration of individual consumer EPGs, namely EPG1, EPG2 and EPG3, by creating ESG IT and Sales that consume the same contract as the VRF vzAny.

Note: You can alternatively configure fine grain contracts for ESG IT and Sales respectively instead. In such a case, you must provide the new contract from ESG Service as well.

- Now migrate the consumer endpoints to ESGs by configuring the selectors on ESG IT and Sales.
- Once the migration of all the consumer EPGs are completed, the administrator can remove the VRF vzAny consumer configuration.

Figure 8: Migration to ESG across VRF



For migration to ESG across VRF

For shared services with EPGs, the administrator would have created a subnet under the provider EPG for route leaking across VRFs via contracts. As shown in the above example figure,

- provider EPG EPG4, part of provBD, has subnet 192.168.1.1/24 configured with *Shared between VRFs* scope enabled.
- Consumer EPGs EPG1, EPG2 and EPG3 are part of ConsBD that has subnet 172.16.1.1/24 with *Shared between VRFs* scope enabled.

The following procedure explains how to migrate EPGs to ESGs in the case of shared service, starting with the provider EPG, EPG4.

1. Create the subnet 192.168.1.1/24 on the ProvBD (provider EPG's BD) if not existing already.
Note: *Shared between VRFs* scope is not required.
2. Create a provider ESG Service .
3. Configure a new contract, with ESG Service as the provider and consumer VRF as vzAny consumer. This will allow the communication between the existing consumer EPG endpoints in EPG1, EPG2 and EPG3 with the newly created ESG Service.
4. Remove the EPG4's subnet 192.168.1.1/24 .
5. Once this subnet is removed, the traffic flow across VRF will stop.
6. Proceed with the creation of ESG route-leak policy in the VRF level between provider and consumer VRF. On the consumer VRF leak subnet 172.16.1.0/24 to the provider VRF. On the provider VRF leak subnet 192.168.1.0/24 to the consumer VRF.
7. Now migrate the provider EPG EPG4's endpoints into ESG Service by configuring the appropriate endpoint selectors.
8. Traffic flow across the VRFs will restore. This completes the migration of the provider EPG EPG4 to ESG.
9. Prepare the migration of individual consumer EPGs, namely EPG1, EPG2 and EPG3, by creating ESG IT and Sales that consume the same contract as the VRF vzAny.
Note: You can alternatively configure fine grain contracts for ESG IT and Sales respectively instead. In such a case, you must provide the new contract from ESG Service as well. These contracts (the fine grain or the one consumed by vzAny) need to have a scope larger than VRF such as global.
10. Now migrate the consumer endpoints to ESGs by configuring the selectors on ESG IT and Sales.
11. Optionally, the administration can remove the *Shared between VRFs* scope on the consumer BD subnet 172.16.1.1/24 if this BD is not participating in any other shared service. Also cleanup the EPG contracts association used for shared services.
12. Once the migration of all the consumer EPGs are completed, the administrator can remove the VRF vzAny consumer configuration.

Guidelines for EPG shared service and ESG shared service

- If the subnet prefix under the provider EPG is exactly the same as the route leaked via ESG route-leak policy in VRF level, the administrator needs to make sure that the EPG subnet (or shared-service contracts) is deleted prior to configuring the ESG route-leak policy in VRF level. This will prevent overlapping configuration for the same prefix.
- If the subnet prefix under the provider EPG is a subset of the route leaked via ESG route-leak policy in VRF level, the administrator can create the ESG route-leak policy prior to deleting the provider EPG subnet. However, the traffic within the provider EPG subnet will not follow ESG security until the provider EPG subnet is deleted.



Note This configuration is valid only when the provider BD has the same subnet as the route leaked via ESG route-leak policy or a subnet larger than that.

- If the subnet prefix under the provider EPG is a superset of the route leaked via ESG route-leak policy in VRF level, the administrator can create the ESG route-leak policy prior to deleting the provider EPG subnet. The traffic within the route leaked via ESG route-leak policy follows the ESG security without deleting the provider EPG subnet.

Configuring Endpoint Security Groups

Creating an Endpoint Security Group Using the GUI

Procedure

- Step 1** On the menu bar, choose **Tenants > All Tenants**. Select the applicable Tenant.
To create a tenant, refer to [Cisco APIC Basic Configuration Guide, Release 4.2\(x\)](#)
- Step 2** In the Navigation pane, choose **Application Profiles > Endpoint Security Groups**
To create an application profiles, refer to [Operating Cisco Application Centric Infrastructure](#).
- Step 3** Right click the **Endpoint Security Groups**, and select **create Endpoint Security Group**.
- Step 4** In **Create Endpoint Security Group** page, enter the following information:
 - a) **Name**: Enter the name of the ESG.
 - b) (Optional) **Description**: Enter the description of the ESG.
 - c) **VRF**: Enter the VRF that will be associated with the ESG.
 - d) **Selector**: Click + to add selector. **Create a Selector** page appears. Enter the following information:
 - **Match Expression**: The **key** field is set to **IP** as default; **operator** field is set to **equals** as default. Enter the IP address in the **value** field. For example, the user can enter a specific IP (/32, /128, or without a subnet mask) or a subnet match with any mask length.
 - (Optional) **Description**: Enter the description of the ESG: Enter the description.
 - Click **OK**.
- Step 5** (Optional) In case you need to block communication within the ESG, in the **Intra ESG Isolation** field, choose **Enforced**. The default is **Unenforced**.
- Step 6** (Optional) In case the ESG needs to be in the Preferred Group, in the **Preferred Group Member** field, choose **Include**. The default is **Exclude**.
When you select **Include**, ensure that the Preferred Group is enabled at the VRF level.
Refer to [Cisco APIC Basic Configuration Guide](#) for more information on Preferred Groups.

- Step 7** (Optional) In case you need to inherit contracts from another ESG, in the **ESG Contract Master** fields, navigate to the far right and click on the + sign. The user can choose ESGs from which contracts are inherited.
- Step 8** Click **Submit**.

Applying a Contract to an Endpoint Security Group Using the GUI

Procedure

- Step 1** On the menu bar, choose **Tenants > ALL TENANTS**. Select the applicable Tenant.
- Step 2** In the Navigation pane expand the **Tenant Name > Application Profiles > Endpoint Security Groups > ESG Name**.
- Step 3** Right click on **Contracts** and choose the action depending on how the contract is to be deployed, such as **Add Provided Contract** or **Add Consumed Contract**.
- Note** A contract that is consumed or provided by an application EPG cannot be used here for an ESG.
- Step 4** In the **Add Contract** dialog box, perform the following actions:
- Enter or select a **Contract Name**.
 - (Optional) Choose a **QOS policy**.
 - (Optional) Choose a **Label**.
- Step 5** Click **Submit**.

Creating Endpoint Security Groups and Applying a Contract Using the REST API

Procedure:

```
<polUni>
  <fvTenant name="t0">
    <fvAp name="ap0">
      <!-- ESG with the name ESG1 and Preferred Group as Exclude -->
      <fvESg name="ESG1" prefGrMemb="exclude">
        <!-- The ESG is associated to VRFA -->
        <fvRsScope tnFvCtxName="VRFA" />

        <!-- provided and consumed contracts -->
        <fvRsProv tnVzBrCPName="provided_contract1" />
        <fvRsCons tnVzBrCPName="consumed_contract2" />

        <!-- Endpoint Selectors for the ESG -->
        <fvEPSelector matchExpression="ip=='192.168.0.1/32'" />
        <fvEPSelector matchExpression="ip=='192.168.1.0/28'" />
        <fvEPSelector matchExpression="ip=='2001:23:45::0:0/64'" />
      </fvESg>
    </fvAp>
  </fvTenant>
</polUni>
```

Configuring Route Leaking of Internal Bridge Domain Subnets using the GUI

Use this procedure to configure route leaking of internal bridge domain subnets.

Before you begin

You must have created the tenant, VRF, bridge domain, and the subnet to be leaked.

Procedure

-
- Step 1** In the Navigation pane, navigate to the **Tenant name > Networking > VRFs > Inter- VRF Leaked Routes for ESG > EPG/BD Subnets**.
- Step 2** Right click on the **EPG/BD Subnets** and select **Configure EPG/BD Subnet to leak**.
- Step 3** In the **Configure EPG/BD Subnet to leak** dialog box, perform the following functions:
- IP**: Enter the bridge domain subnet and its mask to be leaked.
 - (Optional) **Description**: Enter the description of the EPG or bridge domain subnet.
 - (Optional) **Allow L3Out Advertisement**: Set to **True** when this subnet needs to be advertised by L3Outs on another VRF.
- Step 4** In the **Tenant and VRF destinations** field, navigate to the far right and click on the + sign.
- Step 5** In the **Create Tenant and VRF destination** dialog box, perform the following functions:
- Tenant and VRF**: Enter or select the tenant and VRF name.
 - (Optional) **Description**: Enter the description of the destination.
 - Allow L3Out Advertisement**: Set to **True** or **False**, when you need to change the permission per target VRF. By default, this option is set to **inherit** to retain the same configuration as **Allow L3Out Advertisement** in Step 3.
 - Click **OK**.
- Step 6** Click **Submit**.
-

Configuring Route Leaking of Internal Bridge Domain Subnets using the REST API

Before you begin:

You must have configured the BD subnet to be leaked or the BD subnet that includes the leaked subnet.

Procedure:

```
<polUni>
  <fvTenant name="t0">
    <fvCtx name="VRFA">
      <leakRoutes>
        <!--
          leak the BD subnet 192.168.1.0/24 with the Allow L3Out Advertisement
          False (i.e. scope private)
        -->
      </leakRoutes>
    </fvCtx>
  </fvTenant>
</polUni>
```

```

-->
<leakInternalSubnet ip="192.168.1.0/24" scope="private">
  <!--
    leak the BD subnet to Tenant t1 VRF VRFB with the
    Allow L3Out Advertisement configured in the parent
    scope (i.e. scope inherit)
  -->
  <leakTo ctxName="VRFB" tenantName="t1" scope="inherit" />
</leakInternalSubnet>
</leakRoutes>
</fvCtx>
</fvTenant>
</polUni>

```

Configuring Route Leaking of External Prefixes Using the GUI

Use this procedure to configure route leaking of external prefixes.

Before you begin

You must have configured an L3Out in the source VRF and the external prefixes are learned.

Procedure

-
- Step 1** In the Navigation pane, navigate to the **Tenant name > Networking > VRFs > Inter- VRF Leaked Routes for ESG > External Prefixes**.
- Step 2** Right click on the **External Prefixes** and select **Create Leaked External Prefix**.
- Step 3** In the **Create Leaked External Prefix** dialog box, perform the following functions:
- IP**: Enter prefix to be leaked.
 - (Optional) **Description**: Enter the description of the leaked external prefix.
 - (Optional) **Greater than or Equal (Prefix)**: Enter the minimum prefix length to be matched. This is equivalent to “ge” in IP prefix-lists in a normal router.
 - (Optional) **Less than or Equal (Prefix)**: Enter the maximum prefix length to be matched. This is equivalent to “le” in IP prefix-lists in a normal router.
- Step 4** In the **Tenant and VRF destinations** field, navigate to the far right and click on the + sign.
- Step 5** In the **Create Tenant and VRF destination** dialog box, perform the following functions:
- Tenant and VRF**: Enter or select the tenant and VRF name.
 - (Optional) **Description**: Enter the description of the destination.
 - Click **OK**.
- Step 6** Click **Submit**.
-

Configuring Route Leaking of External Prefixes Using the REST API

Before you begin:

You must have configured an L3Out in the source VRF “VRFA” and external prefixes are learned.

Procedure:

```

<polUni>
  <fvTenant name="t0">
    <fvCtx name="VRFA">
      <leakRoutes>
        <!--
          leak the external prefixes in the range of
          10.20.0.0/17 and 10.20.0.0/30
        -->
        <!--
          leak the external prefixes to Tenant t1 VRF VRFB -->
        <leakTo ctxName="VRFB" tenantName="t1" />
      </leakExternalPrefix>
    </leakRoutes>
  </fvCtx>
</fvTenant>
</polUni>

```

Applying Layer 4 to Layer 7 Services to an Endpoint Security Group Using the GUI

All the configurations provided for the deployment of a service graph with EPGs equally apply to the ESGs, the only change required is that instead of associating the contract to EPGs the contract is associated to ESGs. Use this procedure to apply a service graph template for a Layer 4 to Layer 7 service device in unmanaged mode to a contract used by endpoint security groups:

Before you begin

You must have created the following things:

- ESGs
- A service graph template

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 > Services > L4-L7 > Service Graph Templates**.
 - Step 4** In the Navigation pane, right-click on the **Service Graph Template Name** that you want to apply to the ESGs and choose **Apply L4-L7 Service Graph Template**.

The **Apply L4-L7 Service Graph Template To EPGs** dialog box appears. You will be associating a Layer 4 to Layer 7 service graph template to a contract between the endpoint security groups.

- Step 5** Configure a contract in the **Apply L4-L7 Service Graph Template To EPGs STEP 1> Contract** dialog box by entering the appropriate values:
 - Select **Endpoint Security Group** as the endpoint group type.
 - If you are configuring an intra-ESG contract, place a check in the **Configure an Intra-Endpoint Contract** check-box and choose the ESG from the **ESG / Network** drop-down list.
 - If you are using a normal contract instead of intra-ESG contract, select the ESG and network combination for consumer and provider.

- d) Create a new contract or choose an existing one by clicking the appropriate radio button in the **Contract Type** field. If you select **Create A New Contract** and want to configure the filters for it, remove the check from the **No Filter (Allow All Traffic)** check-box. Click + to add filter entries and **Update** when complete.

Step 6 Click **Next**.

The **STEP 2 > Graph** dialog appears.

Step 7 In the **your device name Information** section, configure the required fields represented with a red box.

Step 8 Click **Finish**.

You now have applied a service graph template to a contract used by ESGs.

Note To configure vzAny, select **AnyEPG** as provider and the ESG of interest as consumer, or vice versa in Step 5.c above.

To apply a service graph to a vzAny-to-vzAny contract vzAny-vzAny, select **Endpoint Policy Group (EPG)** as the endpoint group type and select **AnyEPG** as provider and consumer.

Applying Layer 4 to Layer 7 Services to Endpoint Security Groups Using the REST APIs

All the REST API's provided for the deployment of service graph with the EPGs equally apply to ESGs. However, the contract must be associated to the ESGs.

Please refer to [Layer 4 to Layer 7 REST API examples](#) for more information.