

Endpoint Security Groups

This chapter contains the following topic:

- About Endpoint Security Groups, on page 1
- Selectors, on page 5
- Contracts, on page 23
- ESG Shared Service (ESG VRF route leaking), on page 25
- Layer 4 to Layer 7 Services, on page 28
- Operational Tools, on page 28
- Guidelines and Limitations for Endpoint Security Groups, on page 29
- ESG Migration Strategy, on page 31
- Configuring Endpoint Security Groups, on page 34
- Configuring Route Leaking with Endpoint Security Groups, on page 42
- Configuring Layer 4 to Layer 7 with Endpoint Security Groups, on page 44

About Endpoint Security Groups

Endpoint Security Groups (ESGs) are a network security component in Cisco Application Centric Infrastructure (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 and used to define security zones within a bridge domain. This is because the EPGs define both forwarding and security segmentation at the same time. The direct relationship between the bridge domain and an EPG limits the possibility of an EPG to spanning more than one bridge domain. 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) instance instead of a bridge domain. This allows the definition of a security zone that is independent of the bridge domains (the fourth column of *Figure 1*, illustrates this point). Just as the EPGs divide a bridge domain into security zones, the ESGs divide the VRF instance 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 bridge domain subnet should be deployed on, and which subnets to be leaked to which VRF instance 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 bridge domain 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 bridge domains in the associated VRF instance, or a tag associated to endpoint MAC address. For details about these and other supported selector types, see About Selectors, on page 5.

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.





- 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 (13extInstP) 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 endpoint security groups (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 Application Centric Infrastructure (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 bridge domain 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. In some contexts, it is referred to as sclass.
- 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. In some contexts, it is referred to as sclass, S-Class, or source class.

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. Global pcTag numbers range from 16 to 16385.

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. Local pcTag numbers range from 16386 to 65535.

pcTag numbers from 1 to 15 are reserved for system internal use.

Selectors

About Selectors

Selectors are configured under each ESG with a variety of matching criteria to classify endpoints to the ESG. Unlike EPGs, which use VLANs to classify endpoints, ESGs can classify endpoints using much more flexible criteria. This concept is similar to micro segmentation EPG (or useg EPG); however, useg EPGs are still tied to one bridge domain while ESGs can contain endpoints across bridge domains.

The supported ESG selectors are:

- **Tag Selector**: Matches endpoints based on policy tags that are assigned to a variety of attributes such as MAC and IP addresses, virtual machine (VM) tags, virtual machine names [vm name], subnet tags, and static endpoint tags. ESG tag selectors can match only policy tags in the same tenant as the ESG. The tag selector is introduced in Cisco Application Policy Infrastructure Controller (APIC) release 5.2(1).
- **EPG Selector**: Matches all endpoints in a specific EPG, and the ESG will inherit all contracts configured under the EPG. This selector allows users to migrate security configurations from EPG to ESG seamlessly. ESGs can use EPG selectors only for EPGs in the same tenant and same VRF instance as the ESG. The EPG selector is introduced in Cisco APIC release 5.2(1).
- **IP Subnet Selector**: Matches endpoints based on the host IP address or IP subnet. Tag selectors provide the same capability by way of policy tags. The IP subnet selector is introduced in Cisco APIC release 5.0(1).
- Service EPG Selector: The service EPG selector is introduced in Cisco APIC release 5.2(4).

A service EPG is the EPG that Cisco Application Centric Infrastructure (ACI) creates automatically based on the connector of a device selection policy. In most deployments based on service graph redirect, there is no need to configure anything special to allow or deny traffic destined directly to the Layer 4 to Layer 7 device because Cisco ACI redirects traffic to the Layer 4 to Layer 7 device. If you need to send traffic directly to the Layer 4 to Layer 7 device EPG. The service EPG selector allows the mapping of a service EPG to a service ESG to give the administrator greater control about which ESG is allowed to send traffic to a Layer 4 to Layer 7 device deployed through service graph.

About Tag Selectors

A tag selector uses policy tags to classify endpoints to a given ESG. A policy tag consists of a key and a value, such as "key: owner, value: john." Policy tags can be assigned to a variety of user configurable objects, and ACI features can act upon those tags. Security classification using policy tags provides an easy and intuitive operation to add multiple endpoints to the security group (ESG). With policy tags and ESG tag selectors, you can classify your choice of multiple endpoints to an ESG without having to specify each endpoint individually.

ESG tag selectors match only policy tags in the same tenant as the ESG. This isolation ensures that each tenant manages its own resources, and it prevents any unintended policy tag match across tenants. Note, though, that if a user tenant is using a bridge domain or a VRF from the tenant 'common,' the user tenant may not have visibility of some configurations.

Although similar in configuration, policy tags (such as the user-definable tagTag) differ in purpose and usage from annotations (tagAnnotation). For details regarding the differences, see the "Alias, Annotations, and Tags" chapter of the *Cisco APIC System Management Configuration Guide*, *5.2(x)*.

Name	Description	Object
BD Subnet	Subnet under a bridge domain	fvSubnet
IP Endpoint Tags	Endpoint Tags Metadata for a host IP address of an endpoint fvEpIpTag	
MAC Endpoint Tags	Metadata for a MAC address of an endpoint	fvEpMacTag
VMM MAC Endpoint Tags	Metadata derived via VMM integration	fvEpVmmMacTagDef
Static Endpoint	Static endpoint	fvStCEp

ESG tag selectors can match policy tags assigned to the following objects.

The following sections describe the use of policy tags for each type of supported object.

Policy Tags for BD Subnets

By matching a policy tag assigned to a bridge domain subnet, a tag selector can classify all IP endpoints within the subnet to a given ESG. Although similar to an IP subnet selector, a policy tag and tag selector allow you to group multiple IP subnets, in addition to different types of parameters, such as specific MAC addresses.

You can also match a subset of a BD subnet by creating a smaller BD subnet with the **No Default SVI Gateway** option and assigning the policy tag to the smaller subnet. This option allows you to configure a subnet under a BD without deploying the corresponding SVI.

When configuring a tag selector matching a policy tag for BD subnets, consider the following guidelines:

- A tag selector cannot match a policy tag for a BD subnet in another tenant. For instance, if an ESG is in tenant "A" while a BD is configured in tenant "common", a tag selector in tenant "A" cannot match a policy tag for that BD. If subnet-based classification is required in such cases, use an IP subnet selector instead.
- Policy tags under an EPG subnet are not supported for ESG tag selectors. With ESG, there is no need to configure a subnet under an EPG. ESGs are intended to simplify the configuration by decoupling network and security configurations that were formerly combined under EPGs.
- A tag selector matching a policy tag for BD subnets classifies only IP addresses of endpoints to the ESG, not MAC addresses. For this reason, the layer 2 traffic limitation with IP-based selectors applies here. See Layer 2 Traffic Limitation with IP-based Selectors, on page 21 for details.

Policy Tags for IP Endpoint Tags

Because the objects representing endpoints (fvCEp, fvIp) are dynamically created and deleted based on the endpoint learning status on ACI switches, it is not practical to assign policy tags directly to such objects. For that reason, a new user-configurable object, an IP endpoint tag, is introduced with Cisco APIC Release 5.2(1) to represent an IP address of an endpoint. The IP endpoint tag object can be created and maintained even before the IP address is learned as an endpoint. Using this object, you can assign policy tags to an IP address of an endpoint at any given time.

An IP endpoint tag has a scope of VRF and represents the host IP address you configured in the given VRF. The tag is simply a metadata or descriptor of an IP address. Configuring an IP endpoint tag does not deploy an endpoint or the specified IP address. If you need to statically deploy an endpoint and its IP address before the endpoint is learned, configure a static endpoint.

When configuring a tag selector matching a policy tag for an IP endpoint tag, consider the following guidelines:

• A tag selector matching a policy tag for an IP endpoint tag classifies only IP addresses of endpoints to the ESG, not MAC addresses. For this reason, the layer 2 traffic limitation with IP-based selectors applies here. See Layer 2 Traffic Limitation with IP-based Selectors, on page 21 for details.

Policy Tags for MAC Endpoint Tags

Because the objects representing endpoints (fvCEp, fvIp) are dynamically created and deleted based on the endpoint learning status on ACI switches, it is not practical to assign policy tags directly to such objects. For that reason, a new user-configurable object, a MAC endpoint tag, is introduced with Cisco APIC Release 5.2(1) to represent a MAC address of an endpoint. The MAC endpoint tag object can be created and maintained even before the MAC address is learned as an endpoint. Using this object, you can assign policy tags to a MAC address of an endpoint at any given time.

A MAC endpoint tag has a scope of BD and represents the MAC address you configured in the given BD. If the MAC address is unique across BDs, you can specify the scope of BD as "any" ("*") and instead provide a VRF as its scope. The tag is simply a metadata or descriptor of a MAC address. Configuring a MAC endpoint tag does not deploy an endpoint or the specified MAC address. If you need to statically deploy an endpoint and its MAC address before the endpoint is learned, configure a static endpoint.

Policy Tags for VMM MAC Endpoint Tags

APIC automatically populates a read-only VMM MAC endpoint policy tag (fvEpVmmMacTagDef) based on information learned through VMM integration. APIC retrieves information about endpoints through VMM

integration and then maps that information to policy tags for each endpoint. Similar to a MAC endpoint tag object that you manually create, a VMM MAC endpoint tag object is simply a metadata or descriptor of a MAC address to maintain policy tags even when the corresponding endpoint is not learned in the data-plane yet. ESG tag selectors can use these policy tags to classify the endpoints to ESGs.

The following VMM information is supported by ESG tag selectors.

Integration Type	Source Information	Translated Policy Tag Format
VMware vSphere Distributed Switch (vDS)	VM name	Key:vmm::vmname Value: <i>VM name</i>
VMware vSphere Distributed Switch (vDS)	vSphere Tag " <i>Category: Tag Name</i> "	Key: <i>Category</i> Value: <i>Tag Name</i>

VMM MAC endpoint tags and the policy tags translated from the VM's name are automatically populated on APIC under **Tenant > Polices > Endpoint Tags > Endpoint MAC**. To enable this, you must enable **Allow Micro-Segmentation** when associating a VMM domain to EPGs. These tags are displayed with a suffix "(VMM)" to distinguish them from manually configured MAC endpoint tags. Translated policy tags other than a VM's name, such as a VMware tag, are not generated on VMM MAC endpoint tags until matched by an ESG tag selector. You must also enable **Tag Collection** under corresponding VMM domains. Each translated policy tag is assigned to the MAC address of an endpoint.

If a MAC endpoint tag is configured with the same MAC address in the same BD as the VMM MAC endpoint tag, only the policy tags from the MAC endpoint tag are used. In this case, the translated policy tags from the VMM MAC endpoint tags are ignored.

Policy Tags for Static Endpoint

By matching a policy tag assigned to a static endpoint that is configured under an EPG, a tag selector can classify the MAC address of the static endpoint to a given ESG. Policy tag support for static endpoints avoids the need for configuring a MAC endpoint tag for the same MAC address as the static endpoint. In fact, these two configurations are incompatible with each other. In other words:

- If policy tags are assigned to the static endpoint, a MAC endpoint tag with the same MAC address in the same BD cannot be configured.
- If a MAC endpoint tag is assigned to a MAC address, policy tags cannot be assigned to a static endpoint with the same MAC address in the same BD.

The static endpoint tag is supported only for static endpoints of type silent-host.

About EPG Selectors

An EPG selector matches an entire EPG to an ESG. Multiple EPGs can be matched to an ESG using EPG selectors, but only if the EPGs are in the same tenant and the same VRF as the ESG. The EPG selector is ideal for grouping multiple VLANs across bridge domains as a single security group (ESG) to simplify the configuration of contracts.

When an EPG is matched to an ESG by an EPG selector, all endpoints in the EPG belong to the ESG and all security configurations are now handled by the ESG.

EPG selectors have the following characteristics:

• Existing contracts under the EPG are inherited by the ESG.

- · The EPG cannot consume or provide new contracts
- Intra-EPG isolation is overwritten by intra-ESG isolation within the ESG.
- Preferred Group Membership in the EPG is overwritten by the ESG.

When an EPG is matched to an ESG via an EPG selector, intra-EPG/ESG isolation and Preferred Group Membership configuration under the EPG and ESG must be the same. After the match, the ESG settings overwrite the EPG settings.

The contract inheritance from EPG to ESG enables a seamless migration from the existing EPG security design to the new ESG security design. To simplify the configuration and to fully take advantage of ESG, we recommend that you complete the migration and do not retain EPG inherited contracts for EPG to ESG communication as a permanent configuration. When an ESG has contracts inherited by EPG selectors, APIC raises a fault as a warning and a reminder that the EPG to ESG migration has yet to be completed. See the "ESG Migration Strategy" section for details on migration using EPG selectors.

When an EPG is matched to an ESG by an EPG selector, the EPG's policy control tag (pcTag) is replaced by the ESG's pcTag. The pcTag replacement operation may cause a small transient traffic disruption for endpoints in the EPG. This is the same impact as other pcTag update events that occur with other features such as when configuring shared services (route leaking) with EPGs. Note that the pcTag is not specific to ESGs and is not related to policy tags (tagTag) used by tag selectors. The pcTag is an EPG/ESG identifier for applying contracts in the data-plane.

About IP Subnet Selectors

An IP subnet selector classifies endpoints to an ESG based on IP address. You can configure a host IP address to match a specific endpoint or you can configure a subnet to match multiple IP addresses within the subnet.

An IP endpoint tag selector classifies only IP addresses of endpoints to the ESG, not MAC addresses. For this reason, the layer 2 traffic limitation with IP-based selectors applies here. See *Layer 2 Traffic Limitation with IP-based Selectors* for details.

About Service EPG Selectors

Prior to release 5.2(4), you cannot create a contract with a service EPG created through a service graph. There are certain challenges that come with this limitation, such as:

- You can use the **Direct Connect** option to add a permit rule for the traffic from the service EPG to a consumer or provider EPG. However, an EPG that is not a consumer or provider EPG can't communicate with the service EPG unless you also configure a vzAny contract or a preferred group.
- As vzAny includes the service EPG, a vzAny-to-vzAny contract can permit traffic between the service EPG and other EPGs in the VRF. However, this also means that all of the other EPGs in the VRF are able to communicate with the service EPG, whereas you might want to limit only specific EPGs in the VRF to be able to communicate with the service EPGs.

Beginning with release 5.2(4), the service EPG selector for endpoint security groups (ESGs) is now available. This feature allows you to map a service EPG to an ESG and create a contract with that ESG. Using this feature, even if you have a vzAny-to-vzAny permit contract that is configured, you can add a deny contract between the service ESG and other ESGs to allow specific ESGs to communicate with the service ESG.

The following sections provide more information on example configurations with and without using service EPG selectors, and additional information on using service EPG selectors:

• Example Configurations Without Using Service EPG Selectors, on page 10

- Example Configurations Using Service EPG Selectors, on page 14
- Supported and Unsupported Locations for ESGs and Service EPGs, on page 16
- Guidelines and Limitations for Service EPG Selectors, on page 20

Example Configurations Without Using Service EPG Selectors

In order to enable the necessary configurations without using the service EPG selector option introduced in release 5.2(4), you could use the **Direct Connect** option. The following figure shows an example configuration where the **Direct Connect** option is in the default (disabled) setting.



100

The following figure shows an example where the **Direct Connect** option is enabled. A permit rule is added for the traffic from the service EPG to a consumer or provider EPG.

permit

504130

101



However, even with the **Direct Connect** option enabled, an EPG that is not a consumer or provider EPG doesn't have the permit rule with the service EPG, and you cannot add a contract manually.



One possible workaround to this restriction would be to configure a vzAny contract, where the service EPGs are part of the vzAny configuration, as shown in the following graphic.



However, one consideration with this workaround is that the EPG (class ID 300 in the previous example) can also communicate with other EPGs in the VRF.

A second possible workaround is to configure a preferred group, as shown in the following graphic.



However, one consideration with this second workaround is that other EPGs in the preferred group can communicate with each other without a contract. It could also consume more TCAM resources.

If neither of those workarounds provide a workable solution for your situation, you can use the service EPG selector option available beginning in release 5.2(4), as described in the following section.

Example Configurations Using Service EPG Selectors

Using the service EPG selector, available beginning with release 5.2(4), a service device connector representing the service EPG (LifCtx) can be mapped to an ESG, which allows you to add a contract with the ESG. In addition, zoning rules that involve service EPGs are inherited when you use the service EPG selector.

The following figure shows an example configuration using the service EPG selector.

L



Another way that you might use the service EPG selector feature would be to exclude the service device interface in a vzAny-to-vzAny permit contract. In this scenario, vzAny-to-vzAny is used to permit all traffic within a VRF, but you also want to prevent communication with the service device interface, as shown in the following figure.



Supported and Unsupported Locations for ESGs and Service EPGs

This section provides information on the supported and unsupported location for ESGs and service EPGs.

This section is relevant only for designs where the admin needs to allow or deny traffic directed to the Layer 4 to Layer 7 device from the ESGs. Traffic redirected to the Layer 4 to Layer 7 device does not belong to this category, and it is not subject to the restrictions described in this section. This is because, the destination IP address of the redirected traffic is an endpoint, and not the Layer 4 to Layer 7 device IP address.



Note

A service EPG is internally created in the tenant where the Layer 4 to Layer 7 device is defined.

• **Supported**: The Layer 4 to Layer 7 device and the service EPG-to-ESG mapping are defined in the same tenant.



• **Supported**: The Layer 4 to Layer 7 device is in the common tenant and the service EPG-to-ESG mapping is defined in a user tenant. In the example graphic below, the Layer 4 to Layer 7 device in the common tenant is exported to the user tenant tenant1, where the service graph is configured.



• **Unsupported**: The Layer 4 to Layer 7 device is in the common tenant and it's shared across multiple tenants, which means that the service EPG-to-ESG mapping is done in multiple user tenants.



Guidelines and Limitations for Service EPG Selectors

Following are the guidelines and limitations for the service EPG selector feature that is introduced in release 5.2(4):

- Although zoning rules that involve service EPGs are inherited, the class ID of the service EPG will be changed to a global class ID because it's mapped to an ESG that uses a global class ID. Traffic loss will occur when the class ID gets changed for the service EPG.
- All the service device connectors (LifCTx) in the same device using the same bridge domain must be mapped to the same ESG.

For example, assume that you have configured a one-arm mode firewall with PBR service graph, as shown in the following graphic.



In this example, the consumer and provider connectors are in the same bridge domain, using the same service EPG. In this case, both connectors must be mapped to the same ESG. If the connectors using the same service EPG are not mapped to the same ESG, a fault is raised and the service graph deployment will fail.

Note that you can reuse the service device interface for multiple service graph deployments.

- The service EPG and the ESG must be in the same VRF.
- NDO does not support ESGs at this time, so this feature is not supported with NDO.
- Support is available only for Layer 3 PBR with PBR destination in a bridge domain.
 - PBR destination in an L3Out is not supported (contracts can be manually configured with an L3Out EPG)
 - Layer 1/Layer 2 PBR is not supported (Layer 1/Layer 2 device interfaces are not supposed to communicate with servers directly)

Layer 2 Traffic Limitation with IP-based Selectors

With various classification methods in an endpoint security group (ESG), it is important to understand the difference in classification of IP addresses and MAC addresses. This difference is essentially the same as the microsegment (uSeg) EPG criteria.

When a switch routes a packet, the forwarding lookup is based on the IP address. When a switch switches a packet, the forwarding lookup is based on the MAC address even when the packet has an IP header. Similarly, when a switch routes a packet, the contract lookup is based on the IP address. When a switch switches a packet, the contract lookup is based on the MAC address even when the packet has an IP header. This behavior affects the contract application based on the ESG.

IP-based selectors (such as IP subnet selectors, tag selectors matching policy tags for bridge domain subnets or IP endpoint tag objects) classify only the IP addresses. Such classifications do not take effect for switching traffic. On the other hand, other selectors classify MAC addresses, and such classifications take effect for both switching and routing traffic. This means that a MAC-based selector applies also to an IP address associated to the MAC address, unless a separate IP-based selector overrides it. The following three scenarios demonstrate this behavior:

```
Scenario 1:
    MAC A is matched by a selector of ESG 1
    IP A is not matched by any ESG
  Result:
    Both MAC A and IP A are classified to ESG 1
Scenario 2:
    MAC A is matched by a selector of ESG 1
    \mbox{IP}\_\mbox{A} is matched by a selector of \mbox{ESG}\_\mbox{2}
  Result:
    MAC_A is classified to ESG 1
    IP A is classified to ESG 2
Scenario 3:
    MAC A is not matched by any ESG
    IP_A is matched by a selector of ESG_2
  Result:
    MAC A is not classified to any ESG, and still belongs to EPG A.
    IP A is classified to ESG 2
```

In these scenarios, endpoint EP_A is a member of EPG_A and does not initially belong to any ESG. EP_A's MAC address is MAC_A and its IP address is IP_A.

This behavior may cause switching traffic (layer 2 traffic) to bypass ESG contracts when you use IP-based selectors, even if the source and destination IP addresses of the traffic belong to different ESGs. To prevent this issue with IP-based selectors, use the proxy ARP feature in ACI so that all traffic is handled as routed traffic on ACI switches, even if the source and destination IP addresses are in the same subnet. There are three options for using proxy ARP for this purpose:

- Enable intra-EPG isolation along with proxy ARP on all of the EPGs that provide VLAN-to-interface binding for the ESG endpoints.
- Enable an intra-EPG contract with a permit-all filter, such as the common default contract, on all EPGs that provide VLAN-to-interface binding for the ESG endpoints. An intra-EPG contract enables proxy ARP automatically. The reason for a permit-all filter is to ensure that endpoints that are not classified to any ESGs can still communicate with each other within the same EPG. You can use any filters as a default behavior for endpoints that have yet to be classified to ESGs.

• Enable the **Allow Micro-Segmentation** option when associating a VMM domain to the EPGs that provide VLAN-to-interface binding for the ESG endpoints if VMM integration is used. This option automatically enables proxy ARP.

In the case of layer 2 traffic when endpoints in the same subnet (or VLAN) are classified to different ESGs, you may need private VLAN configuration regardless of the layer 2 traffic limitation with IP-based selectors. Private VLAN configuration may be needed when non-ACI switches exist between the endpoints and ACI switches. This is because non-ACI switches may switch the traffic before ACI switches can enforce contracts based on ESGs.



Note

Flood traffic that is not ARP requests, such as Layer 2 multicast, is dropped when it comes from the VLAN with the option to enable proxy-ARP.

Precedence of Selectors

When choosing selector types, consider whether the traffic will be switched or routed. The tables below show the order of precedence of selectors for each type of traffic.

Precedence Order	Selector	
1	Tag Selector (Endpoint MAC Tag)	
	Tag Selector (Static Endpoint)	
2	Tag Selector (Endpoint VMM MAC Tag)	
3	EPG Selector	

Table 1: Precedence Order for Switching Traffic

Table 2: Precedence Order for Routing Traffic

Precedence Order	Selector
1	Tag Selector (Endpoint IP Tag)
	IP Subnet Selector (host IP)
2	Tag Selector (BD Subnet)
	IP Subnet Selector (subnet)
3	Tag Selector (Endpoint MAC Tag)
	Tag Selector (Static Endpoint)
4	Tag Selector (Endpoint VMM MAC Tag)
5	EPG Selector

If an object is matched by multiple tag selectors via the same or different policy tags, the object is associated to the tag selector that matched first. Subsequent tag selectors are then ignored. If an object is matched by

multiple tag selectors when no tag selector had matched the object previously, no tag selectors take effect until the conflict match is resolved. A fault is raised under the ESG and under the object that is matched by multiple tag selectors.

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 \Leftrightarrow ESG
- **2.** ESG \Leftrightarrow L3Out EPG
- **3.** ESG \Leftrightarrow inband-EPG
- **4.** ESG \Leftrightarrow 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.



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.



Note This note is an implementation detail for advanced user information. If a service graph is attached to a contract between ESGs, the Cisco Application Policy Infrastructure Controller (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 a global pcTag.

For more information on Layer 4 to Layer 7 services deployment, see the 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.

- 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.

Guidelines and Limitations for Endpoint Security Groups

The following guidelines and limitations apply when using endpoint security groups (ESGs):

- 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.
- An ESG contract can be applied only for routed traffic when IP-based selectors are used. See details in Layer 2 Traffic Limitation with IP-based Selectors, on page 21.
- When using policy tags that are derived through VMM integrations, such as tags from VMware vCenter, you must have a full VMM integration. A read-only VMM integration is not sufficient.
- Taboo contracts are not supported with ESGs.
- ESGs cannot be specified as a source or destination for SPAN.
- Only the -EX and newer generation of leaf nodes are supported for ESG deployment.
- When classifying endpoints from the same VLAN into different ESGs, a private VLAN with an isolated
 port must be configured in the intermediate non-Cisco ACI switches (if any) to prevent those switches
 from switching traffic before the traffic reaches Cisco ACI. If the EPG is used for VMM VMware DVS
 integration, enable the Allow Micro-Segmentation option that automatically enables private VLAN on
 the VMware port group.



- **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 Cisco ACI leaf switches by using proxy ARP. 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 do not belong to ESGs, but are 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.

Beginning with the 5.2(3) release, the following features or configurations are supported:

- Inter-VRF service graphs between ESGs
- ESG shutdown
- Host-based routing/host route advertisement
- ESGs can be specified as a source or destination of the following features:
 - On Demand Atomic Counter
 - On Demand Latency Measurement
- The following features configured at the bridge domain or EPG level are supported with the specified limitations when endpoints in the bridge domain or EPG are classified to an ESG:
 - Endpoint reachability (static routes on bridge domain/EPG)
 - The MAC or IP address specified by this feature can be classified to an ESG only by using an EPG selector.
 - The static IP address (static route) and its next hop IP address must belong to the same ESG.
 - Anycast service
 - The MAC or IP address specified by this feature can be classified to an ESG only by using an EPG selector.
 - Microsoft NLB
 - The MAC or IP address specified by this feature can be classified to an ESG only by using an EPG selector.
 - When leaking the IP address specified by this feature to another VRF instance using VRF-level route leaking, the /32 or /128 route for the IP address must be explicitly leaked using route leaking for internal bridge domain subnets. For more information, see Configuring Route Leaking of Internal Bridge Domain Subnets using the GUI, on page 42.
 - First hop security (FHS)

- FHS is not supported on uSeg EPGs that match an ESG by using EPG selectors. If FHS is required for endpoints that need to move to an ESG from a uSeg EPG, classify those endpoints to an ESG by using other selectors, such as an IP subnet or tag selector, and remove matching criteria from the uSeg EPG. Then, configure FHS on the base EPG.
- When EPGs are matched to an ESG by using EPG selectors, the FHS binding table and corresponding endpoints are flushed. Traffic will not work until the binding table is refreshed using ARP, DHCP, and so on.

ESG Migration Strategy

Beginning with Cisco Application Policy Infrastructure Controller (APIC) release 5.2(1), EPG selectors allow endpoint security groups (ESGs) to inherit contracts from EPG, simplifying EPG-to-ESG migration. The contract inheritance with EPG selectors enables a seamless and flexible migration by allowing endpoints to keep communicating with other endpoints using inherited contracts even though the other endpoints are not yet migrated to ESGs.

In the following example, we will focus on the EPG to ESG migration of EPG A1 in the following figure. The current communication from EPG A1 is done through contract C1 with EPGs B1, B2, and B3.

Figure 7: Prepare to begin EPG-to-ESG migration



The first step is to create an ESG (ESG A1 in the following figure) and match EPG A1 to it using the EPG selector.





After EPG A1 has been matched to ESG A1, endpoints that belonged to EPG A1 now belong to ESG A1 and contract C1 provided by EPG A1 is inherited by ESG A1. All of the migrated endpoints can still communicate with EPGs B1, B2, and B3 even though these EPGs are not migrated to ESG yet. Remember that without the contract inheritance with EPG selectors, Cisco Application Centric Infrastructure (ACI) does not allow contracts between ESG and EPG. Note that when an ESG inherits contracts via EPG selectors, the original pcTags of the EPGs are replaced by the pcTag of the ESG. This operation may result in a small transient disruption of traffic for endpoints in the EPGs.

At this point, depending on your project schedule, instead of completing the migration of EPG A1, you could configure new contracts between ESG A1 and other ESGs or L3Out external EPGs. However, no more new contracts can be added to EPG A1 because all security configurations should be managed by the ESG. To keep the configuration simple and maintainable, we recommend that you complete the EPG to ESG migration at your earliest convenience. Until EPG A1 stops providing (or consuming) contracts, a fault F3602 is raised as a warning to make you aware of an incomplete migration.

To continue the migration, create ESGs for the EPGs on the other side of contract C1. In this example, EPG A1 is providing contract C1, so those EPGs (EPGs B1, B2, and B3) are consuming contract C1. Migrate these EPGs to new ESGs (ESGs B1, B2, and B3) using EPG selectors. In this example in the following figure, each EPG is mapped to an ESG.



Alternatively, you could combine multiple EPGs into one ESG. For example, you could create one ESG and then configure an EPG selector for both EPG B1 and B2 on the same ESG.

Next, create a new contract (C1' in the following figure) with the same filters as contract C1. Configure the new ESGs as provider and consumer. This is in preparation to stop providing contract C1 from EPG A1, which is the last step of EPG to ESG migration for EPG A1.



Figure 10: Create a new contract

Because contract C1 with the same filters was already inherited by all four ESGs (A1, B1, B2, and B3), the new contract configuration does not deploy any new rules in hardware, so no additional policy TCAM is consumed by creating the new contract.

ESG A1 now has contract C1' that allows the same communication as C1 with ESG B1, B2, and B3. At this point, we can stop providing contract C1 on EPG A1, allowing the ESG A1 to handle all security, as shown in the following figure.





Keep in mind that EPGs B1, B2, and B3 cannot stop consuming contract C1 yet because contract C1 is also provided by EPGs A2 and A3, which are not yet migrated to ESGs. After EPGs A2 and A3 are migrated to ESGs and are providing contract C1', all EPGs (A2, A3, B1, B2, and B3) can stop using contract C1 without traffic disruption.

To complete the migration of EPG to ESG, follow the same procedure for contract C2 and any other contracts on an EPG level.

Configuring Endpoint Security Groups

Creating an Endpoint Security Group Using the GUI

In Cisco APIC Release 5.2(1) and later releases, ESG selectors can be policy tags, EPGs, and IP subnets. In earlier releases, only IP subnets are supported.

Step 1	On the menu bar, choose Tenants and select the applicable Tenant.
Step 2	In the Navigation pane, choose <i>tenant_name</i> > Application Profiles > <i>application_profile_name</i> > Endpoint Security Groups
Step 3	Right click Endpoint Security Groups and select Create Endpoint Security Group.
Step 4	In the STEP 1 > Identity page of the Create Endpoint Security Group dialog box, enter the following information:
	a) Name: Enter a name for the ESG

b) (Optional) **Description**: Enter the description of the ESG.

- c) **VRF**: Enter the VRF that will be associated with the ESG.
- d) Click Next.

The STEP 2 > Selectors page of the Create Endpoint Security Group dialog box opens.

- **Note** In the following steps, you can create selectors based on policy tags, EPGs, and IP subnets. Alternatively, you can click **Next** and configure selectors later as described in Configuring Selectors and Tags, on page 36.
- **Step 5** In the **STEP 2** > **Selectors** page, click the + sign in the **Tag Selectors** bar if you want to use policy tags as an endpoint selector.

The Create a Tag Selector dialog box opens. Follow the procedure in Creating a Tag Selector, on page 36.

Step 6 In the **STEP 2** > **Selectors** page, click the + sign in the **EPG Selectors** bar if you want to specify an EPG as an endpoint selector.

The Create an EPG Selector dialog box opens. Follow the procedure in Creating an EPG Selector, on page 36.

Step 7 In the **STEP 2** > **Selectors** page, click the + sign in the **IP Subnet Selectors** bar if you want to specify an IP subnet as an endpoint selector.

The **Create an IP Subnet Selector** dialog box opens. Follow the procedure in Creating an IP Subnet Selector, on page 37.

Step 8 Click Next.

The **STEP 3 > Advanced (Optional)** page of the **Create Endpoint Security Group** dialog box opens.

- **Step 9** In the **STEP 3** > **Advanced (Optional)** page, you can configure the following options:
 - a) (Optional) To block communication within the ESG, choose **Enforced** in the **Intra ESG Isolation** field, The default is **Unenforced**.

Unenforced allows all endpoints within the same ESG to communicate freely. Alternatively, if you want to allow only a certain type of communications within the same ESG, you can use an intra-ESG contract instead. See Applying a Contract to an Endpoint Security Group Using the GUI, on page 39 for intra-ESG contract configuration.

 b) (Optional) To include the ESGs as preferred group members, choose Include in the Preferred Group Member field. The default is Exclude.

Before you select **Include**, ensure that the Preferred Group is enabled at the VRF level.

Refer to the Cisco APIC Basic Configuration Guide for more information on Preferred Groups.

c) (Optional) To inherit contracts from another ESG, click the + sign in the **ESG Contract Master** bar and choose ESGs from which to inherit contracts.

If you choose an ESG contract master, the ESG that you are creating will inherit all of the contracts of the chosen ESG. Add an ESG contract master if you want the new ESG to have the same security configuration as an existing ESG.

Step 10 Click Finish.

Configuring Selectors and Tags

Creating a Tag Selector

Use this procedure to create a tag selector for an endpoint security group (ESG).

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- **Step 2** In the Navigation pane, expand *tenant_name* > Application Profiles > *application_profile_name* > Endpoint Security Groups > *esg_name* > Selectors.
- Step 3 Right click Tag Selectors and select Create a Tag Selector.
- **Step 4** In the **Create a Tag Selector** dialog box, enter the following information:
 - a) **Tag Key**: Type a tag key or choose an existing tag key from the drop-down list.
 - b) Value Operator: Choose the condition for matching the tag value of an entity for inclusion in the ESG.

The operator choices are:

- Contains: Selects an entity whose tag value contains, but might not fully match, the Tag Value.
- Equals: Selects an entity whose tag value equals the Tag Value.
- Regex: Selects an entity whose tag value matches the regular expression entered in the Tag Value field.
- c) Tag Value: Type a value or a regular expression, or choose an existing value from the drop-down list.

When composing a regular expression, use the following guidelines:

- The allowed characters are: a-z A-Z 0-9 _ . , : ^ \$ [] () {} |+* -
- These characters are not allowed: / \?
- [0-9]+ matches any number (equivalent to \d+)
- a{0,1} matches zero or one of a (equivalent to ?)
- [0-9]{3} matches exactly a 3 digit number
- dev(1)|(2) matches value of dev1 or dev2
- d) Description: (Optional) A description of the selector.
- e) Click Submit.

Creating an EPG Selector

Use this procedure to create an EPG selector for an endpoint security group (ESG).

Step 1 On the menu bar, choose **Tenants** and select the applicable Tenant.

- **Step 2** In the Navigation pane, expand *tenant_name* > Application Profiles > *application_profile_name* > Endpoint Security Groups > *esg_name* > Selectors.
- Step 3 Right click EPG Selectors and select Create an EPG Selector.

Step 4 In the **Create an EPG Selector** dialog box, enter the following information:

- a) **EPGs in ESG VRF**: From the list of EPGs present in the VRF, check the checkboxes of the EPGs to be included in the ESG.
- b) **Description**: (Optional) A description of the selector.
- c) Click Submit.

Creating an IP Subnet Selector

Use this procedure to create an IP subnet selector for an endpoint security group (ESG).

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- Step 2 In the Navigation pane, expand *tenant_name* > Application Profiles > application_profile_name > Endpoint Security Groups > esg_name > Selectors.
- Step 3 Right click IP Subnet Selectors and select Create an IP Subnet Selector.
- **Step 4** In the **Create an IP Subnet Selector** dialog box, enter the following information:
 - a) **IP Subnet: key**: This field is set to **IP**.
 - b) IP Subnet: operator: This field is set to equals. The selector matches only an IP subnet that exactly matches the specified subnet.
 - c) **IP Subnet: value**: Type the IP subnet of the endpoints to be included in the ESG.

You can enter a specific IP (/32, /128, or without a subnet mask) or a subnet match with any mask length.

- d) **Description**: (Optional)
- e) Click Submit.

Creating a Service EPG Selector

Use this procedure to create a service EPG selector for an endpoint security group (ESG).

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- **Step 2** In the Navigation pane, expand *tenant_name* > Application Profiles > *application_profile_name* > Endpoint Security Groups > *esg_name* > Selectors.
- Step 3 Right click Service EPG Selectors and select Create a Service EPG Selector.
- **Step 4** In the **Create a Service EPG Selector** dialog box, enter the following information:
 - a) **Service EPG**: For the service EPG to be included in the ESG, choose from the list of provided service device connectors.

A service device connector (LifCtx), which represents a service EPG, can be mapped to an ESG. The list of service device connectors shown is derived from the connectors defined in the device selection policies, located here:

Tenants > tenant_name > Services > L4-L7 > Device Selection Policies

The service device connectors are presented in the following format:

consumer Of provider

TENANT_NAME/c-CONTRACT_NAME-g-GRAPH_NAME-n-NODE_NAME

For example:

consumer

PBR/c-web-to-app-g-FW-Graph-n-N1

- b) **Description**: (Optional) A description of the selector.
- c) Click Submit.

Creating an Endpoint MAC Tag

Use this procedure to add a policy tag to an endpoint MAC address. The tag can then be used by a tag selector to associate the endpoint MAC address to an endpoint security group (ESG).

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- **Step 2** In the Navigation pane, expand *tenant_name* > **Application Profiles** > *application_profile_name* > **Application EPGs** > *epg_name*.
- **Step 3** In the Work pane, choose the **Operational** > **Client Endpoints** tab.

The **Client Endpoints** table displays the MAC address of each available endpoint along with the IP address associated with it. If an address is already assigned policy tags, those policy tags are displayed in the **Policy Tags** column for the MAC or IP address.

Step 4 Right-click the row with the desired MAC address and select **Configure an Endpoint MAC Tag**.

If the MAC address does not appear in the table, it is not yet learned or visible through VMM integration. In this case, expand *tenant_name* > Policies > Endpoint Tags, right-click Endpoint MAC and select Create an Endpoint MAC Tag.

- **Step 5** In the **Create an Endpoint MAC Tag** dialog box, enter the following information:
 - **Note** If you selected a MAC address from the **Client Endpoints** table, the MAC address and BD fields are already populated.
 - a) Endpoint MAC Address: Enter the MAC address to which the tag will be added.
 - b) BD Name: Select an existing bridge domain or create a new bridge domain.

If you select *, the endpoint MAC tag represents the MAC address in any BDs in the given VRF. In this case, you are also asked to choose the VRF.

c) Annotations: (Optional) Click the + symbol to add an annotation key and value, then click the \checkmark symbol.

You can add more than one annotation.

d) **Policy Tags**: Click the + symbol to add a policy tag key and value, then click the \checkmark symbol.

You can add more than one policy tag.

e) Click Submit

Creating an Endpoint IP Tag

Use this procedure to add a policy tag to an endpoint IP address. The tag can then be used by a tag selector to associate the endpoint IP address to an endpoint security group (ESG).

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- **Step 2** In the Navigation pane, expand *tenant_name* > **Application Profiles** > *application_profile_name* > **Application EPGs** > *epg_name*.
- **Step 3** In the Work pane, choose the **Operational** > **Client Endpoints** tab.

The **Client Endpoints** table displays the MAC address of each available endpoint along with the IP address associated with it. If an address has already been assigned policy tags, those policy tags are displayed in the **Policy Tags** column for the MAC or IP address.

Step 4 Right-click the row with the desired IP address and select **Configure an Endpoint IP Tag**.

If the IP address does not appear in the table, it is not yet learned or visible through VMM integration. In this case, expand *tenant_name* > **Policies** > **Endpoint Tags**, right-click **Endpoint IP** and select **Create an Endpoint IP Tag**.

Step 5 In the **Create an Endpoint IP Tag** dialog box, enter the following information:

If you selected an endpoint from the Client Endpoints table, the IP address and VRF fields are already populated.

- a) **IP**: Enter the IP address to which the tag will be added.
- b) Annotations: (Optional) Click the + symbol to add an annotation key and value, then click the \checkmark symbol.

You can add more than one annotation.

- c) VRF Name: Choose or create the VRF that will contain the endpoint.
- d) **Policy Tags**: Click the + symbol to add a policy tag key and value, then click the \checkmark symbol.

You can add more than one policy tag.

e) Click Submit

Applying a Contract to an Endpoint Security Group Using the GUI

- **Step 1** On the menu bar, choose **Tenants** and select the applicable Tenant.
- **Step 2** In the Navigation pane, choose *tenant_name* > Application Profiles > *application_profile_name* > Endpoint Security Groups > *esg_name*.
- **Step 3** Right click on **Contracts** and choose the action depending on how the contract is to be deployed.

The options are:

- Add Provided Contract
- Add Consumed Contract
- Add Consumed Contract Interface
- Add Intra-ESG 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:
 - a) Enter or select a **Contract Name**.
 - b) (Optional) Choose a QOS policy.
 - c) (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" />
          <!-- Tag Selectors for the ESG -->
          <fvTagSelector matchKey="stage" valueOperator="equals" matchValue="production"/>
          <fvTagSelector matchKey="owner" valueOperator="contains" matchValue="teamA"/>
          <fvTagSelector matchKey=" vmm::vmname" valueOperator="regex"
matchValue="web_[0-9]+"/>
          <!-- EPG Selectors for the ESG -->
          <fvEPqSelector matchEpqDn="uni/tn-TK/ap-AP1/epq-EPG1-1"/>
          <fvEPgSelector matchEpgDn="uni/tn-TK/ap-AP1/epg-EPG1-2"/>
          <!-- IP Subnet 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>
```

Creating Tags and Selectors Using the REST API

Creating an EPG Selector

The EPG selector object (**fvEPgSelector**) matches the DN of a specific EPG.

<polUni>

```
<fvTenant name="ExampleCorp">
<fvAp name="AP">
<fvESg name="esgl">
<fvEPgSelector matchEpgDn="uni/tn-ExampleCorp/ap-app/epg-epgl"/>
<fvRsScope tnFvCtxName="dev"/>
</fvESg>
</fvAP>
</fvTenant>
</polUni>
```

The EPG selector can only match an EPG that belongs to the same tenant and VRF as the ESG.

Creating Tags and a Tag Selector

The tag selector object (**fvTagSelector**) matches tag objects (**tagTag**) discovered under the following objects:

- fvEpIpTag
- fvEpMacTag
- fvSubnet
- fvStCEp



Note

The tag selector object also matches tag objects under **fvEpVmmMacTagDef**. However, policy tags under this object are populated through VMM integration, and are not configurable.

This example shows the location of a **tagTag** object and the **fvTagSelector** object that will find and match the tag.

As an alternative to matching a tag exactly, a tag can be partially matched or matched using a regular expression using the **valueOperator** property of the **fvTagSelector**:

- If the **valueOperator** property is missing or if it is "equals," then only a **tagTag** whose value is an exact match is recognized.
- If the valueOperator property is "contains," a match is recognized if the tagTag's value field contains, but might not fully match, the fvTagSelector's matchValue field.

 If the valueOperator property is "regex," a match is recognized if the tagTag's value satisfies a regular expression contained in the fvTagSelector's matchValue field.

This example shows various matching conditions:

```
<fvTagSelector matchKey="name" matchValue="Blue"/>
<fvTagSelector matchKey="name" matchValue="Blue" valueOperator = "equals"/>
<fvTagSelector matchKey="name" matchValue="prod" valueOperator = "contains"/>
<fvTagSelector matchKey="name" matchValue="prod[0-4]" valueOperator = "regex"/>
```

Special Tag Selector for VMM Endpoints

Using a special key, the tag selector object (**fvTagSelector**) matches VMM endpoints by name. The special **matchKey** is "__vmm::vmname" and the **matchValue** is the name of the VM.

This example shows a tag selector that matches the VM named "vmName-Dev" using an exact match:

```
<polUni>
<fvTenant name="ExampleCorp">
<fvTenant name="AP">
<fvTeSg name="AP">
<fvTeSg name="esgl">
<fvTeSgelector matchKey="type" matchValue="dev"/>
<fvTegSelector matchKey="__vmm::vmname" matchValue="vmName-Dev"/>
<fvRsScope tnFvCtxName="testctx0"/>
</fvESg>
</fvAp>
</fvTenant>
</polUni>
```

Configuring Route Leaking with Endpoint Security Groups

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.

```
      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:
 - a) **IP**: Enter the bridge domain subnet and its mask to be leaked.
 - b) (Optional) **Description**: Enter the description of the EPG or bridge domain subnet.
 - c) (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:

- a) Tenant and VRF: Enter or select the tenant and VRF name.
- b) (Optional) **Description**: Enter the description of the destination.
- c) 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.
- d) 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)
        ___
        <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.

 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:
 - a) **IP**: Enter prefix to be leaked.
 - b) (Optional) Description: Enter the description of the leaked external prefix.
 - c) (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.
 - d) (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:
 - a) Tenant and VRF: Enter or select the tenant and VRF name.
 - b) (Optional) **Description**: Enter the description of the destination.
 - c) 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">
<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
    -->
    <leakExternalPrefix ip="10.20.0.0/16" ge="17" le="30">
        <leakTo ctxName="VRFB" tenantName="t1" />
        <leakTo ctxName="VRFB" tenantName="t1" />
        </leakExternalPrefix>
        </leakRoutes>
        </fvCtx>
        </fvCtx>
        </fvCtnant>
        </polUni>
```

Configuring Layer 4 to Layer 7 with Endpoint Security Groups

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

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:
 - a) Select Endpoint Security Group as the endpoint group type.
 - b) 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.
 - c) 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.