Endpoint Security Groups for ACI, Release 4.2.1

Table of Contents

New and changed information

This table provides an overview of the changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
Nexus Dashboard 4.2.1	There were no major changes from the previous release.	

Endpoint security groups

Endpoint Security Groups (ESGs) are a fundamental network security component in Cisco Application Centric Infrastructure (ACI). They enhance security policy management within the fabric and are integrated with Nexus Dashboard Search and Explore for comprehensive visibility and troubleshooting. While endpoint groups (EPGs) have been responsible for network security in Cisco ACI, they are limited to being associated with a single bridge domain and defining security zones within that domain. This limitation arises because EPGs handle both forwarding and security segmentation simultaneously. Consequently, an EPG cannot span multiple bridge domains. To overcome this limitation, the new ESG constructs are introduced.

ESG constructs

ESGs are introduced to overcome the limitations of EPGs. They are logical entities containing a collection of physical or virtual network endpoints. ESGs are associated with a single VRF (Virtual Routing and Forwarding) instance, not a bridge domain, enabling them to create a security zone separate from bridge domains. Similar to how EPGs divide a bridge domain into security zones, ESGs divide the VRF instance into security zones.

EPG policy and ESG functionality

The EPG policy combines forwarding and security logic. For instance, EPGs establish security zones based on VLAN and bind VLANs on leaf node interfaces. Additionally, contracts on EPGs ensure security and manage subnet deployment and VRF route leaking. In contrast, ESGs enforce security through contracts, while forwarding logic is managed by other components at the VRF level. The VLAN binding on leaf node interfaces remains at the EPG level.

ESG security constructs

An ESG is a security construct that has specific match criteria, known as ESG selectors, to determine endpoint membership. These selectors are based on attributes like IPv4/IPv6 addresses or endpoint MAC address tags that span across bridge domains in the associated VRF instance.

ESG selectors and endpoint classification

Selectors are configured under each ESG with a variety of matching criteria to classify endpoints. Unlike EPGs, which use VLANs to classify endpoints, ESGs can classify endpoints using more flexible criteria.

ESG supports these selectors in multi-fabric deployments:

- EPG Selector: All endpoints in the EPG will be assigned to the ESG.
- IP Subnet Selectors: IP subnet (network/mask). All endpoints belonging to this subnet will be assigned the ESG.
- Tag Selector: The tag selector includes key, operator, value, and description fields. Supported operators are equals, contains, and regex (regular expression). A set of tags can be configured in the ESG.

As of APIC 6.1(4), external EPG and external subnet selectors are supported for ACI fabrics:

- External EPG selectors are all subnets configured as external subnets for the external EPG.
- External IP Subnet Selectors: IP subnet (network/mask). All endpoints belonging to this subnet will be assigned the ESG. External IP Subnet Selectors are recommended over External EPG selectors for new deployments.

ESG contracts and communication rules

Contracts in Cisco ACI are similar to access control lists (ACLs). ESGs can only communicate with other ESGs based on the contract rules. Administrators use contracts to determine the types of traffic that can pass between ESGs, including allowed protocols and ports. An ESG can act as a provider, consumer, or both for a contract, and can consume multiple contracts at the same time. ESGs can also be part of a preferred group, allowing them to freely communicate with other ESGs in the same group.

Route leaking and security configuration in ESG

When an endpoint requires a service that is shared by another VRF, there are two requirements for communication to occur. The first requirement is routing reachability, and the second requirement is security permission. In an EPG, these two requirements are closely coupled in one set of configurations, such as the EPG subnet and contracts. However, in ESG, these two requirements are decoupled and require two different configurations, such as

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

With these two configurations completely decoupled, there is no need to configure a subnet or a subset of the subnet under the ESG, as is required for an EPG. The following sections explain how to configure route leaking for the bridge domain subnets and external prefixes learned from external routers. Once route leaking is configured, you can establish a contract between two ESGs, or between an ESG and L3Out EPG, to enable communication. It is important to use a contract with a scope larger than VRF, such as tenant or global.

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. This is performed by specifying a subnet to leak and the target VRF instance in the source VRF instance at the VRF level. 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.

Bridge domain subnet scopes and L3Out advertisement

VRF-level route leaking for an Endpoint Security Group (ESG) does not require the Advertised Externally and Shared between VRFs bridge domain subnet scopes. To advertise a leaked bridge domain subnet through L3Outs in the target VRF instance, set Allow L3Out Advertisement to True in the VRF-level route leaking configuration. The VRF-level route leaking configuration takes precedence over subnet scopes configured under a bridge domain when leaking the subnet to the target VRF instance. Bridge domain subnet scopes are still honored for other configurations, such as:

- Advertising the subnet from an External EPG in the same VRF instance.
- Route leaking to another VRF instance through a traditional configuration (EPG contracts).

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, differentiating it from the shared L3Out for EPGs. In order to leak routes learned from a L3Out for ESG communication, the administrator must configure the route leaking for external prefixes at the VRF level. This is done using an IP prefix-list style configuration. The user can configure a specific prefix or specify a range of prefixes using "le" (less

than or equal to) or " ge" (greater than or equal to), similar to 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. Due to the lack of this restriction, a leaked external prefix can specify a range to leak multiple prefixes with one configuration. In the configuration, it is also necessary to specify the target VRF.

To configure prefixes for an ESG shared L3Out, you have two options. This definition is required in addition to configuring route leaking in the VRF and applying a contract with L3Out.

- Configure an L3Out subnet with the External Subnets for the L3Out and Shared Security Import Subnet scopes.
- Alternatively, for shared L3Out use cases with ESGs, use External Subnet Selectors. For these configurations, enable the shared setting directly on the External Subnet Selector.

Implementing ESG

This section summarizes how the Cisco Application Policy Infrastructure Controller (APIC) programs leaf nodes when you configure endpoint security groups (ESGs):

- Each ESG is associated with a VRF instance, and the ESG selectors define which endpoints within the VRF instance belong to the ESG.
- Cisco Application Centric Infrastructure (ACI) applies the ESG configuration to all leaf nodes where the associated VRF instance is deployed.
- When an ESG is configured, the system programs all Bridge Domain subnet routes on all leaf switches where the VRF is present. This occurs because ESGs have a VRF-wide scope. In a VRF where only EPGs are present, the system programs Bridge Domain subnets on leaf switches based on contract configuration.
- When an ESG is configured with an External EPG or External Subnet selector, it must be deployed with **Immediate** option in **Deployment Immediacy**. For all other configurations, ESGs are deployed with the default **On Demand** option, and their associated contract rules are programmed only after an endpoint matching the ESG selectors is learned on the given leaf node.
- The contracts between ESGs are programmed as policy-cam rules in the leaf node TCAM, similar to EPGs.
- The Class ID used by the ESG is a global Policy Class Tag (pcTag), also known as class ID in some contexts.
- Unlike EPGs, contracts between ESGs only create security rules. ESGs are not used for network deployment such as subnet deployment or route leaking. Thus, for inter-VRF communication between ESGs, route-leaking configuration is required in addition to an inter-VRF contract.
- Even when ESGs are used for security enforcement instead of EPGs, EPGs are still required to configure VLAN bindings on leaf node interfaces.

After you implement ESGs, their communication and policy enforcement can be verified and monitored using the Search and Explore feature in Nexus Dashboard. This provides insights into traffic flow and security posture.

EPG to ESG migration

EPG selectors facilitate the migration of endpoint security groups (ESGs) by allowing them to inherit contracts from endpoint groups (EPGs). This inheritance streamlines the migration process, ensuring that endpoints can continue to communicate using the inherited contracts, even if the other endpoints have not yet been migrated to ESGs. For more information about ESG migration, see the **Cisco APIC Security Configuration Guide**.

Intra-VRF migration

1. Initial Setup: Identify and plan the migration of EPGs (EPG1 and EPG2) to corresponding ESGs (ESG1 and ESG2).
2. Configure ESG1: Deploy ESG1 with a selector for EPG1 in site1. Nexus Dashboard should reference the existing contract (C1), which defines communication between EPG1 and EPG2 as a consumer for ESG1, which will create a shadow ESG1 in site2.
3. Configure ESG2: Deploy ESG2 with a selector for EPG2 in site2. Nexus Dashboard should reference the existing contract (C1) as a provider for ESG2, which will create a shadow ESG2 in site1.
4. Contract Deployment: Deploy a cloned contract, C1-1, to both site1 and site2, configuring ESG1 as the consumer and ESG2 as the provider.
5. Remove Old Contract: Remove the original contract C1 from EPG1 and EPG2, ensuring the new contract C1-1 is operational between ESG1 and ESG2 for seamless migration.

Inter-VRF migration

1. Initial Setup: Identify consumer EPG1 in BD1 under VRF1 and provider EPG2 in BD2 under VRF2, and plan their migration to ESG1 in VRF1 and ESG2 in VRF2, respectively.
2. Configure ESG1: Deploy ESG1 with a selector for EPG1 in site1, ensuring contract C1 is inherited as a consumer and creating a shadow ESG1 in site2.
3. Configure ESG2: Deploy ESG2 with a selector for EPG2 in site2, inheriting contract C1 as a provider and creating a shadow ESG2 in site1.
4. Contract Deployment: Deploy cloned contract C1-1 to both site1 and site2, configuring ESG1 as the consumer and ESG2 as the provider.
5. Configure Route Leaking: Set up route leaking between VRF1 and VRF2, ensuring Subnet2, which is associated with EPG2 in BD2 in VRF2, is leaked to VRF1 to enable communication with EPG1 in BD1. Use the same subnet settings as configured under EPG2.
6. Transition to ESG or Route-Leak Path: Remove the contract C1 from EPG1 and EPG2, eliminating their shadows and directing traffic through the ESG/route-leak path.
7. Cleanup: Remove Subnet2 from EPG2, as it is no longer needed, completing the migration with all traffic routed through the ESG and route-leak configuration.

Guidelines and limitations for endpoint security groups

- In Nexus Dashboard 4.2.1, IP and tag selectors are never shadowed, while external IP and external EPG selectors are always shadowed. EPG selectors are shadowed only if the EPG itself is shadowed to the fabric.
- Objects are shadowed by ESGs if they are not already present in the fabric:
  - BDs: Shadowed if they share the same VRF as the ESG in fabrics where the ESG is present.
  - External EPGs: Shadowed whenever their External EPG selector is used by ESG, regardless of ESG stretch or shadow status, provided they are not already present in the fabric.
- When an Endpoint Security Group (ESG) references an Endpoint Group (EPG) that is not local to the ESG's fabric or not already shadowed to its site, the EPG's shadow is not created on the remote site for ESG's reference. This design allows ESG to function as forward references, enabling their presence on sites even without the direct presence of the referenced EPG. Consequently, the APIC may display a configuration error fault for the ESG, with a warning message. This fault does not impede traffic flow, as the underlying bridge domain is correctly shadowed across sites, ensuring connectivity.
- When validating subnet selector configurations between Nexus Dashboard and APIC using the API, the shared flag applies only to extSubnets and does not apply to ipSubnetSelectors.
- When you enable **Intra-ESG Isolation** from Nexus Dashboard, it may not work as intended. However, if you deploy intra-ESG isolation in APIC, it will modify both intra-ESG and intra-EPG isolation properties, resulting in inconsistent behavior. Although the template remains fully synchronized after deployment, running a 'Reconcile config drift' may report differences in intra-EPG property values.
- During EPG to ESG migration for inter-VRF route leaking, the provider EPG subnet must be configured under the bridge domain (BD) with the same subnet and flags as originally configured under the EPG.
- Endpoints are not automatically learned, so you may need to update them manually.
- Direct communication between EPGs and ESGs through contracts is not supported. Migration from EPGs to ESGs requires proper planning to avoid traffic loss.
- A provider ESG with an L3Out selector can participate in only one multi-fabric PBR contract. Combining ESGs with L3Out selectors and External EPGs in the same PBR contract is not supported.

List of supported scenarios

- ESG can be stretched within an intra-tenant and intra-VRF configuration.
- Inter-ESG can occur within an intra-tenant and intra-VRF setup.
- Inter-ESG operations can take place within an intra-tenant but across different VRFs.
- Inter-ESG can function across tenants and VRFs.
- External EPG can interact with ESG within an intra-tenant and intra-VRF environment.
- External EPG can connect to ESG within an intra-tenant but across different VRFs.

- External EPG can link to ESG across tenants and VRFs.
- Inter-ESG can be configured from a user tenant to a VRF in the common tenant, crossing VRFs.
- L3InstP-ESG can be set up from a user tenant to an External EPG and VRF in the common tenant, involving inter-VRF communication.

Create the ESG

Follow these steps to create the endpoint security group.

1. Log in to your Nexus Dashboard and open the Orchestration service.

   Manage > Orchestration > Tenant Templates

2. In the **Applications** page, choose the schema from the list and click the *schema*.

   The **Overview** page appears.

3. From the **Overview** drop-down list, choose a *template*.

   The **Template Properties** page appears.

4. From the **Create Object** drop-down list, choose **ESG**.

   The **Create ESG** page appears.

5. In the **Application Profile** drop-down list, choose a *profile*.
6. Click **Add**.

   The **untitled ESG** page appears.

7. Enter the name and description for the ESG template in the **Display Name** and **Description** fields.
8. From the **VRF** drop-down list, choose a *VRF object*.
9. Click **Ok**, to save the changes.

ESG is a security construct, so no specific fabric-level properties exist, unlike EPG.

Configure the ESG selectors

Follow these steps to configure the endpoint security group selector.

1. Navigate to **Template Properties**.
2. In the **Template Properties** page, from ESGs, click the *ESG object*.

   The *ESG object* page appears.

3. In the *ESG object* page, perform the following functions.
   - **Template properties**: Enter or edit the **Display name** and (optional) enter the **Description**.
   - **VRF**: Choose the *VRF* from the drop-down list.
   - **ESG Admin State**: (Optional) Choose from the available options. **Admin up** is chosen by default.
   - **Deployment Immediacy**: (Optional) Choose from the available options. **On demand** is chosen by default.
   - **Deployment Immediacy Operator**: (Optional) Displays fields based on the previous selection.

- **Contracts**: (Optional) Click **Add contract**, to choose the **Contract** and **Type** from the drop-down list.
- **Include in Preferred Group**: (Optional) Check this check box for preferred group.
- **Contracts via EPG Selectors**: (Optional) Displays the contracts with the EPG selectors.
- **Contracts via External EPG Selectors**: (Optional) Displays the contracts with the external EPG selectors.
- **Intra ESG Contracts**: (Optional) Choose from the available options in **Intra ESG Isolation**. **Unenforced** is chosen by default.
- **EPG selectors**: (Optional) Click **Add selector**, to synchronize the endpoint configurations of the EPG with the ESG.
- **External EPG selectors**: (Optional) Click **Add selector**, to add the external EPG.
- **Subnet Selectors**: (Optional) Click **Add selector**, to add the **IP Subnet** address and description.
- **External subnet selectors**: (Optional) Click **Add selector**, to add the IP subnet. Check the **Shared** check box to leak this prefix to other VRFs.
- **Tag selectors**: (Optional) Click **Add selector**, to attach the tag to your subnet or any MAC or IP address.
- **Annotations**: (Optional) Click **Create annotations**, in the **Create annotations** page, add the key and value.

4. Click **Ok**, to save the changes.

Configure the ESG contracts

Follow these steps to configure the endpoint security group contracts.

1. Navigate to the **Template Properties** page.
2. In the **Template Properties** page, from **Contacts**, click the *Contract object*.

   The *Contract object* page appears.

3. In the *Contract object* page, perform the following functions.
   - **Template properties**: Enter or edit the **Display name** and (optional) enter the **Description**.
   - **Annotations**: (Optional) Click **Create annotations**, in the **Create annotations** page, add the key and value.
   - **Scope**: Choose from the drop-down list.
   - **Apply both directions**: Enabled by default.
   - **Filter Chain**: Click edit icon, and perform the following in the **Update Filter Chain** page.
     a. **Directives**: Choose **Enable Policy Compression** from the drop-down list to optimize the use of policy CAM.
     b. **Action**: Click **Deny** to choose the **Priority** from the drop-down list. **Permit** is enabled by default.
   - **QoS Level**: (Optional) Choose the QoS level from the drop-down list.

- **Target DSCP**: (Optional) The drop-down list is based on the previous selection.
- **Service Chaining/Service Graph**: (Optional) **Service Chaining** is chosen by default. Click add + icon to add the Layer 4 or Layer 7 devices from the **Device Settings** page and click **Add**.
- Click **Ok**, to save the changes.

Configure the route leaking for internal subnets

Follow these steps to configure the route leaks for internal subnets.

1. Navigate to the **Template Properties** page.
2. In the **Template Properties** page, click the *VRF object*.
3. In the *new VRF* page, click **Add Leak Route** from **Template Properties**.
4. In the **Add Leak Routes** page perform the following.
   a. **Type**: **Internal Subnets** is selected by default.
   b. **IP**: Choose or type the IP address in the drop-down list.
   c. **Description**: (Optional) Enter the description for leak routes.
   d. **Allow L3Out Advertisement**: (Optional) Check the **Enable** check box to allow External EPG advertisement.
   e. **Tenant and VRF Destinations**: Click **Add Tenant** and **VRF Destinations** to include the following.
      - **Tenant Name**: Choose the destination tenant from the drop-down list.
      - **VRF Name**: Choose the destination VRF from the drop-down list.
      - **Allow L3Out Advertisement**: Click from the available options. **False** is chosen by default. The **Inherit** option uses the setting at the subnet level. At the destination, tenant controls advertising of the subnet out L3Outs in the destination VRF.
      - **Description**: (Optional) Enter the description.
5. Click **Save**, to save the changes.

You must configure the routing for inter-tenant communications between ESG1 to Tenant1 or ESG2 to Tenant2 or ESG that belongs to VRF1 and VRF2.

Configuring route leaking for external prefixes

Follow these steps to configure the route leaks for external prefixes.

1. Navigate to the **Template Properties** page.
2. In the **Template Properties** page, click the *VRF object*.
3. In the *new VRF* page, click **Add Leak Route** from **Template Properties**.
4. In the **Add Leak Routes** page perform the following.
   a. **Type**: Click **External Prefixes**.
   b. **IP**: Choose or type the IP address in the drop-down list.

c. **Description**: (Optional) Enter the description for leak routes.

d. **Greater Than or Equal (Prefix length)**: (Optional) Enter the minimum prefix length to be matched. This is equivalent to “ge” in IP prefix-lists in a normal router.

e. **Less Than or Equal (Prefix length)**: (Optional) Enter the maximum prefix length to be matched. This is equivalent to “le” in IP prefix-lists in a normal router.

f. **Tenant and VRF Destinations**: Click **Add Tenant and VRF Destinations** to include the **Tenant Name** and **VRF Name** from the respective drop-down lists. Click from the available options in **Allow External EPG Advertisement**.

5. Click **Save**, to save the changes.

PBR with ESG

Policy-Based Redirect (PBR) with multi-fabric Endpoint Security Groups (ESGs), introduced in Cisco ACI 6.1.4 and Nexus Dashboard 4.1.1, enhances traffic steering by redirecting network traffic to Layer 4-Layer 7 service devices (such as firewalls or load balancers) for ESG based consumers and providers across multiple ACI fabrics.

Before ACI 6.1(4) and Nexus Dashboard 4.2.1, multi-fabric PBR supported the following use cases.

Category	**EPG-to-EPG**	**EPG-to-External EPG, External EPG-to-EPG**	**vzAny-to-vzAny**	**VzAny-to-EPG**	**vzAny-to-External EPG**	**External EPG-to-External EPG**
Redirection	Site of the provider EPG	Site of the EPG	Both sites	Site of the provider EPG	Both sites	Both sites
Service node	1-node or 2-node, 1-arm or 2-arm	1-node or 2-node, 1-arm or 2-arm	1-node, 1-arm	1-node, 1-arm	1-node, 1-arm	1-node, 1-arm
VRF	Intra-VRF, Inter-VRF	Intra-VRF, Inter-VRF (External EPG must be the provider)	Intra-VRF	Intra-VRF	Intra-VRF	Intra-VRF and Inter-VRF
Symmetric PBR hash option	SIP-only, DIP-only or default	SIP-only, DIP-only or default	Default only	SIP-only, DIP-only or default	Default only	Default only
PBR threshold down action	Supported	Supported	Not supported	Supported	Not supported	Not supported

Starting with ACI 6.1(4) and Nexus Dashboard 4.1.1, ESG support is added for multi-fabric service integration, and it now supports following use cases:

- In the table, hash '#' indicates the new support for ESG starting from ACI release 6.1(4) and Nexus Dashboard 4.2.1

Category	**# ESG-to-ESG**	**# ESG-to-External EPG, External EPG-to-ESG**	**vzAny-to-vzAny**	**# VzAny-to-ESG**	**vzAny-to-External EPG**	**External EPG-to-External EPG**
Redirection	Site of the provider EPG	Site of the EPG	Both sites	Site of the provider EPG	Both sites	Both sites

Category	# ESG-to-ESG	# ESG-to-External EPG, External EPG-to-ESG	vzAny-to-vzAny	# VzAny-to-ESG	vzAny-to-External EPG	External EPG-to-External EPG
Service node	1-node or 2-node, 1-arm or 2-arm	1-node or 2-node, 1-arm or 2-arm	1-node, 1-arm	1-node or # **2-node**, 1-arm or # **2-arm**	1-node, 1-arm	1-node, 1-arm
VRF	Intra-VRF, Inter-VRF	Intra-VRF, Inter-VRF (# **External EPG must be the provider**)	Intra-VRF	**Intra-VRF**#	Intra-VRF	Intra-VRF and Inter-VRF
Symmetric PBR hash option	SIP-only, DIP-only or default	SIP-only, DIP-only or default	Default only	SIP-only, DIP-only or default	Default only	Default only
PBR threshold down action	Supported	Supported	Not supported	Supported	Not supported	Not supported

- For ESG and vzAny, you do not need to configure the EPG subnet due to conversational learning.
- ESG can include External EPGs and vzAny includes ESGs.

Traffic redirection overview of PBR with ESG

Intra-VRF, ESG-to-ESG/vzAny-to-ESG: In the intra-VRF scenario consumer and provider ESGs belong to the same VRF. The VRF stretches across multiple fabrics. Service devices must be deployed on both the consumer and provider fabric within the same VRF.

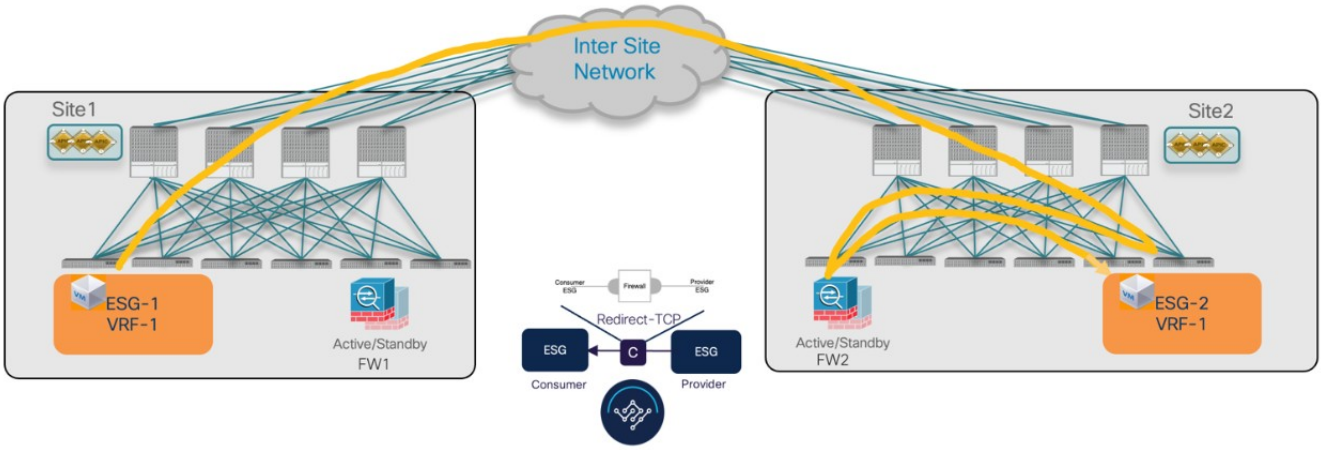

Intra-VRF, ESG-to-ESG/vzAny-to-ESG

Inter-VRF, ESG-to-ESG/vzAny-to-ESG: In the inter-VRF scenario consumer and provider ESGs belong to different VRFs. Service devices must be deployed on both the consumer and provider fabric and must belong to either the consumer or the provider VRF.

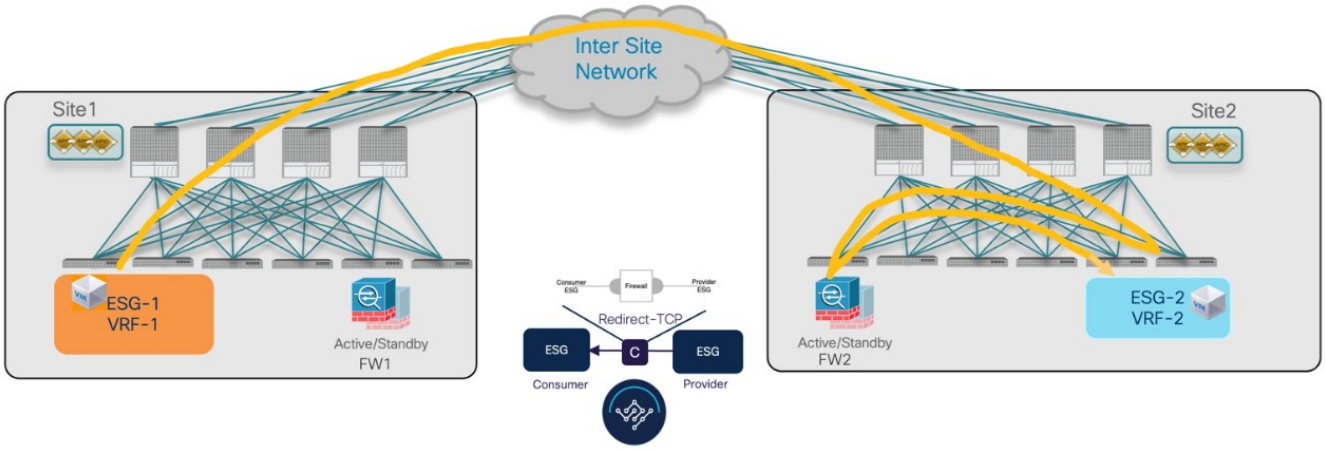

Inter-VRF, ESG-to-ESG/vzAny-to-ESG

For both intra-VRF and inter-VRF deployments, any communication between the consumer and provider is steered through firewalls provisioned on the provider's fabric because the redirection happens in the fabric of the provider ESG.

Even when the consumer initiates communication to the provider, the policy is enforced after reaching the provider fabric, and packets are steered through the provider fabric's firewalls. When the provider initiates communication to the consumer, if the policy is not enforced in the provider fabric (due to the consumer ESG's pcTag not resolving), that traffic traverses to the consumer fabric and redirects back to the provider fabric's firewalls. To conserve WAN bandwidth, the consumer leaf announces the consumer endpoint to the provider fabric to optimize this traffic path. This ensures packets destined for the consumer always derive the consumer ESG's pcTag, allowing packets to be directed to the provider fabric's firewalls directly from the provider or service leaf.

Conversational learning is used for ESG-to-ESG PBR, vzAny-to-ESG PBR, vzAny-to-EPG and vzAny-to-vzAny PBR to ensure efficient traffic steering, even for traffic across sites. For more information on conversational learning, see **Initial Consumer-to-Provider Traffic Flow and Conversational Learning**.

Policy rule expansion for inter-VRF and TCAM compression

A contract from a consumer vzAny (VRF1) to a provider ESG (VRF2) enables communication between endpoints in VRF1 and endpoints in the provider ESG in VRF2. If this contract just adds permit rules between any to this provider ESG in VRF2, it could enables communication from endpoints in VRF2 to endpoints in this provider ESG in VRF2 without an explicit contract. To prevent this unintended communication, Cisco ACI automatically performs policy TCAM rule expansion in the provider VRF.

To optimize TCAM usage, you can enable policy compression for contracts with PBR enabled. Support for this feature begins with ACI 6.1(4) and Nexus Dashboard 4.2.1. For more information on policy compression, see **Cisco ACI contracts**.

Provider fabric enforcement for symmetric firewalls

Before ACI 6.1(4) and Nexus Dashboard 4.2.1, External EPG-to-External EPG and vzAny-to-External EPG scenarios only supported a 1-node, 1-arm firewall service chain configuration. This model forwarded consumer-to-provider traffic through both consumer and provider fabric firewalls, resulting

in symmetrical enforcement.

Beginning with ACI 6.1(4) and Nexus Dashboard 4.2.1, these scenarios support two-node, two-arm, and inter-VRF service chain configurations, provided you migrate the External EPG to an ESG. This enhanced design steers traffic through the provider fabric's service chain, eliminating the previous symmetrical forwarding behavior.

Migrating external EPGs to ESGs for PBR

If you opt to migrate External EPGs to ESG and redirect traffic to the provider site in vzAny-to-External EPG or External EPG-to-External EPG scenarios with multi-fabric PBR enabled, follow these recommendations:

- Ensure all APICs are upgraded to release 6.1(4) or later before you use the L3Out selector for ESG in Orchestration.
- If a vzAny-to-vzAny contract shares a service device with a vzAny-to-External EPG or External EPG-to-External EPG, create a new logical service device or a new device interface (for example, by using a different VLAN).
- Identify all External EPGs who are the consumer or the provider for a contract with PBR before migration.
- For those External EPGs, identify the filters used in the contracts with PBR.
- Create a new contract with the identified filters from the existing contracts with PBR.
- Associate the service device with the new contract to establish the service chain.
- Create the ESGs in immediate mode and associate it with the new contract.
- Configure the L3Out selector for this ESG. This configuration might cause brief traffic interruption as the External EPG is migrated to the ESG.
- Delete contracts used in External EPG-to-External EPG or vzAny-to-External EPG redirection, or both.

L3Out selector limitations for PBR with ESG

- A provider ESG with an L3Out selector (external subnet, or External EPG) can participate in only one multi-fabric PBR contract.
- When an ESG with an L3Out selector is configured as the provider of a PBR contract, all other providers for that contract must match at least an L3Out selector.
- If an ESG with an L3Out selector is configured as a consumer or provider of a contract, you cannot also configure an External EPG as a consumer or provider of the same contract.
- If vzAny-to-vzAny with multi-fabric PBR is enabled, then an ESG with an L3Out selector cannot be the provider of the same contract. It must be the provider of a different contract, and the service device interface must be different for these contracts.
- If vzAny(VRF1)-to-vzAny(VRF1) with multi-fabric PBR is enabled, and if an ESG with an L3Out selector (VRF1) is configured as the provider of another contract, then you must configure vzAny(VRF1)-to-ESG with an L3Out selector (VRF1).

Configuration workflow for PBR with ESG

Policy-Based Redirect (PBR) setup (redirect policies, service graph, and contract action) is identical whether you use EPGs or ESGs. ESG only changes endpoint classification; the underlying PBR configuration steps do not change.

Before you begin:

- Enable Fabric-Aware Policy Enforcement on the VRF and Layer 3 multicast for conversational learning.
- Consider ESG with L3Out selector limits (for example, a provider ESG with an L3Out selector can participate in only one multi-fabric PBR contract).

1. Create the ESG and selectors (EPG, IP, tag, external EPG, or subnet).
2. Define the Layer 4-Layer 7 service device and interfaces (for example, firewall or load balancer) and enable redirection on the interface where traffic should be redirected.
3. Create a contract,and choose filters.
4. Configure the service chaining option within the contract by inserting the service devices in the desired sequence and configuring their corresponding device settings.
5. Associate the contract to ESGs (consumer or provider).
6. If you use ESG with L3Out selectors, ensure the provider ESG's L3Out selector constraints are met (for example, one PBR contract and no mixing with External EPG in the same contract).

General limitations and considerations for PBR with ESG

- ESG support is only available through the new service chain workflow.
- You must enable the **Fabric Aware** option on the VRF.
- You must enable Layer 3 multicast on the VRF for conversational learning. It is not enabled by default.
- The service device supports only Layer 3 device types.
- Multi-fabric and Autonomous templates are supported.
- Nexus Dashboard does not support intra-ESG contracts with PBR
- Nexus Dashboard does not support Service EPG selectors (ESGs that include a service EPG created by a service graph).

Copyright

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706

USA
https://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883