White paper
Cisco public

Cisco ACI Multi-Site and Service Node Integration for ACI version 6.1(4) or later

© 2026 Cisco and/or its affiliates. All rights reserved.

Contents

Introduction 3

Prerequisites 3

Executive summary 4

Service node integration with Cisco ACI Multi-Site architecture 7

Design options and considerations 7

Overview of the recommended design 11

Multi-Site service graph with PBR use cases 16

PBR to a firewall service 17

Load balancer with Source Network Address Translation (SNAT) 33

Load balancer without SNAT (use of PBR for the return traffic) 40

PBR to a firewall and a load balancer without SNAT (two nodes service graph) 47

Advanced design examples 54

Example 1: Insert firewall for most (but not all) inter-ESG traffic in a VRF. 54

Example 2: Use different firewalls for north-south and east-west traffic 57

Configuration examples 58

Overview 58

Create a tenant policy template for IP-SLA monitoring policy 58

Create a service device template 62

Application template 73

GUI and CLI output example for verification 81

Overview 81

Check that a service graph is deployed 82

Check if the traffic is redirected 83

FAQ 92

Conclusion 95

For more information 97

© 2026 Cisco and/or its affiliates. All rights reserved.

Introduction

This document describes the deployment considerations for integrating Layer-4 through Layer-7 (L4-L7) network services in a Cisco® Application Centric Infrastructure (Cisco ACI®) Multi-Site fabric using ACI release 6.1(4) or later. The document specifically focuses on stateful firewalls (FWs) and load balancers. The following use cases are considered:

- Layer-3 firewall design
- Layer-3 load-balancer design
- Layer-3 firewall and load-balancer service chain
- North-south and east-west service insertion design
- Independent clustered service nodes in each site

The assumption of this document is that ESGs (Endpoint Security Groups) are used for contract configuration, which means EPGs and external EPGs (L3Out EPGs) are migrated to ESGs that requires ACI release 6.1(4) or later with Cisco Nexus Dashboard Release 4.1(1) or later. If it's not applicable to you, please refer Cisco ACI Multi-Site and Service Node Integration for ACI versions 6.1(3) or earlier White Paper.

Prerequisites

To best understand the design presented in this document, you should have basic knowledge of the Cisco ACI Multi-Site solution, the deployment of L3Out connectivity between the Multi-Site fabric, ESG, contract, and the functionality of service graphs with Policy-Based Redirect (PBR).

Starting from release 3.0(1) of the Cisco ACI software, Cisco offers the Cisco ACI Multi-Site solution, which allows you to interconnect multiple Cisco ACI sites, or fabrics, under the control of the different Cisco Application Policy Infrastructure Controller (APIC) clusters. This solution provides an operationally simple way to interconnect and manage different Cisco ACI fabrics that may be either physically collocated or geographically dispersed. For more information about the Cisco Multi-Site architecture, please refer to the following white paper: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html. In this document, we use the terms "site" and "fabric' interchangeably to refer to a single "APIC domain," which could represent a single pod or a Cisco ACI Multi-Pod fabric deployment.

Cisco ACI offers the capability to insert L4-L7 services, such as firewalls, load balancers, and Intrusion Prevention Services (IPSs), using a feature called a service graph. For more information, please refer to the Cisco ACI service-graph-design white paper: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-2491213.html.

The service-graph functionality can then be enhanced by associating to it one or more Policy-Based Redirection (PBR) policies. For more detailed information on Cisco ACI contracts and PBR, please refer to the Cisco ACI PBR white paper: https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html.

For information about ESG and contract, see the Cisco APIC Security Configuration Guide, Release 6.2(x) and Cisco ACI Contract Guide White Paper.

© 2026 Cisco and/or its affiliates. All rights reserved.

Executive summary

As of Cisco ACI Release 6.2(1), the recommended option for integrating L4-L7 services into a Cisco ACI Multi-Site architecture calls for the deployment of independent service nodes in each site (Figure 1).

This is the logical consequence of the fact that the ACI Multi-Site architecture has been designed to interconnect separate ACI fabrics, at both the network fault domain and management levels. The focus in this document, therefore, will be exclusively on this deployment model.

The service-node High Availability options considered in this paper are the following ones:

- Active/standby service-node pair in each site
- Active/active cluster in each site
- Independent active service nodes in each site

It is possible to mix and match each HA option in the different fabrics that are part of the Multi-Site domain:

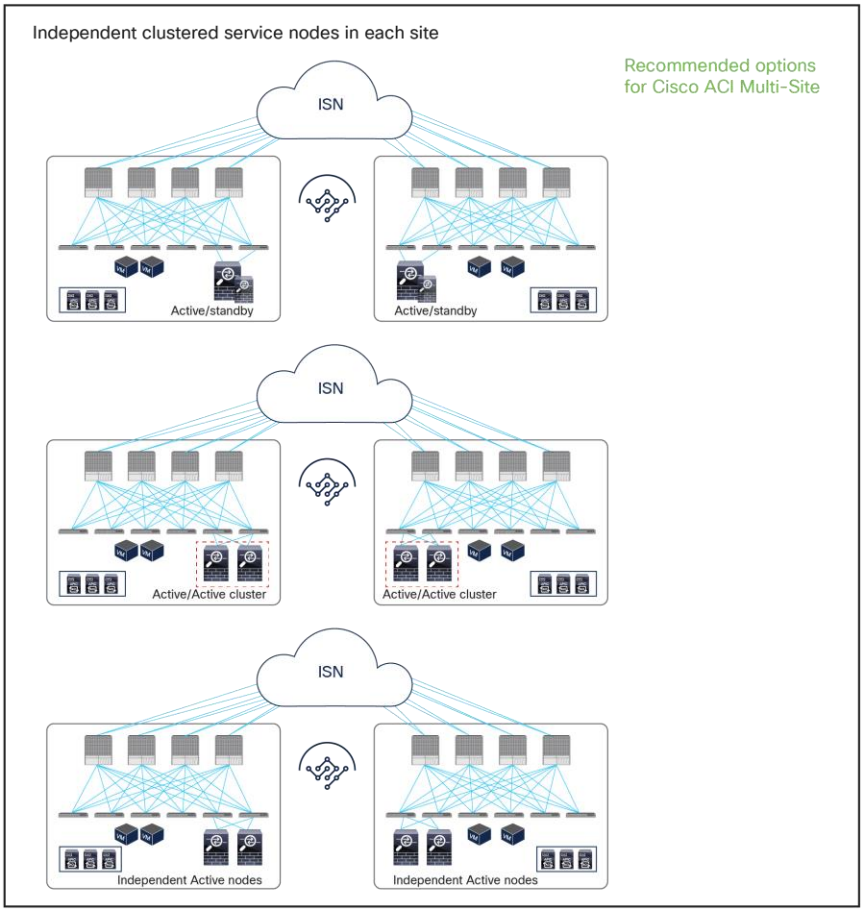


Figure 1.
Recommended network services deployment options with the Cisco ACI Multi-Site solution

Note: This white paper uses an active/standby service-node pair in each site mainly in the figures, though the other HA options shown in Figure 1 are also supported.

© 2026 Cisco and/or its affiliates. All rights reserved.

This model mandates that symmetric traffic flows through the service nodes be maintained, because the connection state is not synchronized between independent service nodes deployed in different sites. This requirement can be achieved with the following approaches:

- Use of host-route advertisement for north-south communication with stateful firewall nodes connected through L3Out: this allows connecting independent firewall nodes deployed between the border leaf nodes and the external WAN edge routers because inbound traffic is always optimally steered toward the site where the destination endpoint resides, whereas outbound traffic usually goes back through the same local L3Out connection. This approach, while fully supported and useful in many cases, relies on a more traditional routing design and only applies to north-south communication; this document therefore focuses on the second approach, described below, which leverages the advanced service insertion capabilities offered by an ACI network infrastructure.
- Use of service graph with PBR for both north-south and east-west communication: you can deploy service graph with Policy-Based Redirect (PBR) for both north-south and east-west security policy enforcement. This approach is the most flexible and recommended solution. It consists of defining a PBR policy in each site that specifies at least a local active service node, but it is also possible to deploy multiple active service nodes in the same site by leveraging symmetric PBR. The Cisco Nexus® 9000 Series Switches, used as leaf nodes, would then apply the PBR policy, selecting one of the available service nodes for the two directions of each given traffic flow (based on hashing). Different deployment models are supported for the service nodes: L3-routed mode (which has been supported from the beginning) but also L1/L2 inline/transparent mode as well.

Figure 2 and Figure 3 illustrate the other two models for the deployment of clustered service nodes between sites.

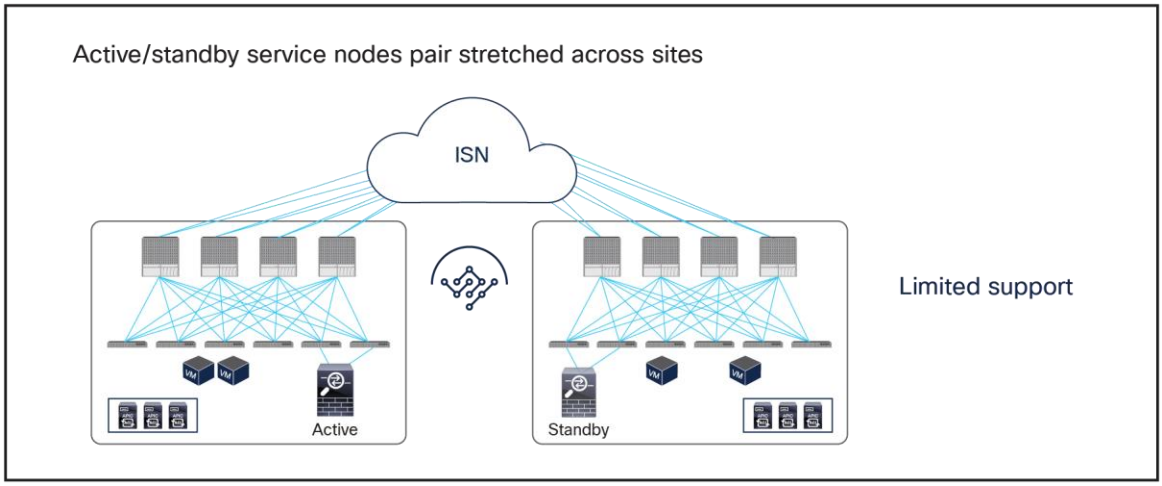


Figure 2.
Limited support network services deployment options with the Cisco ACI Multi-Site solution

© 2026 Cisco and/or its affiliates. All rights reserved.

- Active/standby service nodes pair stretched across sites: this model can be applied to both north-south and east-west traffic flows. This fail-safe model does not allow the creation of an asymmetric traffic path that could lead to communication drops. At the same time, because of the existence of a single active service node connected to the Multi-Site fabric, this option has certain traffic-path inefficiencies, because by design some traffic flows will hair-pin across the Inter-Site Network (ISN). Therefore, you should be sure to properly dimension the bandwidth available across sites and consider the possible latency impact on application components connected to separate sites. Also, this approach is only supported if ACI only performs Layer-2 forwarding (firewall as the default gateway for the endpoints or firewall in transparent mode) or when the active/standby firewall pair is connected to the fabrics via L3Out connections.

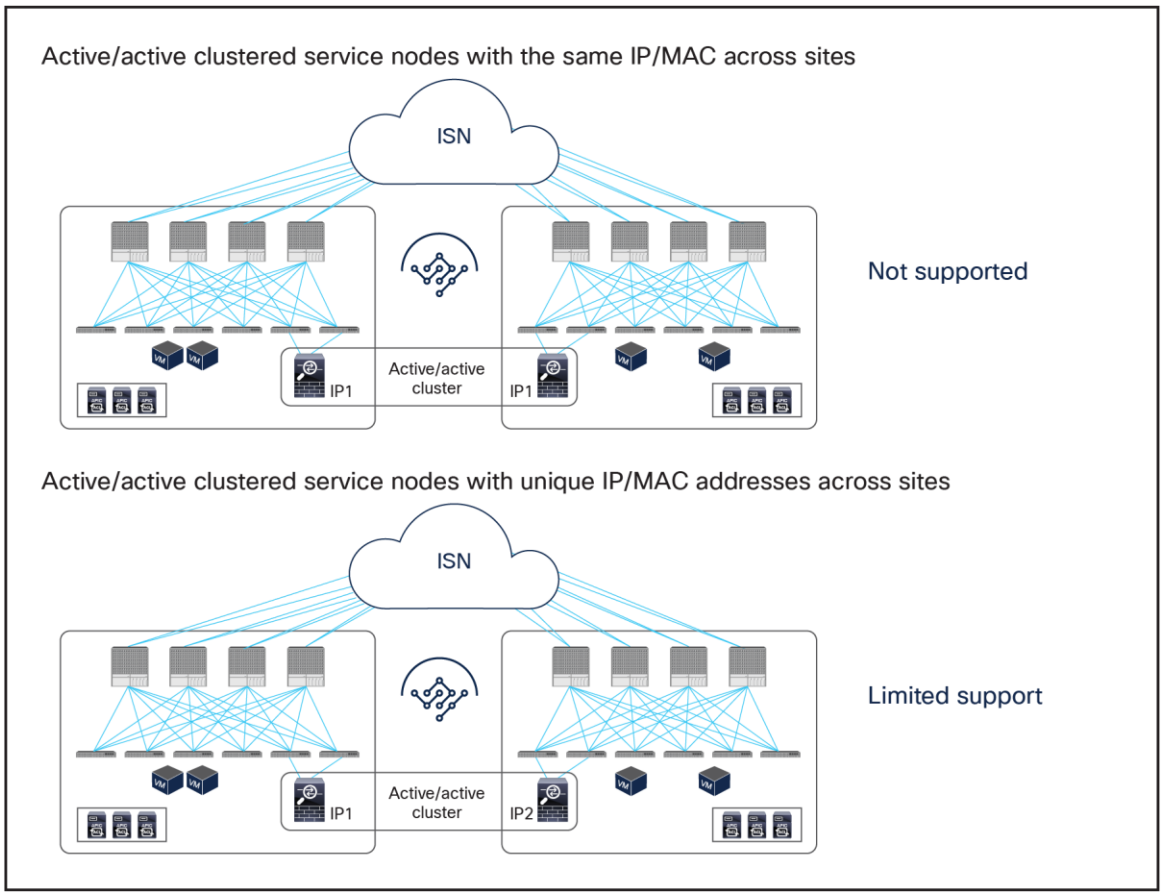


Figure 3.
Active-active service node cluster deployment options with the Cisco ACI Multi-Site solution

- Active/active clustered service nodes that use the same virtual MAC and virtual IP addresses stretched across sites: this model cannot be applied to a Multi-Site environment as of ACI Release 6.2(1), though an active/active firewall cluster can be stretched across pods in a Cisco ACI Multi-Pod environment. In the Cisco firewall implementation, this deployment model takes the name of Split Spanned EtherChannel cluster: all the firewall nodes that are part of the same cluster are seen as a logical distributed firewall reachable through a single virtual MAC and virtual IP. ACI Multi-Site does not currently support the capability of discovering the same virtual MAC and virtual IP pair across different sites. This situation, where the same endpoint is continuously learned in different locations, is considered to be an endpoint flapping scenario.

© 2026 Cisco and/or its affiliates. All rights reserved.

- Active/active clustered service nodes that use unique MAC and IP addresses across sites: this represents a second implementation option of an active/active firewall clustering, where each firewall node that is part of the same cluster owns its unique MAC and IP addresses. This option, supported by Cisco firewalls and some third-party implementations, can work today with an ACI Multi-Site architecture for some use cases[1]. Although this option works, the deployment models illustrated in Figure 1 are still primarily recommended, because clustering across sites potentially consumes more firewall resources, and connection sync across firewalls in different sites is not really required when deploying Cisco ACI PBR functionalities. The Cisco ACI fabric forwarding behavior is the same as the one used for independent clustered service nodes in each site, which is explained as part of the "Multi-Site service graph with PBR use cases" section in this paper.

Note: Cisco ACI Multi-Pod remains the recommended architectural approach for the deployment of active/standby service-node pairs across data centers and active/active clustered service nodes with the same virtual IP and virtual MAC addresses across data centers. For more information,
Please refer to the following white paper: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739571.html.

Service node integration with Cisco ACI Multi-Site architecture

Design options and considerations

Several deployment models are available for integrating network services in a Cisco ACI Multi-Site architecture. To determine the best options to choose, you should consider all the specific requirements and characteristics of the design:

- Service-node insertion use case
  - North-south service node (or perimeter service node), for controlling communications between the data center and the external Layer-3 network domain.
  - East-west service node, for applying policies for traffic flows within the data center and across sites. For the east-west enforcement, there are two cases to consider: in the first one, the service node (or its virtual context) is used to apply policies between ESGs that are part of the same Virtual Routing and Forwarding (VRF). The second scenario, very commonly deployed, is the one where a service node frontends each tenant/VRF, so as to be able to apply security policies to all of the inter-VRF traffic.
- Service-node appliance form factor
  - Physical appliance
  - Virtual appliance

[1] **Note:** For vzAny-to-vzAny PBR use cases that redirect traffic in both source and destination site, an active/active clustered firewalls across sites shouldn't be used regardless it's spanned etherchannel or individual modes because the cluster is not supposed to receive the same flow multiple times and it will potentially drop traffic depending on the cluster implementation.

© 2026 Cisco and/or its affiliates. All rights reserved.

- Service-node type
  - Inline (Layer 1 [L1]), transparent (Layer 2 [L2]), or routed (Layer 3 [L3]) mode firewall/IPS with PBR
  - Routed (Layer 3) mode load balancer with SNAT or without SNAT
- Service-node high-availability model
  - Active/standby HA pair in each site
  - Active/active cluster in each site
  - Independent active nodes in each site
- Connectivity to the external Layer-3 network domain
  - Traditional L3Outs deployed on the border leaf nodes

This document focuses on the service-node insertion use cases discussed below, describing north-south and east-west traffic flow examples and associated deployment considerations for each option in detail:

- Intra and inter-VRF traffic between ESGs.
- Intra and inter-VRF traffic between vzAny as the consumer and ESG(s) as the provider(s).
- vzAny-to-vzAny traffic in the same VRF.

Note: All the use cases listed above will be discussed in the section "Multi-Site service graph with PBR use cases". The internal endpoints and the L3Out connections with the external network domain can be either part of the same fabric or spread across different fabrics.

The figures below show the use cases covered in this document. Although the following examples and examples in this document mainly use separate ESGs for external subnets behind L3Outs and internal endpoints, both can co-exist in the same ESG.

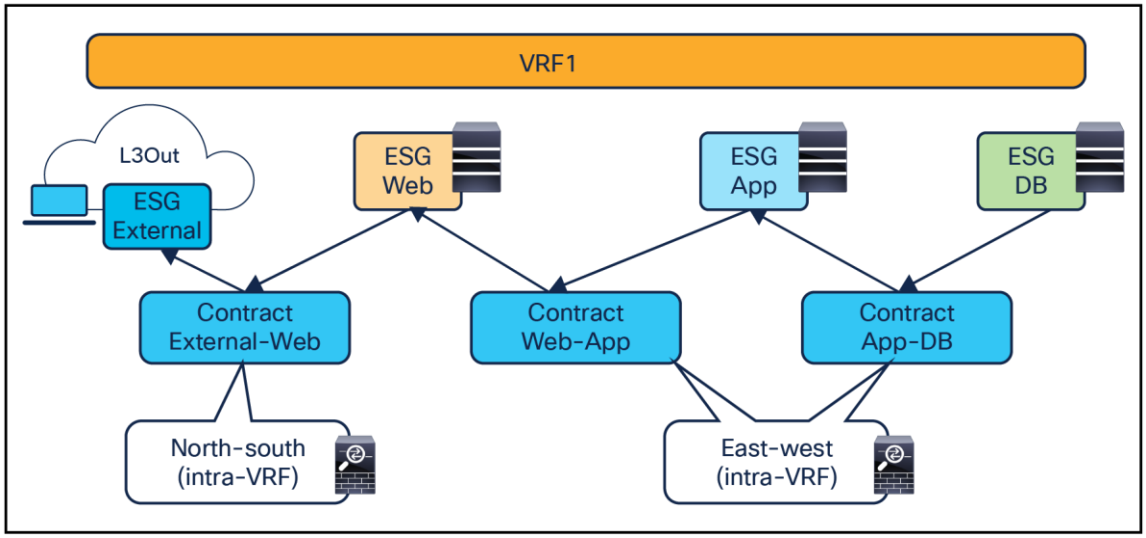


Figure 4.
ESG-to-ESG PBR for north-south and east-west service nodes (intra-VRF)

© 2026 Cisco and/or its affiliates. All rights reserved.

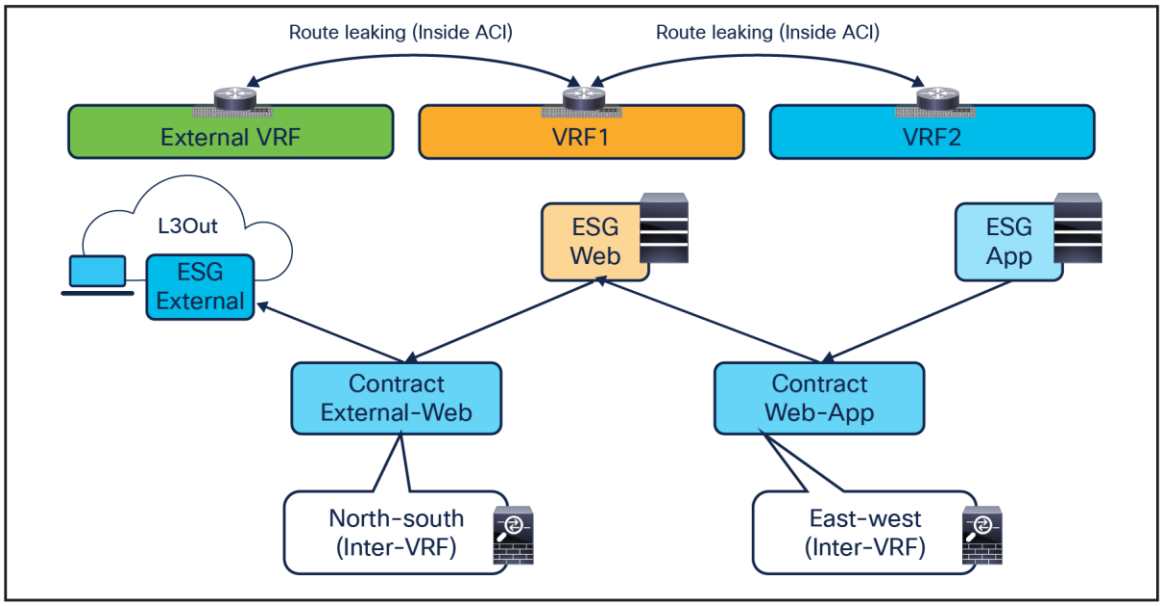


Figure 5.
ESG-to-ESG PBR for north-south and east-west service nodes (inter-VRF)

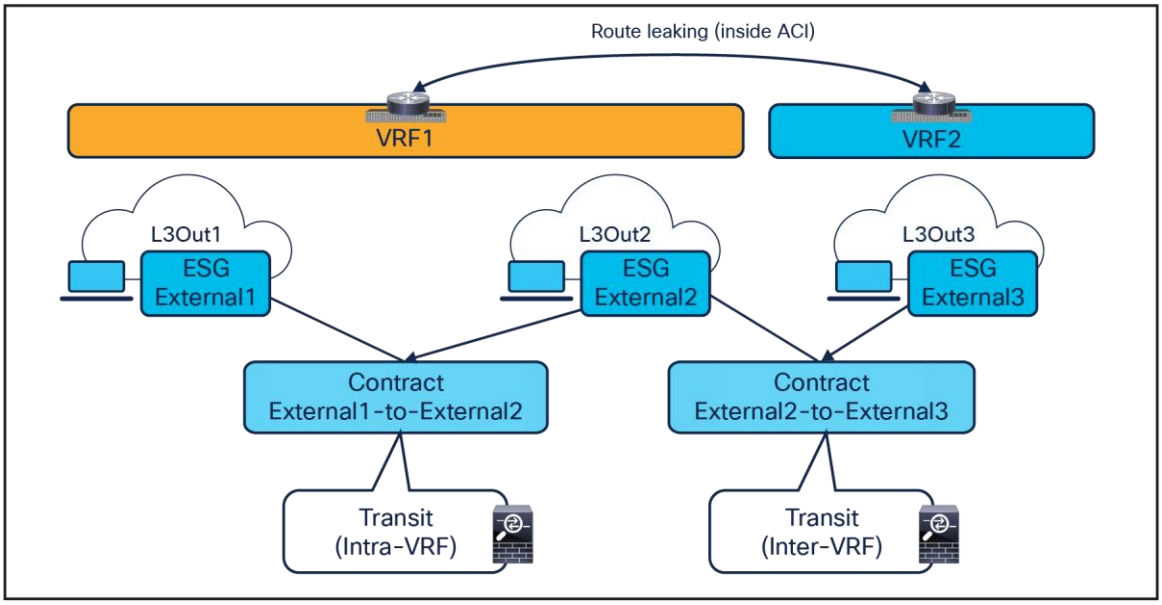


Figure 6.
ESG-to-ESG for transit service nodes (intra-VRF and inter-VRF)

© 2026 Cisco and/or its affiliates. All rights reserved.

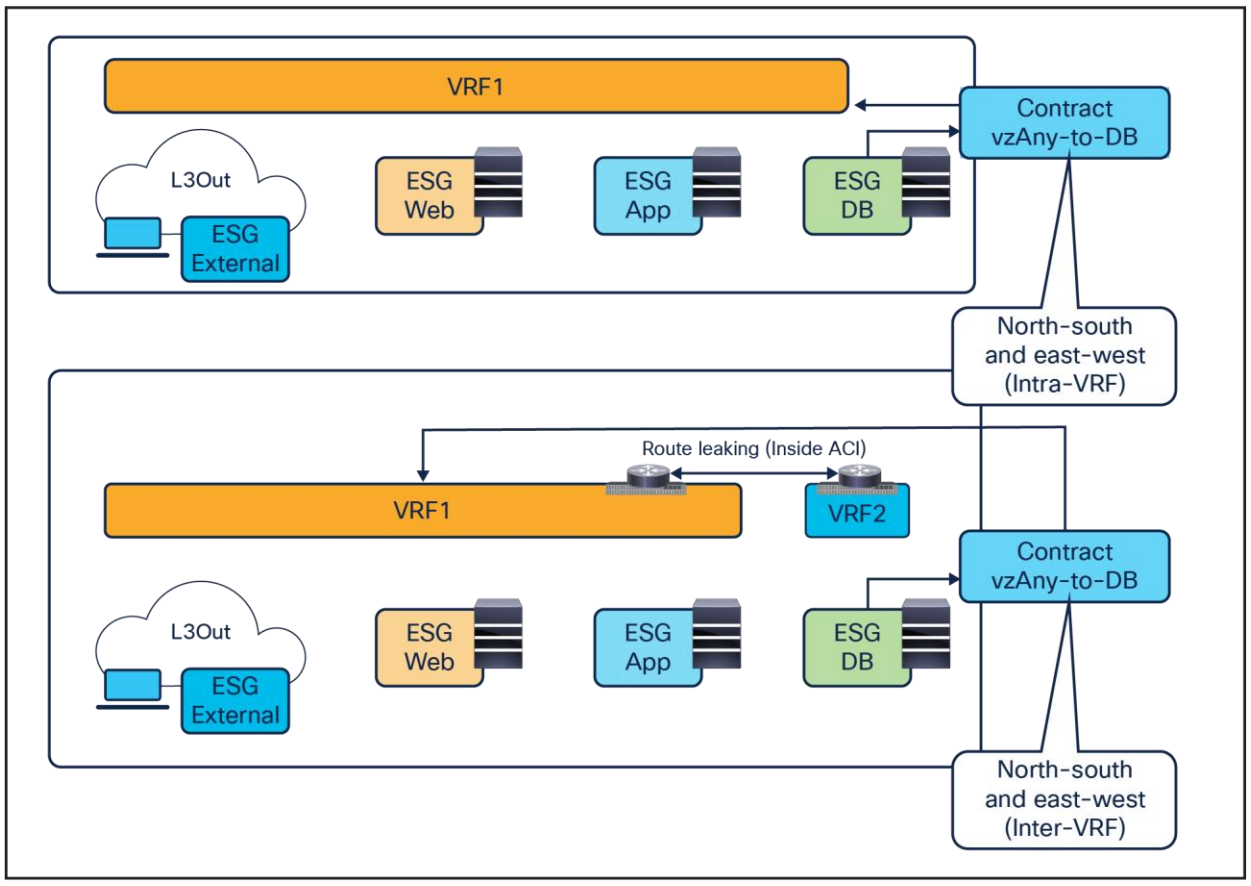


Figure 7.
vzAny-to-ESG for north-south and east-west service nodes (intra-VRF and inter-VRF)

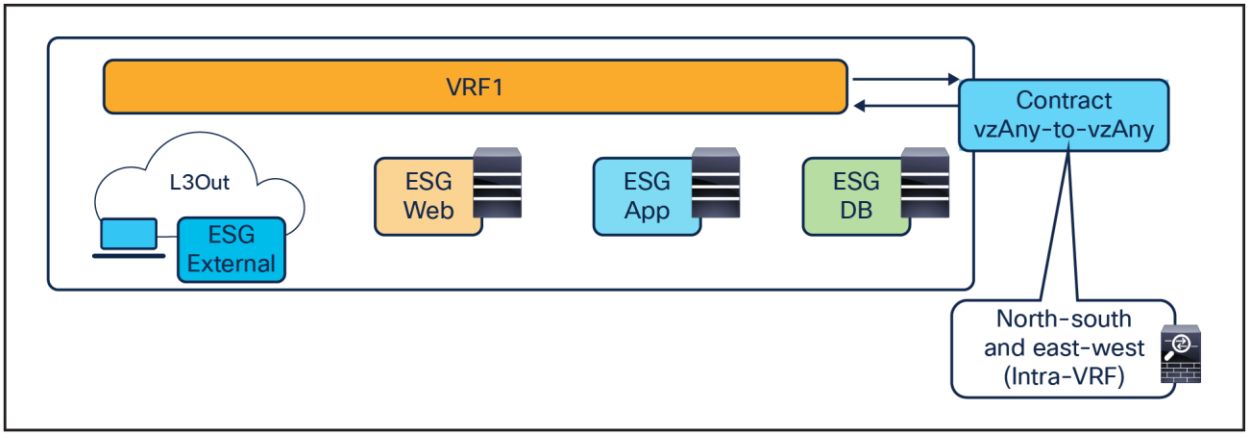


Figure 8.
vzAny-to-vzAny for north-south and east-west service nodes (intra-VRF)

© 2026 Cisco and/or its affiliates. All rights reserved.

Overview of the recommended design

Figure 9 shows a high-level view of the topology representing the recommended deployment option with independent clustered service nodes in each site. We are going to use routed mode firewall, routed mode load balancer, and traditional L3Outs as examples in this document.

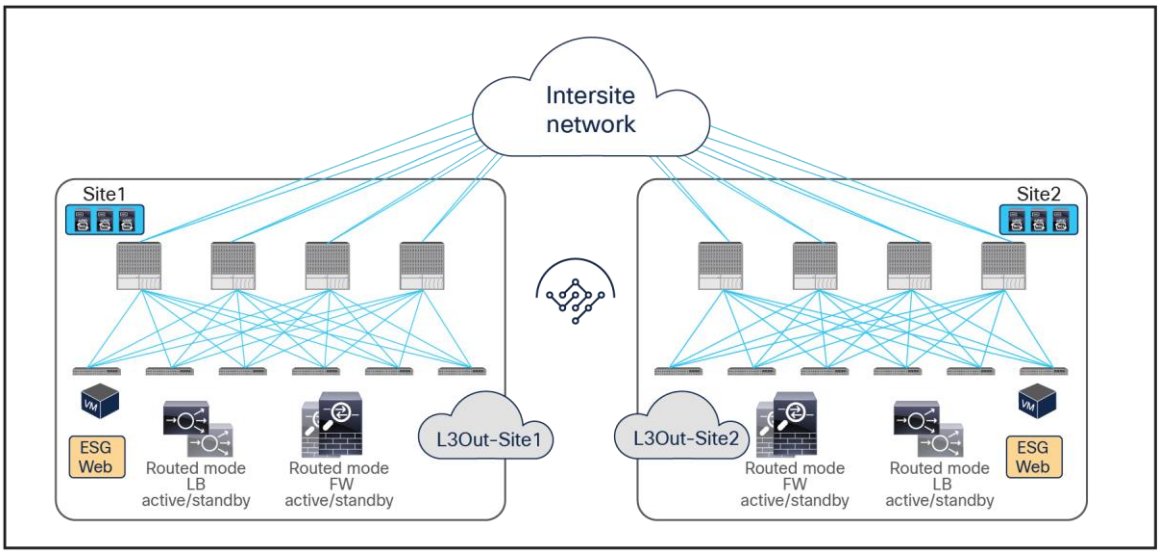


Figure 9.
Independent clustered service nodes in each site

The deployment of independent service nodes across sites raises an operational concern about how to maintain policy configuration consistency across them. In the specific example of Cisco® firewalls, some options are available:

- **Cisco Security Manager for Adaptive Security Appliances (ASAs):**
  For more information, see https://www.cisco.com/c/en/us/products/security/security-manager/index.html.
- **Cisco Firepower® Management Center (FMC) for Cisco Firepower Next-Generation Firewall (NGFW) devices:**
  For more information, see https://www.cisco.com/c/en/us/products/security/firesight-management-center/index.html.

When planning for the deployment of this model, it is important to keep in mind a few important design requirements:

- The policy to be applied (the 'intent") is defined directly on Cisco Nexus Dashboard (ND), for example, specify that any communication between the Web ESG and App ESG must be sent through a service node (or a chain of service nodes). Each specific service node is then mapped, at the site level, to the specific physical or virtual service appliances locally deployed.
- In the current implementation, the PBR policy applied on a leaf switch can only redirect traffic to a service node deployed in the local site. As a consequence, it becomes paramount to improve the resiliency of the local service nodes. This can be achieved with the different options shown in Figure 9.

© 2026 Cisco and/or its affiliates. All rights reserved.

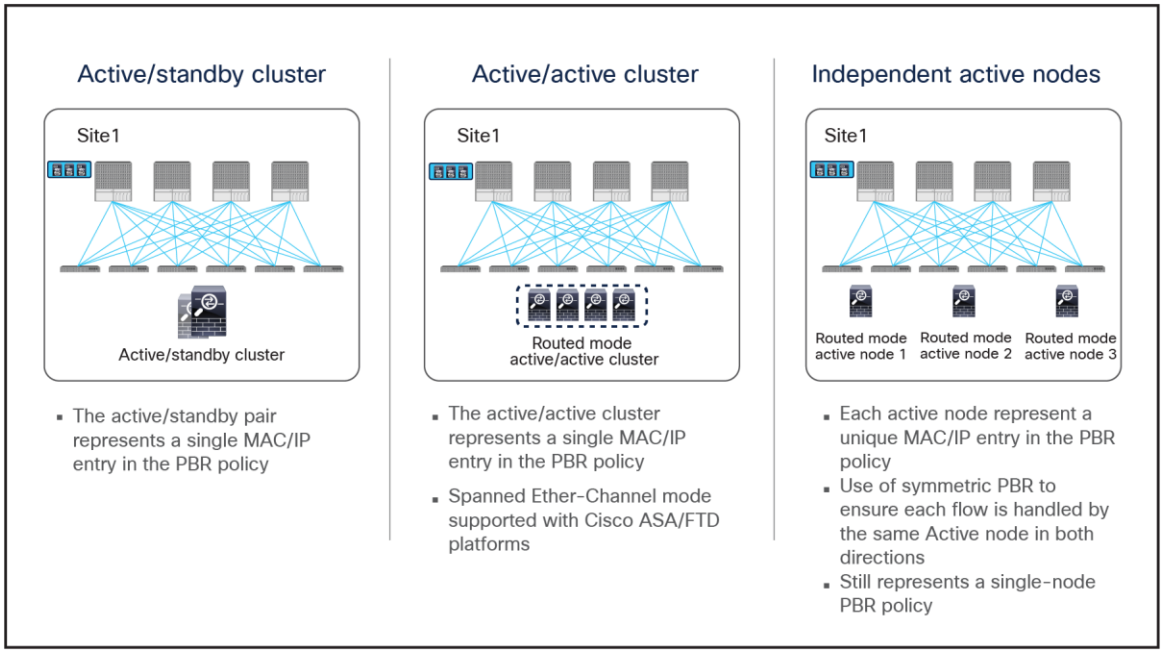


Figure 10.
Deployment options to increase the resiliency of service nodes in a single site

The first two models are obvious, as they both ensure that the service node is seen by the fabric as a single entity, so the PBR policy would only contain a single MAC/IP pair. With the third option, multiple MAC/IP pairs are instead specified in the same PBR policy, so that a given traffic flow can be redirected to a service node. Use of symmetric PBR ensures that both the incoming and return directions of the same flow are steered through the same service node.

The definition and behavior of an active/active cluster differ depending on the vendors. In the case of Cisco ASA and a Cisco Firepower Threat Defense (FTD) active/active cluster, service nodes in the same cluster can use the same MAC and IP addresses, which is the second option in the figure above, whereas service nodes in the same Palo Alto Networks active/active HA use unique IPs, which is an enhanced version of the third option.

Note: Cisco ASA can also support an active/active cluster where each firewall node owns a unique MAC/IP address pair.

As previously mentioned, service graph with PBR can be used to handle service-node insertion for north-south, east-west, and transit traffic flows, as illustrated in Figure 11.

© 2026 Cisco and/or its affiliates. All rights reserved.

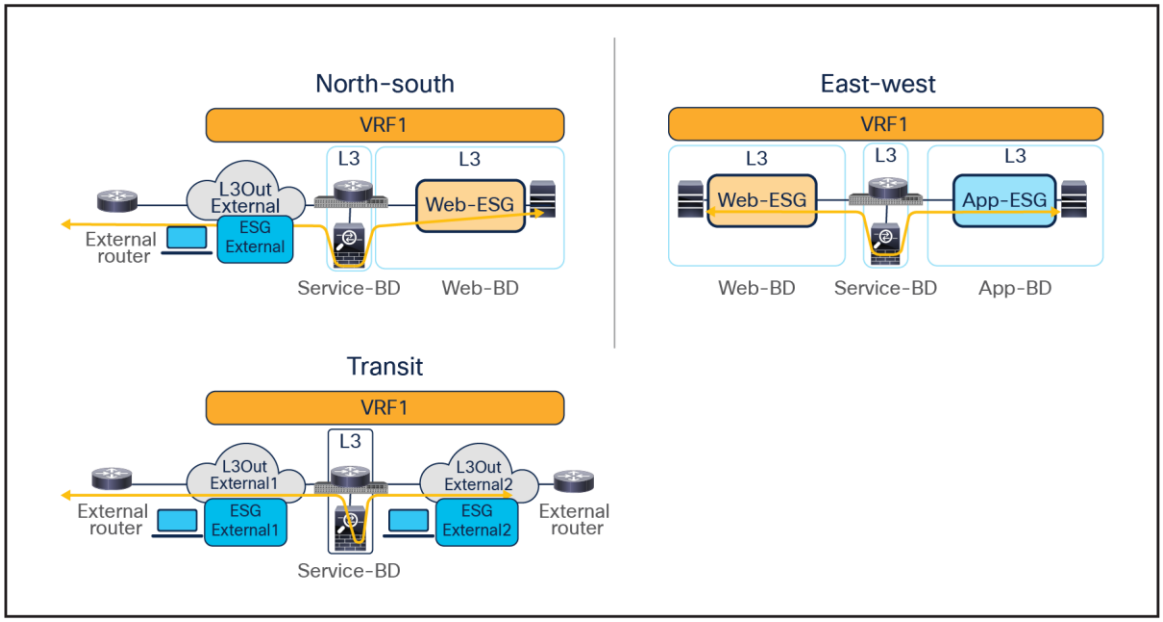


Figure 11.
Service-node insertion for north-south, east-west, and transit traffic flows (one-arm example)

Several considerations apply when deploying service graph with PBR in an ACI Multi-Site architecture:

- Service graph with PBR integration with Multi-Site is only supported when the service node is deployed in unmanaged mode. This implies that ACI only takes care of steering the traffic through the service node; the configuration of the service node is, instead, not handled by the APIC. As such, there is no requirement to support any device package, and any service node (from Cisco or a third-party vendor) can be integrated with this approach.
- In the example in Figure 11, the service node is deployed in one-arm mode, leveraging a single interface to connect to a dedicated service Bridge Domain (BD) defined in the ACI fabric. It is worth being reminded that in order to leverage service graph with PBR with ACI Multi-Site, the service node must be connected to a BD and not to an L3Out logical connection, which essentially means that no dynamic routing protocol can be used between the service node and the ACI fabric. The deployment of one-arm mode is therefore advantageous, because it simplifies the routing configuration of the service node, which requires only the definition of a default route pointing to the service BD IP address as next-hop. That said, two-arm deployment models (with inside and outside interfaces connected to separate BDs) are also fully supported, as shown in Figure 12.
- The service BD(s) must be L2-stretched across sites. This means that the interfaces of the service nodes in different sites must be in the same service BD. The recommendation is to do this without extending BUM flooding, to avoid spreading broadcast storms outside a single fabric.
- The consumer and provider ESGs such as Web ESG for internal endpoints and External ESG for external network behind an L3Out can be stretched across sites or locally confined in a site (or a combination of the two).
- vzAny-to-vzAny PBR must use one-arm service node instead of two-arm. For more guidelines and deployment considerations for those new cases, please refer to the section "Multi-Site service graph with PBR use cases."

© 2026 Cisco and/or its affiliates. All rights reserved.

- Consumer endpoints of an east-west contract with PBR must not be connected under the border leaf node where an inter-site L3Out resides.
- The following per host configurations under EPG are not supported for a consumer or a provider ESG of a contract with PBR if the VRF is stretched across sites:
  - IP Data-plane Learning: Disable (default is Enabled)
  - Type Behind Subnet: EP Reachability, MSNLB and Anycast MAC (default is None)

  If needed, you can disable data plane IP learning by using prefixes smaller than /32 for IPv4 and /128 for IPv6 at BD level instead of at EPG level. You can also disable data plane IP learning at VRF level. This workaround is not applicable to “Type Behind subnet” because it’s always /32 for IPv4 and /128 for IPv6.

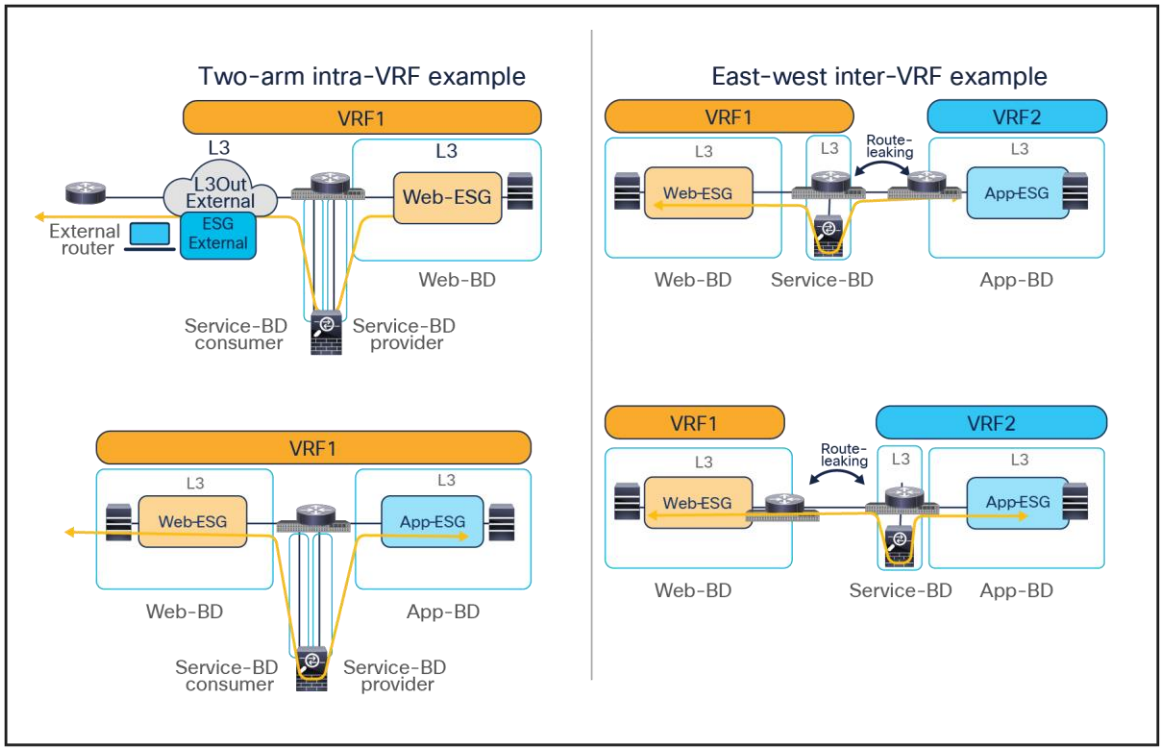


Figure 12.
Other service-insertion examples (one-arm and two-arms examples)

Though this document uses mainly a two-arm mode design in its examples, both one-arm and two-arm are valid options, except for the vzAny-to-vzAny PBR, which mandate one-arm mode service nodes.

The following are general L3 PBR design considerations that are applied to PBR in Multi-Site:

- In a Multi-Site design, redirection to a PBR node is only supported to an interface that is connected to a bridge domain (that is, the interface cannot be connected to an L3Out). However, the same physical interface connecting a PBR node to the fabric could be used for both types of connectivity when leveraging different logical interfaces (each associated to a separate VLAN tag). For example:
  - The PBR node is connected to Leaf1 Ethernet 1/1 and uses VLAN 10 to connect to the service bridge domain. This is the logical interface used for PBR for east-west communication between ESGs.

© 2026 Cisco and/or its affiliates. All rights reserved.

    - The PBR node uses, instead, VLAN 20 on the same interface Ethernet 1/1 on Leaf1 to connect to an L3Out (the L3Out must use SVIs in that case). This is the logical interface that could be used for north-south traffic, using the service node as a perimeter Firewall (FW).
- The PBR node interfaces can be part of the same bridge domain used by the consumer/provider ESG, or you can define different dedicated service bridge domains.
- The PBR node can be deployed in two-arm mode or in one-arm mode with a single interface connected to a service bridge domain. As already mentioned, this is not valid for the vzAny-to-vzAny PBR.
- Prior to Cisco ACI Release 5.2(1), the deployment of an active/standby service node pair is only supported if the active device always uses the same virtual MAC (vMAC) address. This is because those older ACI releases do not support dynamic PBR destination MAC address detection, and traffic redirection is performed by statically configuring the MAC address associated to the active service-node Virtual IP (VIP). This requirement implies that when a service node failover occurs, the standby unit that is activated must inherit both the VIP and vMAC addresses of the failed active unit (this is, for example, the case with Cisco ASA and Cisco Firepower models). Depending on the service node vendor, this might not be the default behavior, but it might have the vMAC address as a configuration option. Starting from Cisco ACI Release 5.2, this consideration is no longer applicable if dynamic PBR destination MAC detection is used instead of static PBR destination MAC configuration.

While not the main focus of this document, the following are general L1/L2 PBR design considerations that are also applied to PBR in Multi-Site:

- Cisco ACI Release 4.1(1) or later is required.
- The PBR node interfaces must be part of dedicated bridge domains.
- The PBR node can be deployed in two-arm mode, not one-arm mode.

Note: The term “PBR node” refers to the network services node (firewall, load balancer, etc.) specified in the PBR policy.

vzAny-to-vzAny PBR in an ACI Multi-Site architecture is supported under the following considerations:

- Cisco ACI Release 6.0(4c) or later is required.
- Cisco Nexus Dashboard Orchestrator Release 4.2(3e)* or later is required.
- Single-node service graph is supported, not multiple-nodes service graph.
- Only one-arm mode is supported for the service device.

ESG-to-ESG and vzAny-to-ESG PBR in an ACI Multi-Site architecture are supported under the following considerations:

- Cisco ACI Release 6.1(4) or later is required.
- Cisco Nexus Dashboard Release 4.1(1) or later is required
- Multiple-node service graph is supported.
- One-arm mode and two-arm mode are supported for the service device
- Intra-VRF and inter-VRF contracts are supported

*Nexus Dashboard 4.2(3e) is a service that runs on Nexus Dashboard 3.0(1). Nexus Dashboard 4.1 is deployed as a unified platform which includes the orchestrator feature. There is no separate orchestrator release number on Nexus Dashboard 4.1.

© 2026 Cisco and/or its affiliates. All rights reserved.

For more information about generic PBR design considerations and configurations, please refer to the document below: https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html.

Multi-Site service graph with PBR use cases

The following sections describe different use cases where service graph with PBR can be used to redirect north-south and east-west traffic flows to a service node (or to a chain of service nodes). The specific scenarios that will be considered are:

- The deployment of a single-node service graph, for redirecting traffic flows either to a firewall service or to a load-balancer service.
- The deployment of a multi-nodes service graph for redirecting traffic flows to a service chain built with firewall and load-balancer services.

As previously discussed in the "Recommended design overview" section, in an ACI Multi-Site architecture each specific type of service is implemented by leveraging a distributed model, with a separate set of service nodes in each fabric. For example, a firewall service can be represented by the deployment in each site of an active/standby pair of firewalls or of one of the other redundancy options previously shown in Figure 1. The same considerations apply to the deployment of a load-balancer service.

The critical requirement for integrating distributed stateful service nodes into an ACI Multi-Site architecture is avoiding the creation of asymmetric traffic paths for the incoming and return directions of flows, because doing so would cause communication drops due to the stateful nature of those service nodes. Figure 13 illustrates an example. For incoming traffic from an external client to an internal endpoint in site2, traffic may be steered toward the L3Out in site1, depending on the routing design. However, the outbound traffic from the internal endpoint goes out (by default) through the local L3Out in site2. The return traffic would, hence, be dropped by the external firewall connected to site2 since the firewall does not have the connection state information for the traffic flow that was created earlier on the external firewall connected to site1.

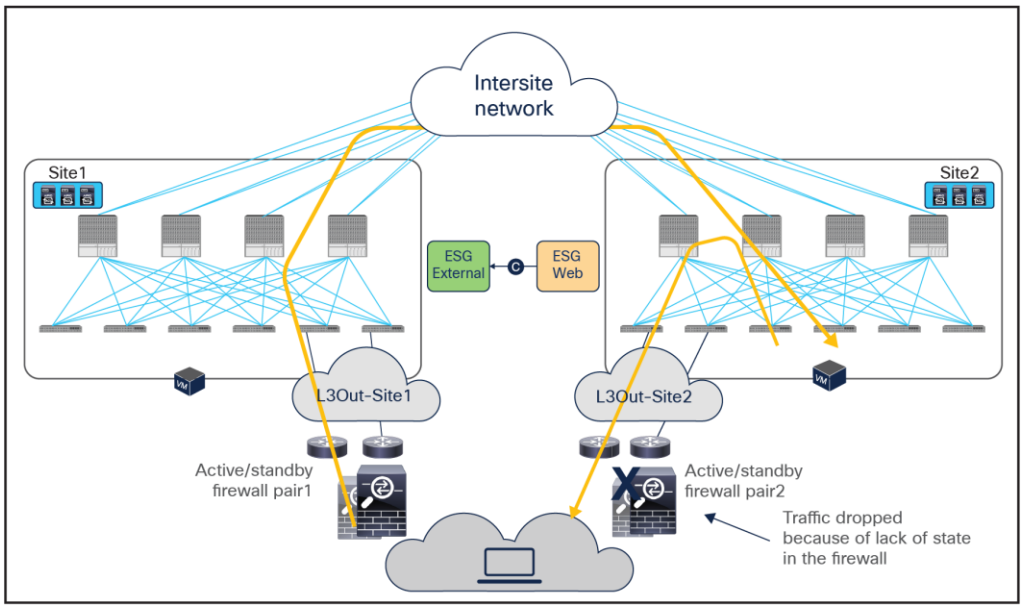


Figure 13.
Why traffic symmetricity is important in multilocation data centers

© 2026 Cisco and/or its affiliates. All rights reserved.

Even if the external firewall connected to site2 has an Access Control List (ACL) to permit outgoing traffic, the external firewall connected to site2 will drop the asymmetric outgoing traffic because firewalls are generally stateful regardless of traffic direction. For example, Cisco ASA and FTD firewalls only match the first packet of a connection to an ACL. For Transmission Control Protocol (TCP), any new connection initiation segment that is not a SYN will be dropped by an implicit stateful check and will never be matched against an ACL permit-rule by default. Only User Datagram Protocol (UDP) connections may be permitted in an asymmetrical fashion with bidirectional ACLs.

A solution is therefore required to keep both directions of traffic flowing through the same service node. The asymmetric traffic path shown in the previous figure for traffic destined to endpoints that are part of bridge domains that are stretched across sites, can be avoided by leveraging host-route advertisement to optimize the traffic path for ingress communication, but this approach to avoid asymmetricity can be used for a north-south traffic path only.

Advanced logic has been built into the ACI Multi-Site implementation to provide an elegant answer to such a requirement. As a result, the PBR policy will be enforced on specific fabric-leaf nodes that may be different depending on the specific traffic flow considered (intra-site vs. inter-site) and on the type of contract defined between the endpoints (ESG-to-ESG, vzAny-to-ESG or vzAny-to-vzAny). For example, we'll see how the provider leaf node is always used to enforce the PBR policy when using an ESG-to-ESG or vzAny-to-ESG contract.

Table 1. PBR policy enforcement in Multi-Site (Cisco ACI Release 6.1(4) or later is required.)

VRF design	ESG-to-ESG	vzAny-to-ESG	vzAny-to-vzAny
Intra-VRF	Provider leaf	Provider leaf	Intra-site traffic: either source or destination leaf node Inter-site traffic: both source and destination leaf nodes
Inter-VRF	Provider leaf	Provider leaf	Not supported

PBR to a firewall service

This section explains firewall insertion with PBR for north-south and east-west traffic flows for the following PBR use cases:

- vzAny-to-vzAny
- vzAny-to-ESG
- ESG-to-ESG

© 2026 Cisco and/or its affiliates. All rights reserved.

vzAny-to-vzAny use cases

The deployment of this service graph with PBR use case requires the use of Cisco ACI Release 6.0(4c) or later and Cisco Nexus Dashboard Orchestrator Release 4.2(3e) or later. Figure 14 shows a sample Cisco ACI network design with vzAny-to-vzAny PBR. vzAny is both the consumer and the provider of a contract with a firewall service graph attached to it with PBR enabled in both directions.

Although this example shows just three ESGs (ESG for External ESG, Web ESG, and App ESG), the VRF could have more ESGs in the same or different BDs, and the firewall could be inserted for all inter-ESGs communications because of the vzAny-to-vzAny contract with PBR.

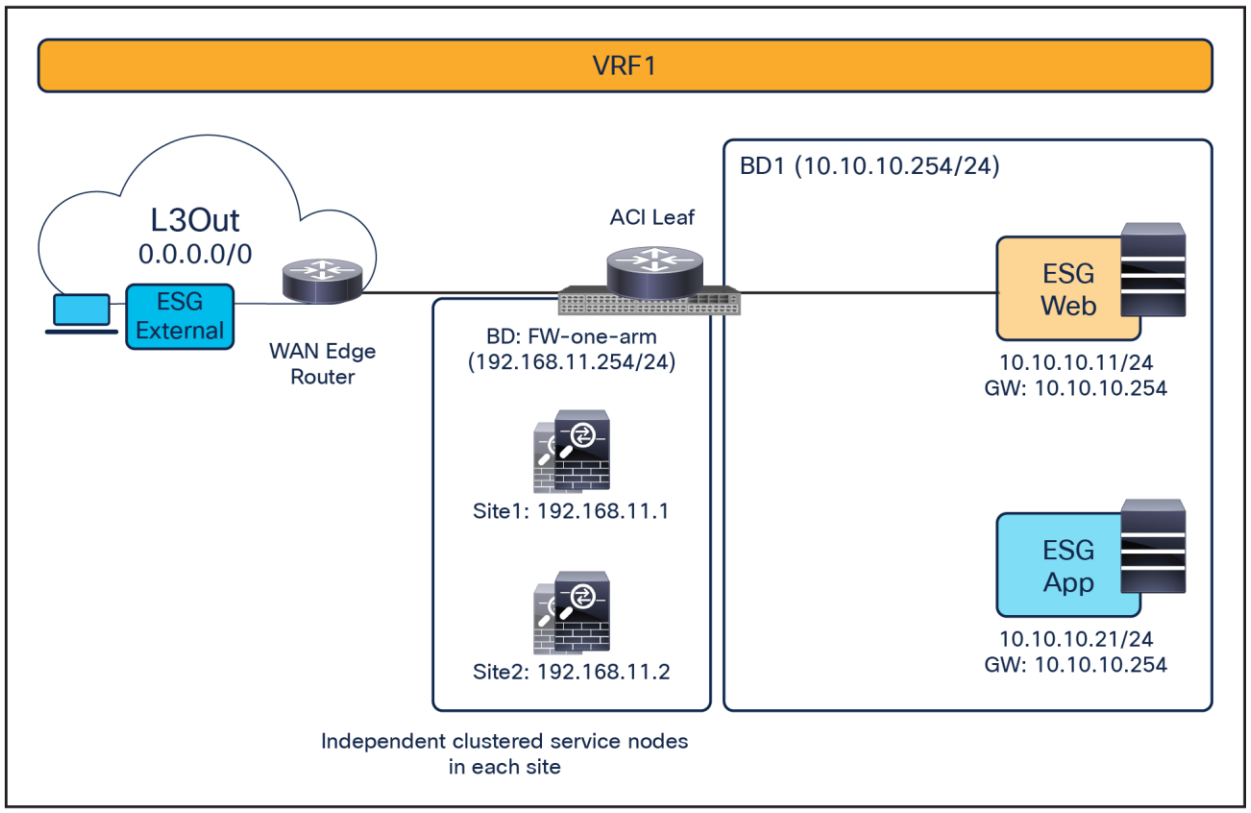


Figure 14.
North-south and east-west firewall with PBR design example (vzAny-to-vzAny)

When considering the use of vzAny-to-vzAny PBR for inter-site north-south and east-west traffic between ESGs, the difficulty of avoiding the creation of an asymmetric traffic path through the independent firewall services deployed across fabrics becomes immediately clear. This is because, differently from ESG-to-ESG PBR and vzAny-to-ESG PBR use cases, when applying a vzAny-to-vzAny PBR contract, it is not possible to distinguish the role of the consumer and the provider of the contract.

A different approach is therefore required in this case, and the chosen solution has been to redirect all north-south and east-west traffic flows to both firewall services deployed in the source and in the destination fabric (Figure 15).

© 2026 Cisco and/or its affiliates. All rights reserved.

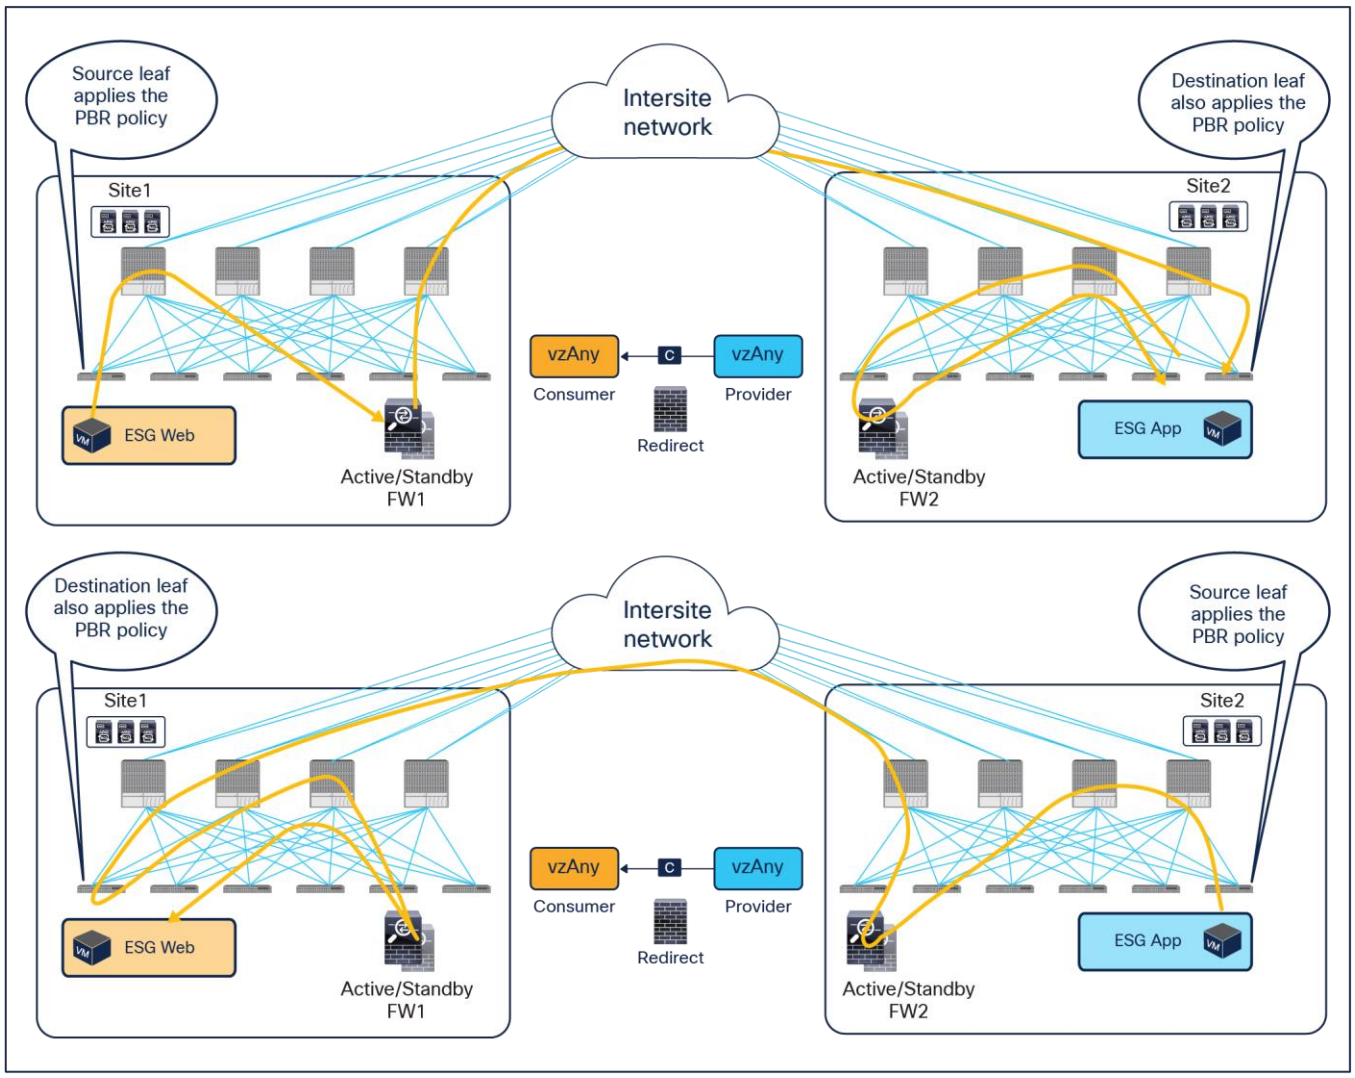


Figure 15.
Use of ACI PBR to keep inter-site traffic symmetric for the vzAny-to-vzAny PBR use case

The behavior shown above can be achieved if the PBR is applied on the source and destination leaf node for both directions of the same traffic flow. But for this to be possible, those leaf nodes should always know the class ID for the destination endpoint, and this cannot always be guaranteed under normal circumstances, hence some innovative functionalities have been introduced into ACI Multi-Site to achieve that.

© 2026 Cisco and/or its affiliates. All rights reserved.

Figure 16 illustrates an example of the ideal behavior with a vzAny-to-vzAny PBR contract where inter-site communication between an endpoint in Web ESG and an endpoint in App ESG is steered through the firewall services in both the source and the destination sites.

- When the Web endpoint sends traffic toward the App endpoint, the ingress leaf in site1 redirects the traffic to the local active firewall node. As mentioned, for this to be possible we are assuming here that the source leaf node has all the required information (that is, the source and destination class IDs) to enforce the PBR policy.
- Once the firewall in the source site has applied the locally configured security policies, the traffic is sent to the destination leaf. The service leaf in site1 sets special flags[2] in the VXLAN header, to indicate that the local firewall has been inserted and has applied its security policy.
- When the traffic arrives to the destination leaf in site2, the special flags setting and the specific source VTEP address convey the information to the leaf that the firewall in the remote source site has already seen the traffic. The leaf can therefore just apply the PBR policy to redirect the traffic through the local active firewall node.
- After the local firewall has applied its security policy, the traffic is sent to the destination leaf node again. The service leaf in site2 also set the special flags in the VXLAN header (as was done by the service leaf nodes in the source site) to indicate the fact that the local firewall has seen the traffic. However, this information is now ignored by the destination leaf because this is intra-site VXLAN traffic (that is, it originated from a local service leaf node), so the destination leaf must simply forward it to the destination endpoint.

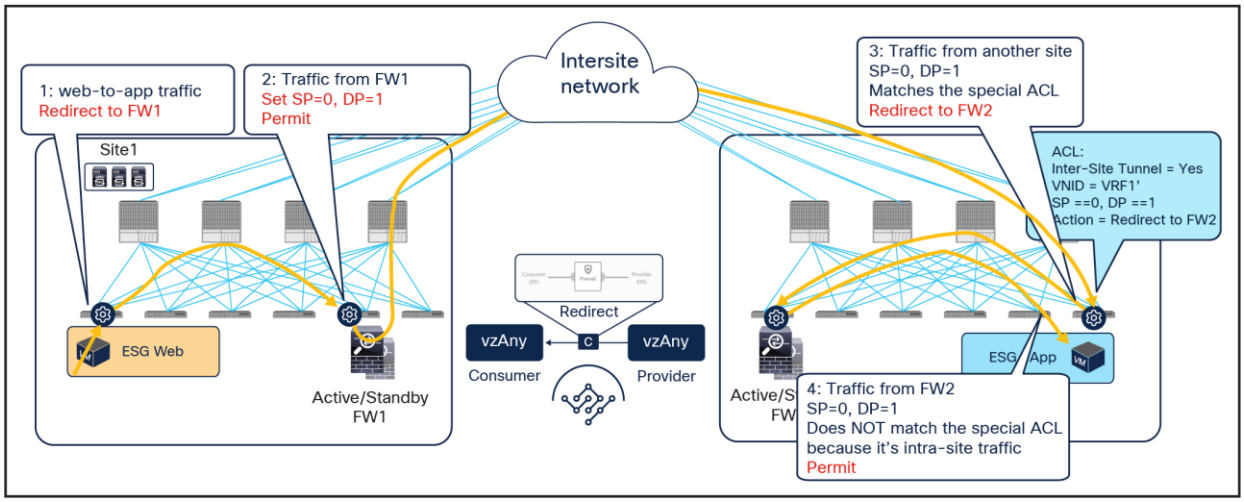


Figure 16.
Use of PBR for Web-to-App traffic flows (vzAny-to-vzAny)

[2] Note for advanced readers: When an ingress leaf applies a contract policy, the SP (source policy) and DP (destination policy) bits are set to 1 in the VXLAN header so that the destination leaf can identify whether the policy was already applied or not. If SP=1 and DP=1, the destination leaf does not apply the policy again. A new behavior has been introduced to support the vzAny PBR use cases: a service leaf node will set SP=0 and DP=1 for traffic received from the firewall node that is, the traffic with the source class ID of the service EPG and destined to the consumer/provider ESG. This is to indicate that the service node (firewall) has already been inserted in the source site.

© 2026 Cisco and/or its affiliates. All rights reserved.

The insertion of the firewall services in both sites must also be done for the return traffic flow (Figure 17)

- When the App endpoint sends traffic toward the Web endpoint, the ingress leaf in site2 redirects traffic to the local active firewall node. Again, we are assuming that the ingress leaf knows the destination class ID information to be able to locally enforce the PBR policy.
- Once the firewall has applied its locally configured security policy, the traffic is sent to the destination leaf. The service leaf in site2 encapsulates the traffic with special flags properly set in the VXLAN header.
- When the traffic arrives to the leaf in site1, it is again redirected through the local active firewall node because of the special flags setting in the traffic received from the remote site.
- After the firewall has applied the locally configured security policies, the traffic is sent back to the destination leaf. The service leaf in site1 also set the special flags in the VXLAN header, but the destination leaf does not redirect traffic again because it is intra-site traffic.

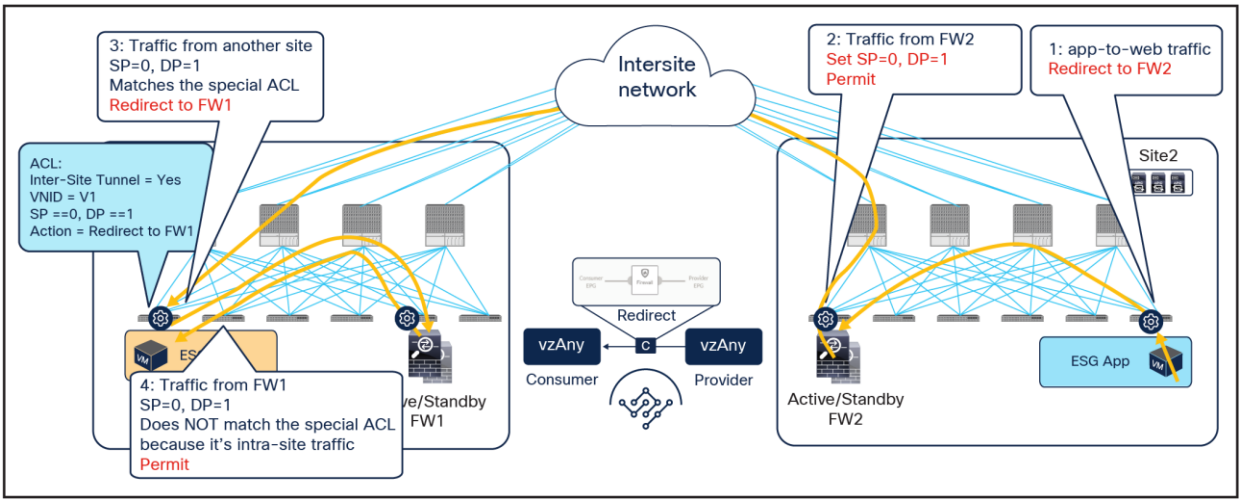


Figure 17.
Use of PBR for App-to-Web traffic flows (vzAny-to-vzAny)

For both directions of the flow, the redirection on the ingress leaf node is predicated on the leaf's knowledge of the class ID of the destination endpoint. If, for whatever reason, that is not the case, the ingress leaf cannot apply the policy, and a different mechanism is required to ensure redirection of traffic to the firewall services in both the source and the destination site.

As shown in Figure 18, if the ingress leaf cannot apply the PBR policy, the traffic is implicitly permitted, and the Policy-Applied (PA) bit in the VXLAN header is not set (PA = 0). The traffic is forwarded across sites and received by the destination leaf in the remote site, which will redirect traffic back to the active firewall node in the source site. This is because setting the PA bit to 0 indicates that the policy was not applied by the ingress leaf (and consequently not sent to the firewall in that site), and the destination leaf redirects the flow back to the firewall in the source site. After the firewall in site1 has applied its locally configured policy, the traffic is sent back to the destination leaf. The remaining flow (Figure 19) is then the same as already shown in Figure 16 (the traffic is redirected to the local firewall in site2 before reaching the destination endpoint).

© 2026 Cisco and/or its affiliates. All rights reserved.

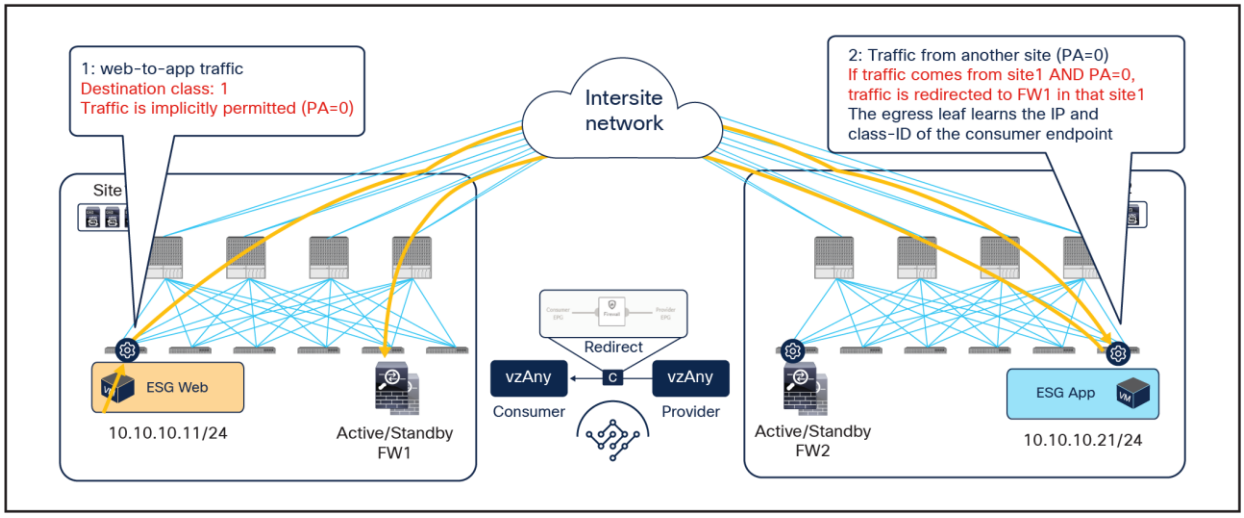


Figure 18.
Hair-pinning of traffic when the consumer leaf cannot apply the PBR policy

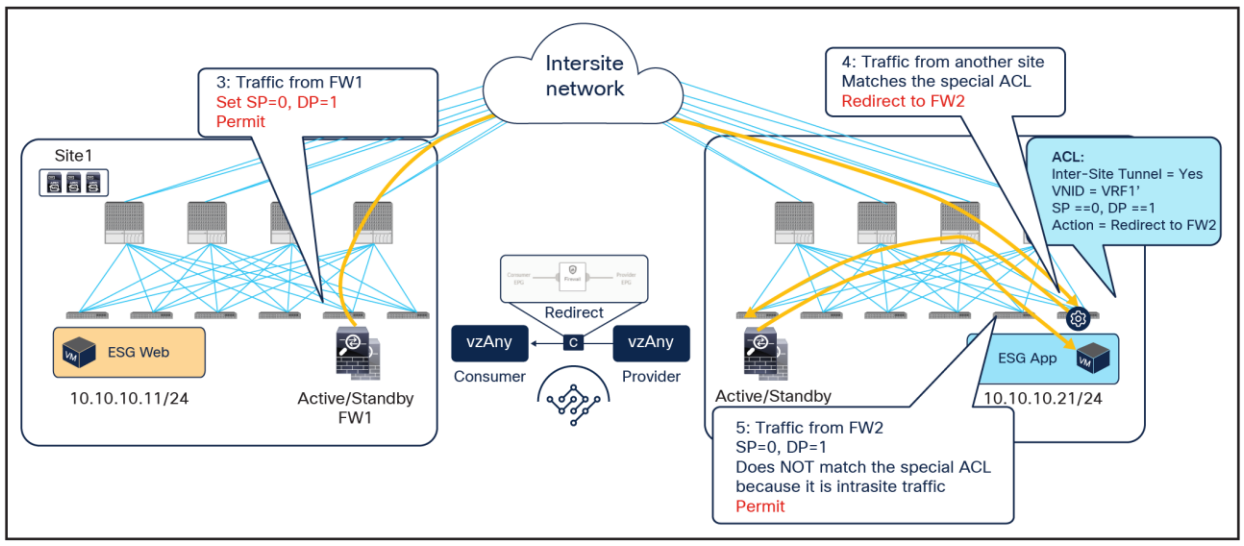


Figure 19.
Traffic forwarded back to the destination site

The traffic hair-pinning shown above, while not creating asymmetricity through the different firewall nodes, represents suboptimal data-path behavior. In order to eliminate it, an additional functionality named "conversational learning" has been implemented in Cisco ACI fabric for this vzAny-to-vzAny PBR use case.

Figure 20 shows how the reception of traffic with the PA bit set to 0 on the destination leaf triggers (in parallel to the data-plane traffic redirection shown in Figure 18 the origination of a control packet containing information about the destination endpoint IP address and class ID. This control packet is sent to the source leaf in site1, which receives it and installs the destination endpoint information on the source leaf in site1.

© 2026 Cisco and/or its affiliates. All rights reserved.

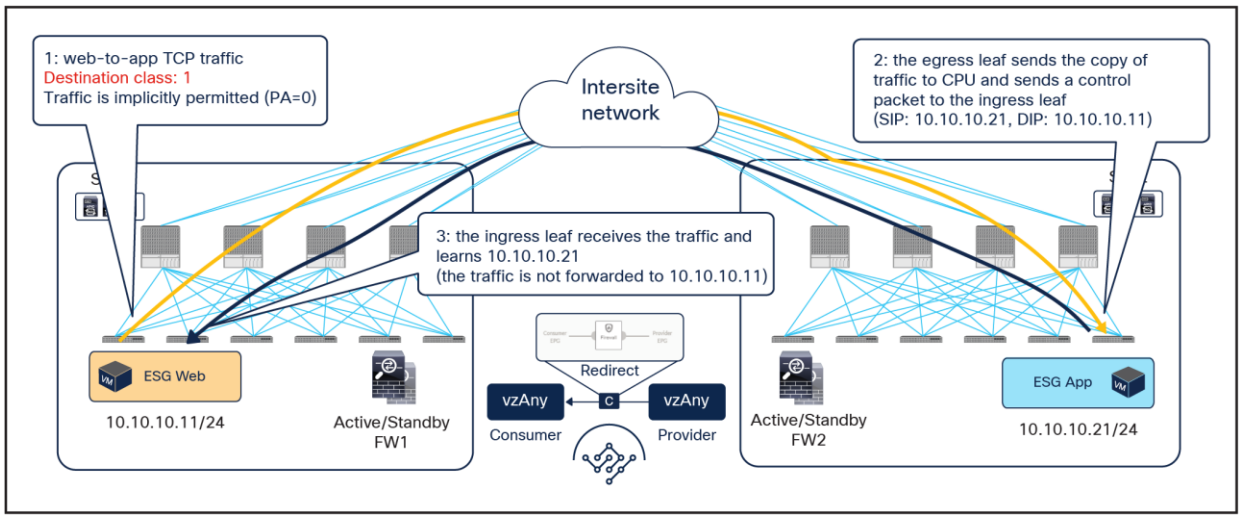


Figure 20.
Conversational learning

After the ingress leaf has learned the destination endpoint information, traffic is forwarded optimally from the source to the destination site, as previously shown in Figure 15.

It is worth noticing that if the east-west flow is between two endpoints connected to the same fabric, the traffic is redirected by either the source or the destination leaf. The endpoint IP learning status does not matter, because the local PBR destination is always used regardless of which leaf applies the PBR policy.

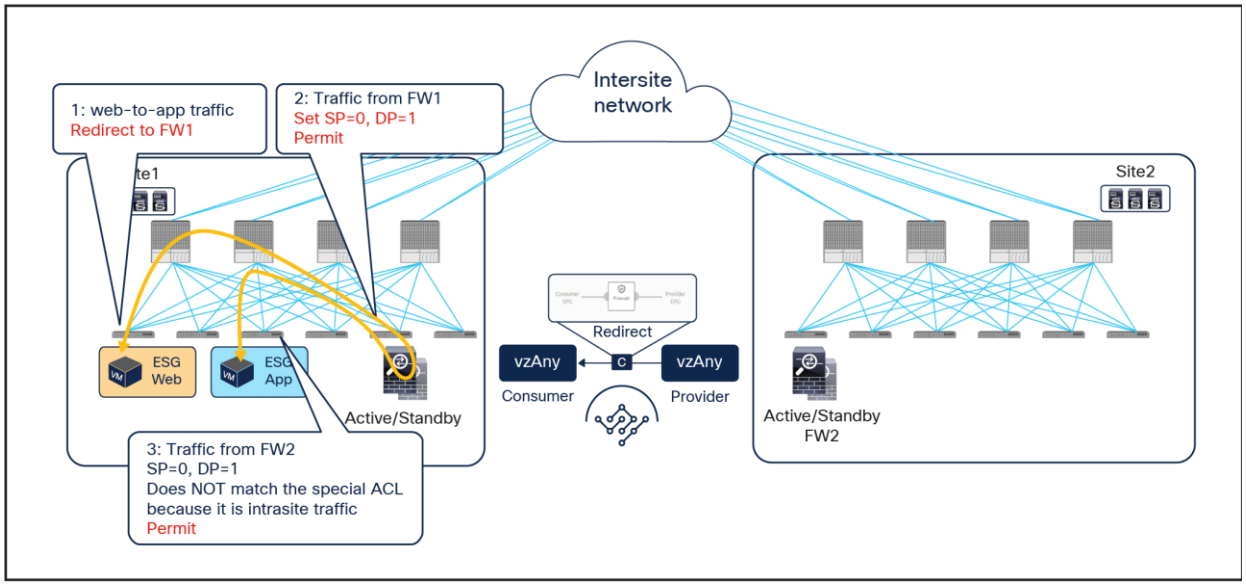


Figure 21.
Intra-site traffic (vzAny-to-vzAny)

© 2026 Cisco and/or its affiliates. All rights reserved.

Applying the policy on the ingress leaf node, which is a requirement for east-west communication in the vzAny-to-vzAny PBR use case, may represent a problem for the redirection of north-south traffic flows without inbound traffic path optimization. To better understand this issue, let's consider the scenario depicted in the figure below.

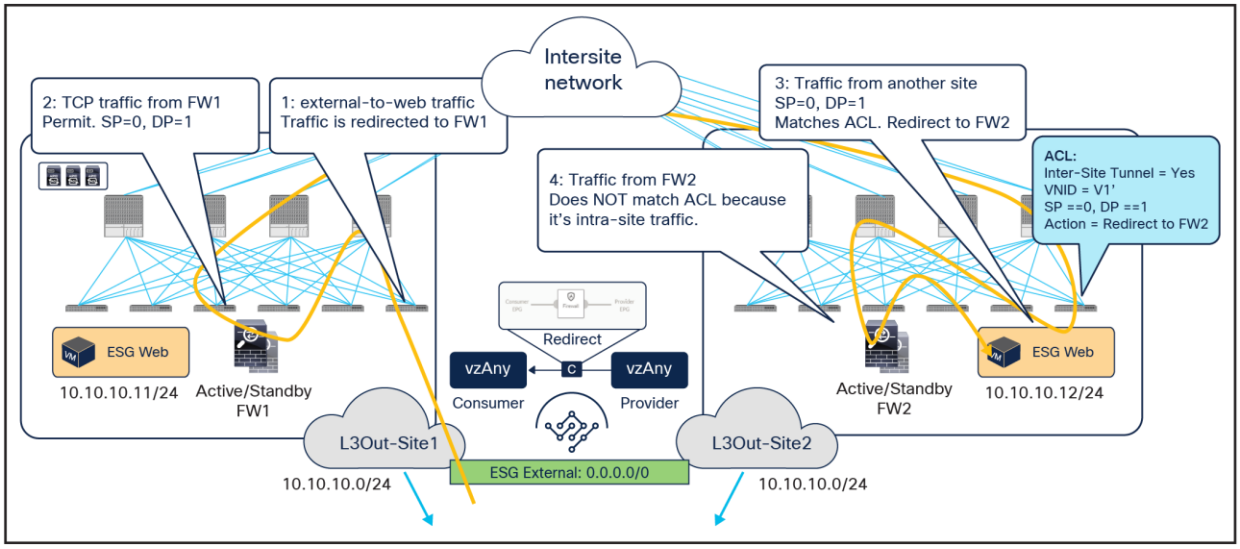


Figure 22.
Inbound flow redirected to both firewall services

Inbound traffic is received in site1 even if the destination is the Web ESG endpoint connected to site2. The border leaf node in site1 applies the PBR policy (because it is the ingress leaf), and traffic is redirected first to the firewall in site1 and then to the firewall in site2, as expected for the vzAny-to-vzAny use case.

The below shows instead the return traffic flow, from the Web ESG endpoint in site2 to the external destination.

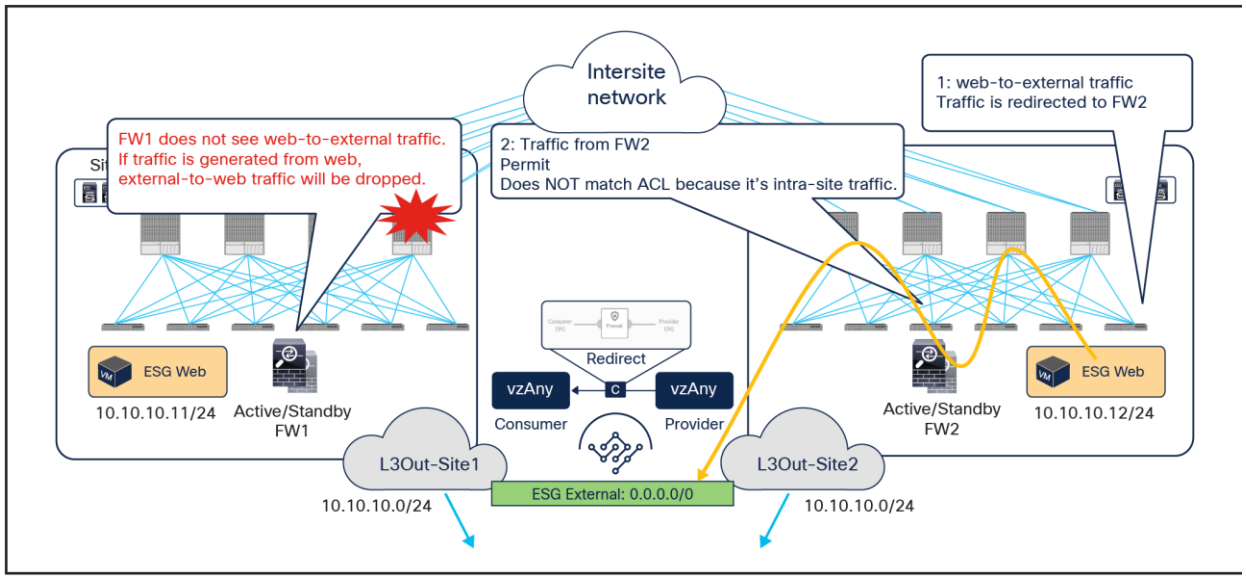


Figure 23.
Outbound flow redirected only to the firewall service in site2

© 2026 Cisco and/or its affiliates. All rights reserved.

The outbound flow can only be redirected to the firewall service in site2, and this causes an asymmetric behavior that will cause traffic drop when the north-south communication is initiated by the Web endpoint.

The solution to this problem, highlighted in the figure below, consists in enabling host-based routing advertisement so that inbound traffic paths are optimized. Notice that this means that, for north-south traffic redirection with the vzAny-to-vzAny PBR use case, the redirection should only happen to the firewall located in the site where the internal endpoint is connected.

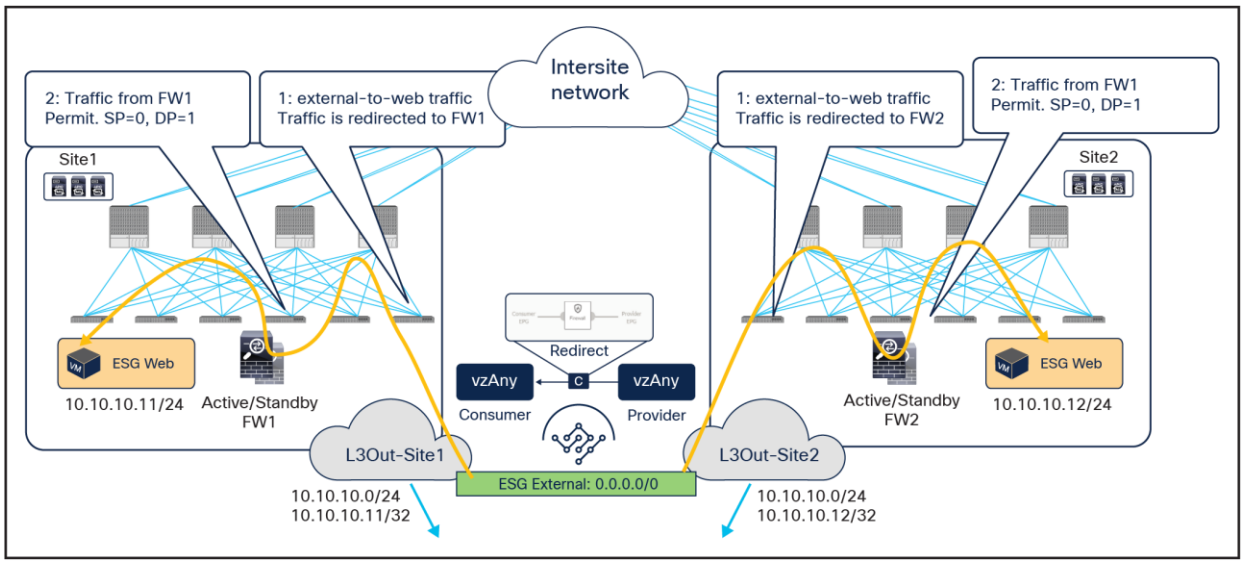


Figure 24.
Inbound traffic optimization

vzAny-to-ESG use cases

Figure 25 shows a sample Cisco ACI network design for east-west and north-south firewall insertion with vzAny-to-ESG PBR. App ESG and vzAny have a contract with a firewall service graph attached to it, with PBR enabled in both directions.

Although the figure below shows just three ESGs (External ESG, Web ESG, and App ESG) and one-arm firewall, the VRF could have more ESGs in the same or in different BDs, and one or more service devices could be inserted by using one-arm or two-arm for the communication between all the ESGs in the VRF and the App ESG, as a result of the vzAny-to-App contract with PBR.

Note: Though these examples use an intra-VRF contract, inter-VRF contract is also supported.

© 2026 Cisco and/or its affiliates. All rights reserved.

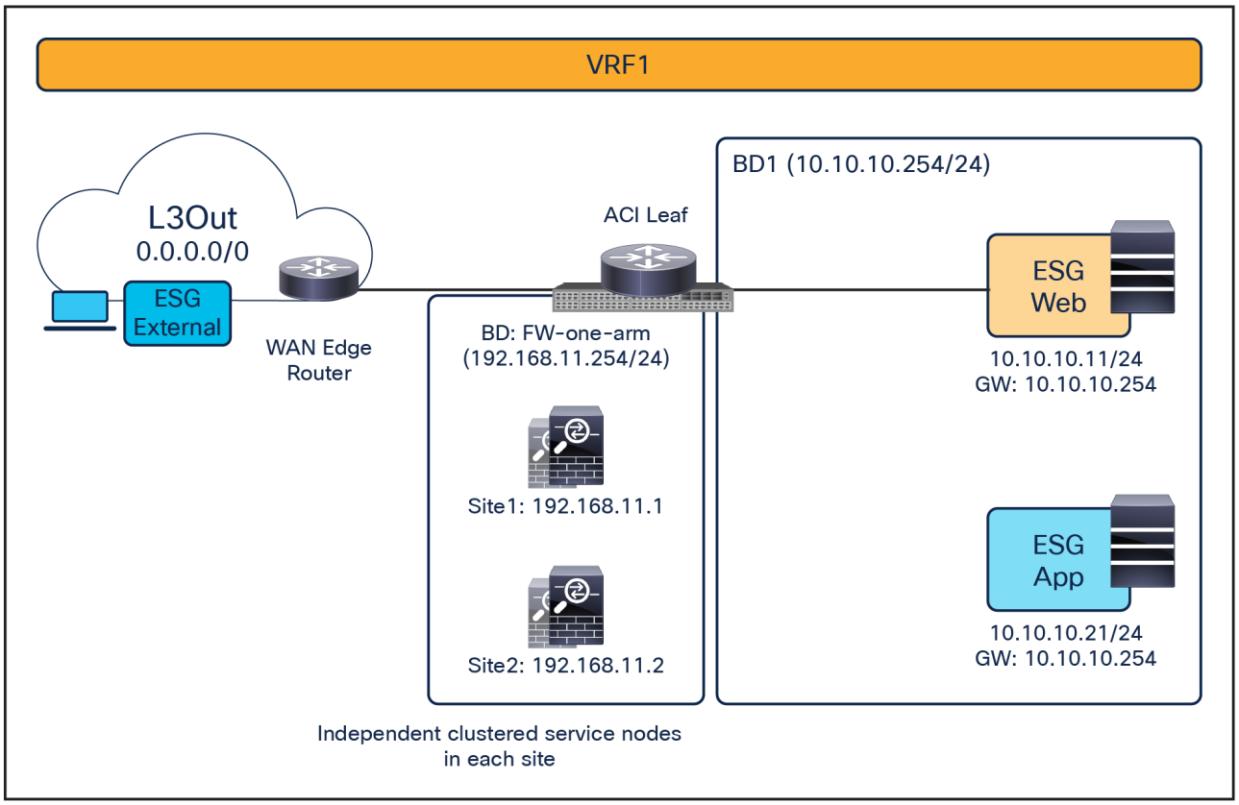


Figure 25.
North-south and east-west firewall with PBR design example (vzAny-to-ESG)

Figure 26 illustrates an example of a service-graph PBR deployment for the inter-site communications between an endpoint in Web ESG and an endpoint in App ESG. In this case, in order to avoid the creation of an asymmetric path across separate firewall nodes, the policy must be applied on the provider leaf node for both directions of the same flow.

- When the Web endpoint (representing a consumer ESG part of vzAny) sends traffic toward an App endpoint, the consumer leaf just forwards the traffic toward the provider leaf where the App endpoint has been discovered. The consumer leaf is programmed for not applying the PBR policy unless the consumer leaf is the provider leaf where the destination is connected, which uses a “redirect override” flag.
- The PBR policy kicks in on the provider leaf, and the traffic gets redirected through the local active firewall node.
- Once the firewall has applied its locally configured security policy, the traffic is sent back toward the fabric and forwarded to the App endpoint.

© 2026 Cisco and/or its affiliates. All rights reserved.

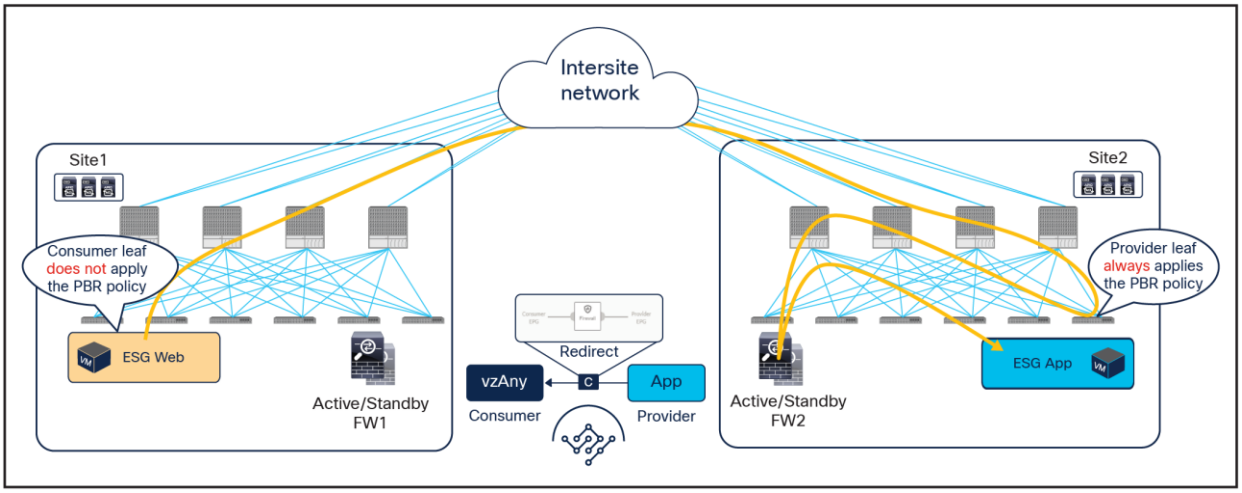


Figure 26.
Use of PBR for consumer-to-provider traffic flows (vzAny-to-ESG)

The same firewall is inserted for the return traffic flow (Figure 27).

- The PBR policy is applied on the provider leaf, and the traffic is steered through the same firewall node that built the connection state by receiving the incoming traffic. This is under the assumption that the provider leaf knows the class ID for the consumer endpoint. If that was not the case, the traffic would be sent directly to the consumer leaf in site1, which would redirect the flow to the remote firewall in site2 and generate the control-plane packet required for conversation learning similar to the example shown in Figure 20 (the same behavior already discussed for the previous vzAny-to-vzAny PBR use case).
- Once the firewall has applied its local security policy, the traffic is sent back toward the remote site and forwarded to the Web endpoint.
- The consumer leaf does not apply the policy, because this was already done on the provider leaf.

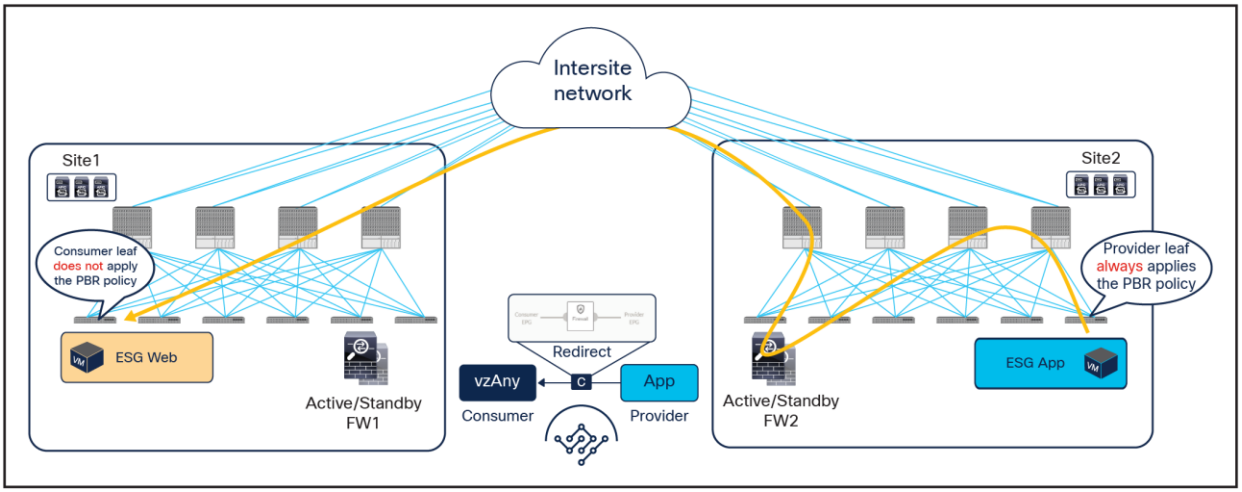


Figure 27.
Use of PBR for provider-to-consumer traffic flows (vzAny-to-ESG)

If the source endpoint and the destination endpoints are located in the same site, the traffic is always redirected to a local firewall node, and there is no traffic hair-pinning across sites as shown the figure below.

© 2026 Cisco and/or its affiliates. All rights reserved.

ESG-to-ESG use cases

The figures below show Cisco ACI network design for north-south and east-west routed firewall insertion with an intra-VRF ESG-to-ESG contract with PBR. A contract with a service graph attached is applied between ESGs. The service graph is configured with PBR enabled in both directions to steer the traffic to a firewall service device that can be one-arm or two-arm.

The key point here is that both north-south and east-west traffic use cases can be covered by using ESG-to-ESG contract because external EPG selector and external subnet selector are introduced for ESG classification.

Note: Though these examples use an intra-VRF contract, inter-VRF contract is also supported.

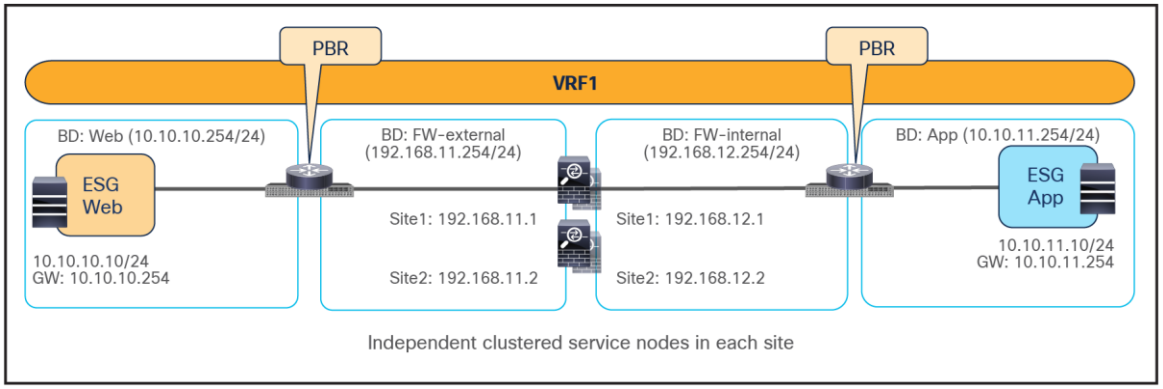


Figure 28.
Example of firewall insertion with PBR for east-west (intra-VRF)

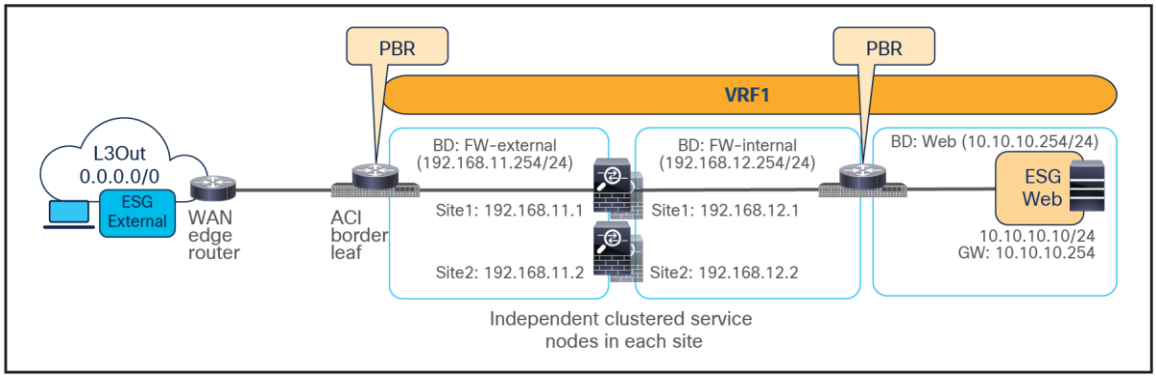


Figure 29.
Example of firewall insertion with PBR for north-south (intra-VRF): External ESG containing external EPG or external subnet

© 2026 Cisco and/or its affiliates. All rights reserved.

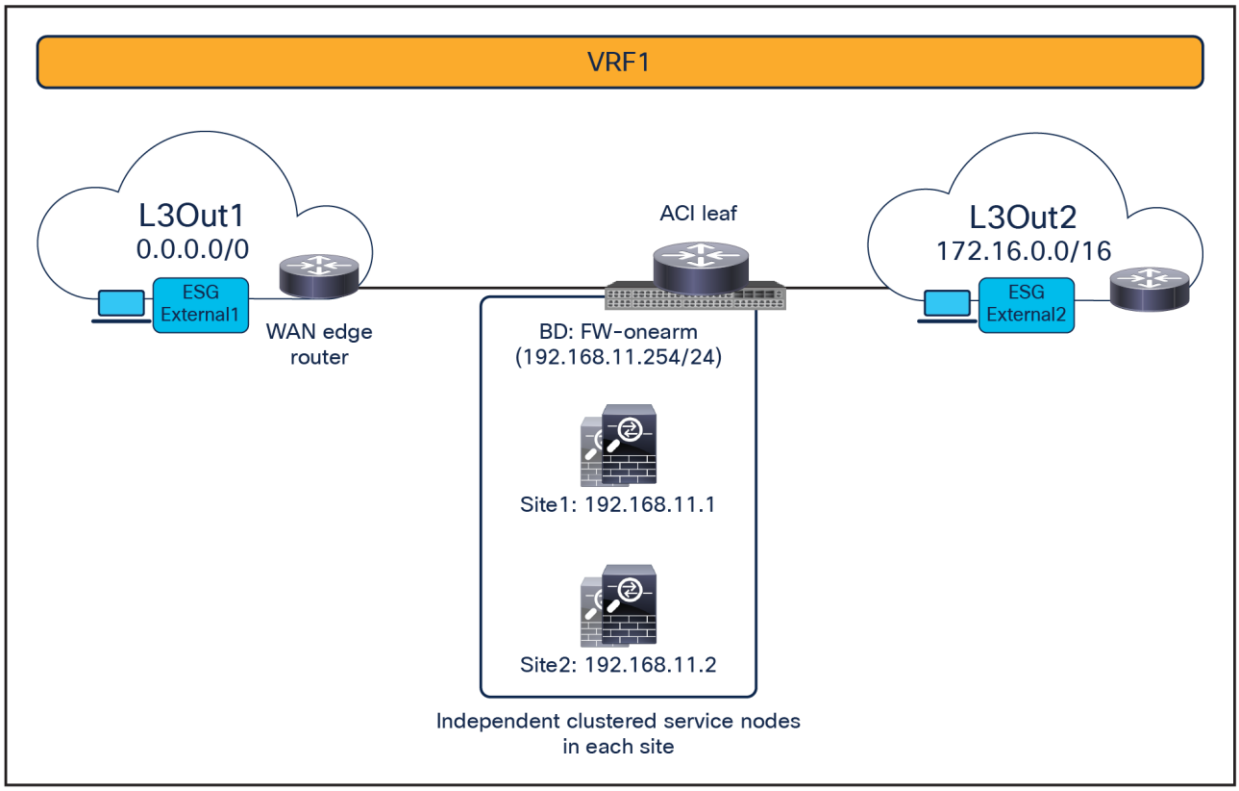


Figure 30.
Example of firewall insertion with PBR for transit traffic (intra-VRF): External ESG containing external EPG or external subnet

The figures below illustrate an example of a service-graph PBR deployment for the inter-site communications between an endpoint in Web ESG and an endpoint in App ESG. In this case, in order to avoid the creation of an asymmetric path across separate firewall nodes, the policy must be applied on the provider leaf node for both directions of the same flow, which is the same behavior already discussed for the previous vzAny-to-ESG PBR use case.

- When the Web endpoint sends traffic toward an App endpoint, the consumer leaf just forwards the traffic toward the provider leaf where the App endpoint has been discovered. The consumer leaf is programmed for not applying the PBR policy unless the consumer leaf is the provider leaf where the destination is connected, which uses a “redirect override” flag.
- The PBR policy kicks in on the provider leaf, and the traffic gets redirected through the local active firewall node.
- Once the firewall has applied its locally configured security policy, the traffic is sent back toward the fabric and forwarded to the App endpoint.

© 2026 Cisco and/or its affiliates. All rights reserved.

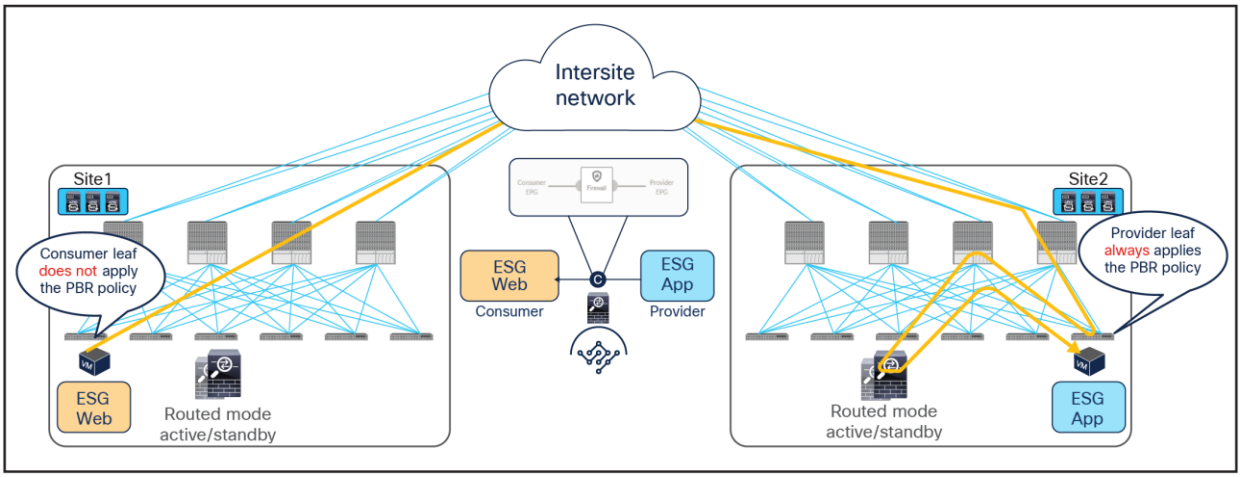


Figure 31.
Use of PBR for consumer-to-provider traffic flows (east-west)

The same firewall is inserted for the return traffic flow (Figure 32).

- The PBR policy is applied on the provider leaf, and the traffic is steered through the same firewall node that built the connection state by receiving the incoming traffic. This is under the assumption that the provider leaf knows the class ID for the consumer endpoint. If that was not the case, the traffic would be sent directly to the consumer leaf in site1, which would redirect the flow to the remote firewall in site2 and generate the control-plane packet required for conversation learning similar to the example shown in Figure 20 (the same behavior already discussed for the previous vzAny-to-vzAny PBR use case).
- Once the firewall has applied its local security policy, the traffic is sent back toward the remote site and forwarded to the Web endpoint.
- The consumer leaf does not apply the policy, because this was already done on the provider leaf.

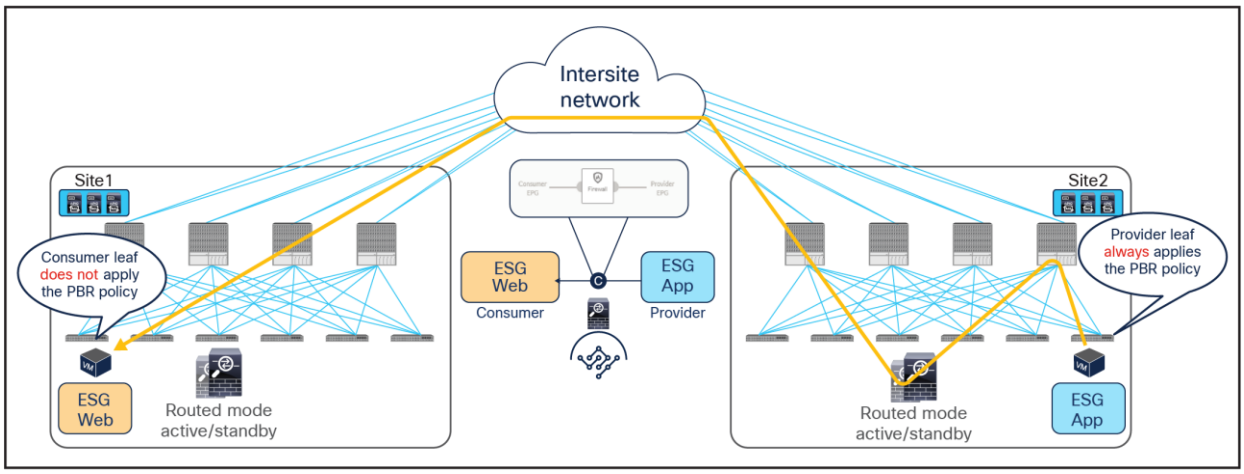


Figure 32.
Use of PBR for provider-to-consumer traffic flows (east-west)

© 2026 Cisco and/or its affiliates. All rights reserved.

If the source endpoint and the destination endpoints are located in the same site, the traffic is always redirected to a local firewall node, and there is no traffic hair-pinning across sites as shown the figure below.

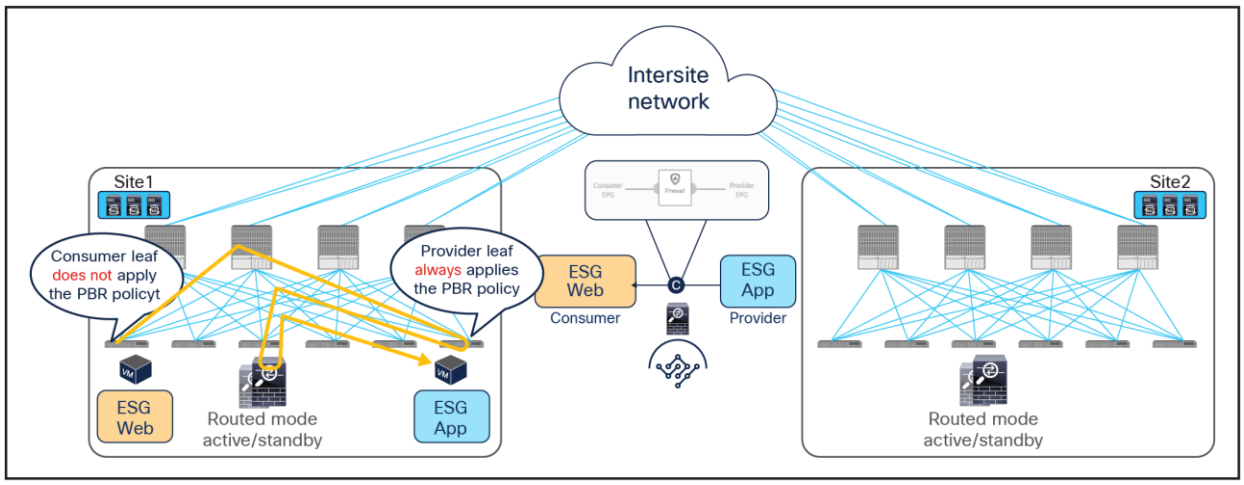


Figure 33.
East-west traffic within a site

The behavior of the provider leaf always applying the PBR policy is same for ESG-to-ESG contract for north-south traffic as well. The figures below illustrate an example of a service-graph PBR deployment for the inter-site communications between an endpoint in Web ESG and an external endpoint behind an L3Out.

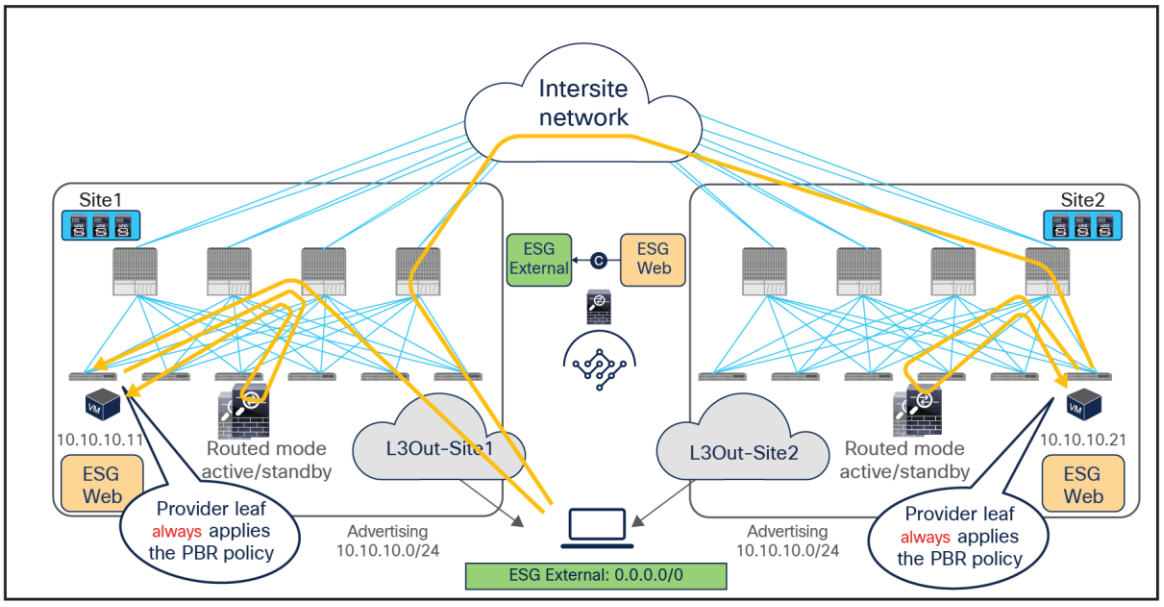


Figure 34.
Use of PBR for inbound traffic flows (north-south)

© 2026 Cisco and/or its affiliates. All rights reserved.

The outbound flows toward the external endpoint doesn't require conversational learning because the destination consumer class ID can be always resolved based on the IP addresses.

- The destination endpoints send traffic back to the external destination, and the PBR policy is again applied on the same provider leaf where it was applied for the inbound direction.
- Once the firewalls have applied the locally configured security policies, the traffic is then sent back to the fabric and forwarded to the external endpoint through the local L3Out connection. This is the default behavior, unless specific routing policies are configured to ensure the outbound flow is sent through an L3Out connection deployed in a remote site.

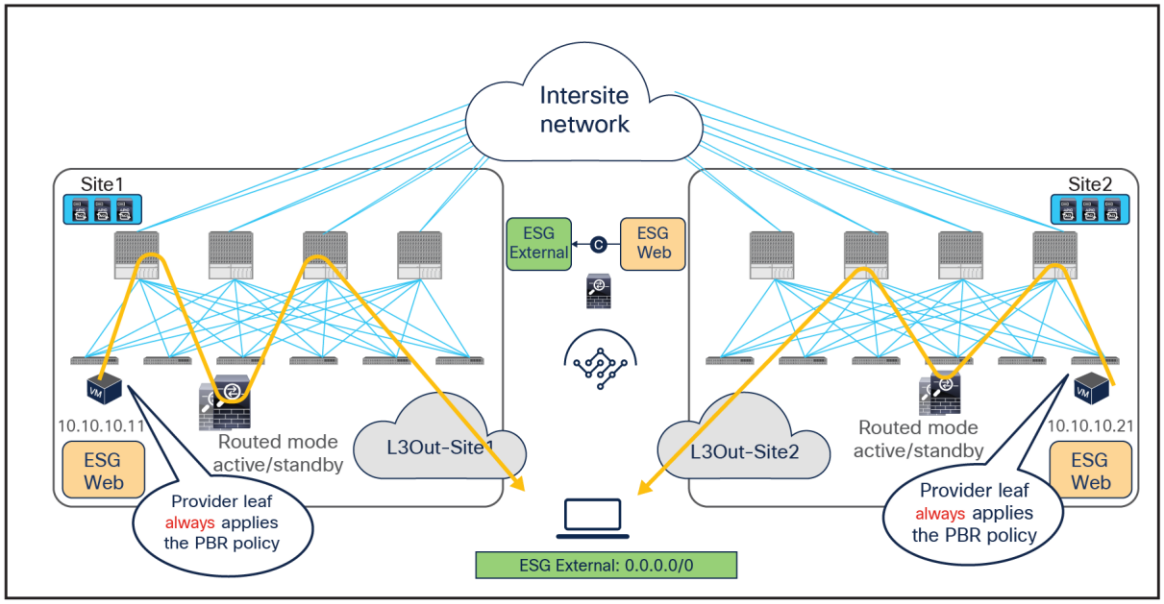


Figure 35.
Use of PBR for outbound traffic flows (north-south)

When comparing the two previous figures, it is evident, regarding the endpoint sitting in site2, that there may be an "asymmetric" use of the L3Out connection (that is, inbound traffic uses L3Out-Site1, whereas outbound traffic is sent through L3Out-Site2), but there may be a "fully symmetric" use of the same service node for both directions of the communication. This is always the case for ESG-to-ESG contract with PBR in Cisco ACI Multi-Site where the provider leaf applied the policy. This is same for both intra-VRF and inter-VRF use cases. Table 1 summarize the policy enforcement in the different use cases.

When you have available L3Out connections in both fabrics, the web server subnet stretched across sites is advertised through the border leaf nodes in both sites. As previously discussed, depending on the specific routing metric design, incoming traffic may be steered to the border leaf nodes of one of the sites. This suboptimal inbound traffic can be avoided by leveraging host-route advertisement to optimize the traffic path for ingress communication. With the use of service graph and PBR, such an approach represents only an optimization, but it is not necessary to avoid the establishment of asymmetric traffic across stateful services (as the previous example in Figure 22 and Figure 23 describes). Figure 36 illustrates how to optimize the inbound traffic flows: the destination IP address is the endpoint 10.10.10.11 located in Site1, and, because of the host route advertisement function, traffic originating from an external client can be selectively steered to Site1 and reach the destination leaf where the 10.10.10.11 endpoint is located. The destination leaf in Site1 then selects the local active PBR node, which sends traffic back to the destination. Similar behavior is achieved for traffic destined for the endpoint 10.10.10.21 in Site2.

© 2026 Cisco and/or its affiliates. All rights reserved.

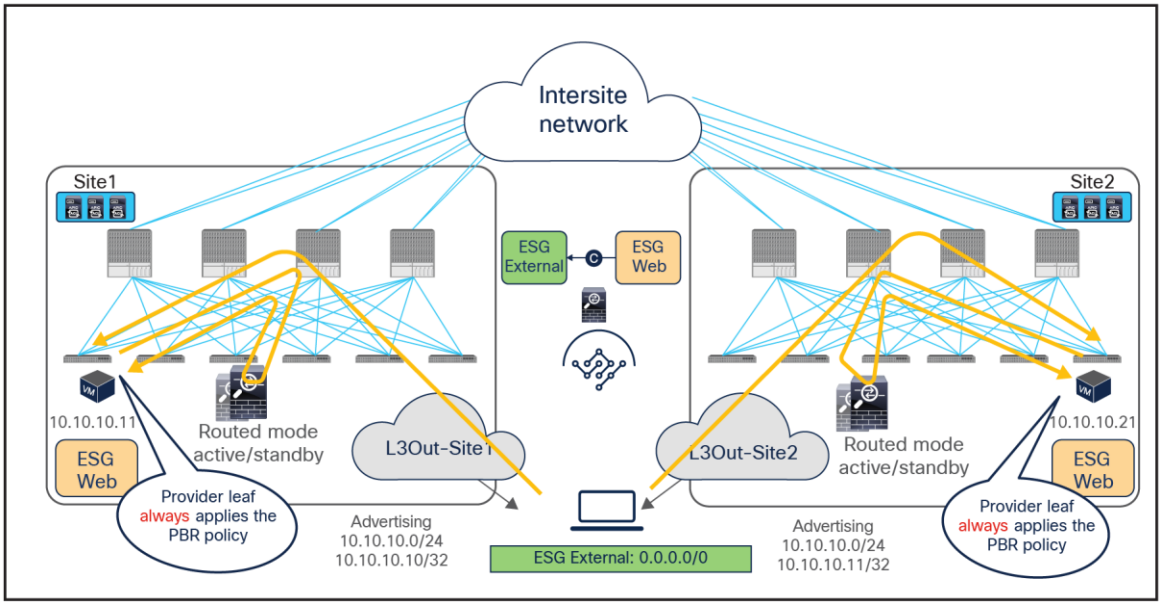


Figure 36.
Use of host route advertisement for ingress traffic optimization (north-south)

Note: In order to keep the configuration simple and to be able to apply a single ESG-to-ESG contract, the External EPG associated to the L3Out (that is, used for the External ESG classification) must be deployed as a "stretched" object associated to each site-specific L3Out. For more information, please refer to the ACI Multi-Site paper below: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html.

Load balancer with Source Network Address Translation (SNAT)

This section explains load-balancer insertion with Source Network Address Translation (SNAT) for north-south and east-west traffic use cases. In this deployment model, applying a PBR policy is not required because both incoming traffic flows (toward the VIP address) and return traffic flows (toward the IP address translated by SNAT) are destined to the load balancer and do not need redirection services.

Though this document uses a contract with a load-balancer service graph as an example, a service graph is not mandatory for this design. The main differences between the use of a service graph without PBR and the non-use of a service graph are the following:

- Use of service graph without PBR

  As shown on the left side of Figure 37, using the service graph without PBR brings the advantage of being able to simply define a contract between the consumer (clients) and the provider (server farm). The service EPGs for the load-balancer interfaces are automatically created through the service graph, together with the required contracts to ensure traffic can flow in both directions.

- Non-use of service graph

  In this case, two different contracts are required. The first one is between the consumer ESG (clients) and the ESG for the interface of the load balancer facing the clients. The second is between the ESG for the interface of the load balancer performing SNAT and the provider ESG (server farm) associated to the Virtual IP (VIP). If there is no contract security requirement, use of the same ESG for clients and the load balancer, and servers and the load balancer is also an option.

© 2026 Cisco and/or its affiliates. All rights reserved.

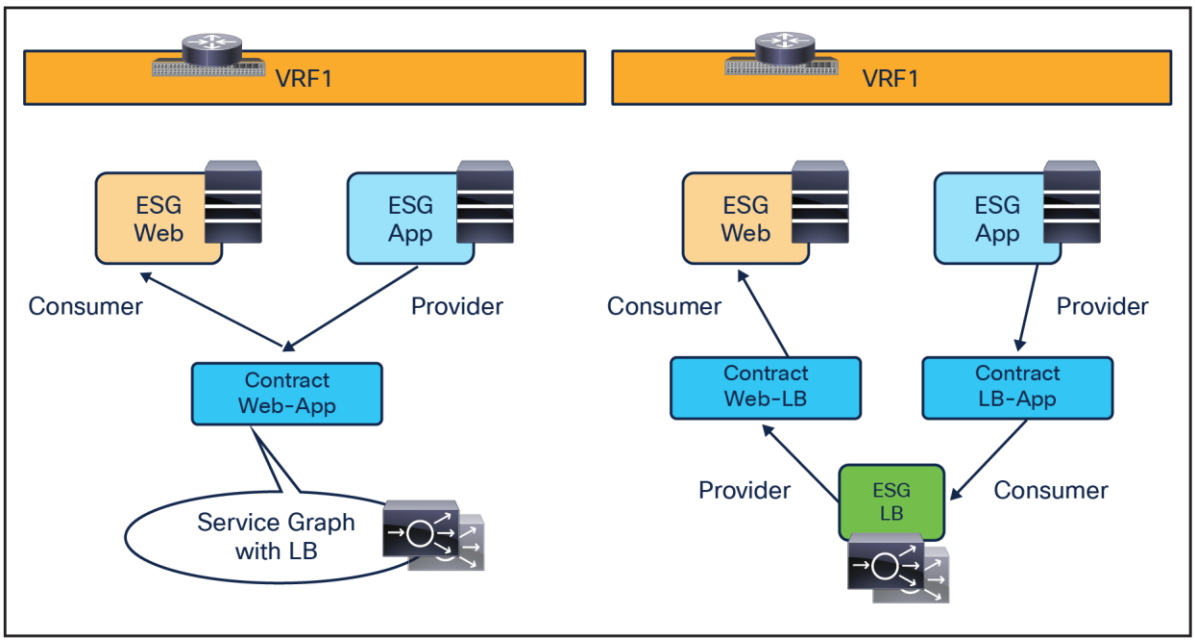


Figure 37.
Load-balancer insertion with (left) and without (right) service graph

Although this section uses intra-VRF ESG-to-ESG contract as an example, inter-VRF contract and vzAny-to-ESG contract can be used as well.

North-south traffic use case (ESG-to-ESG and vzAny-o-ESG)

The figure below shows a sample Cisco ACI network design for north-south and east-west routed load-balancer insertion with SNAT. The consumer External ESG and the provider Web ESG have a contract with a load-balancer service graph (without PBR). The endpoints in the Web ESG are the real servers that are part of the server farm associated to the VIP of the load balancer. You can have multiple load balancers, which can be represented by multiple high-availability pairs deployed in separate sites.

The assumption here is that each load balancer pair has assigned a unique VIP address that is part of the same service BD, as shown in the example below. In this scenario, Global Server Load Balancing (GSLB) can be used for load balancing access to a specific application through multiple VIPs.

Note: If a service graph is not defined, using the same service BD for each load balancer pair is not mandatory. Each load-balancer pair can use a unique VIP address in different service BDs. Also, without a service graph, an inter-VRF design is also possible.

© 2026 Cisco and/or its affiliates. All rights reserved.

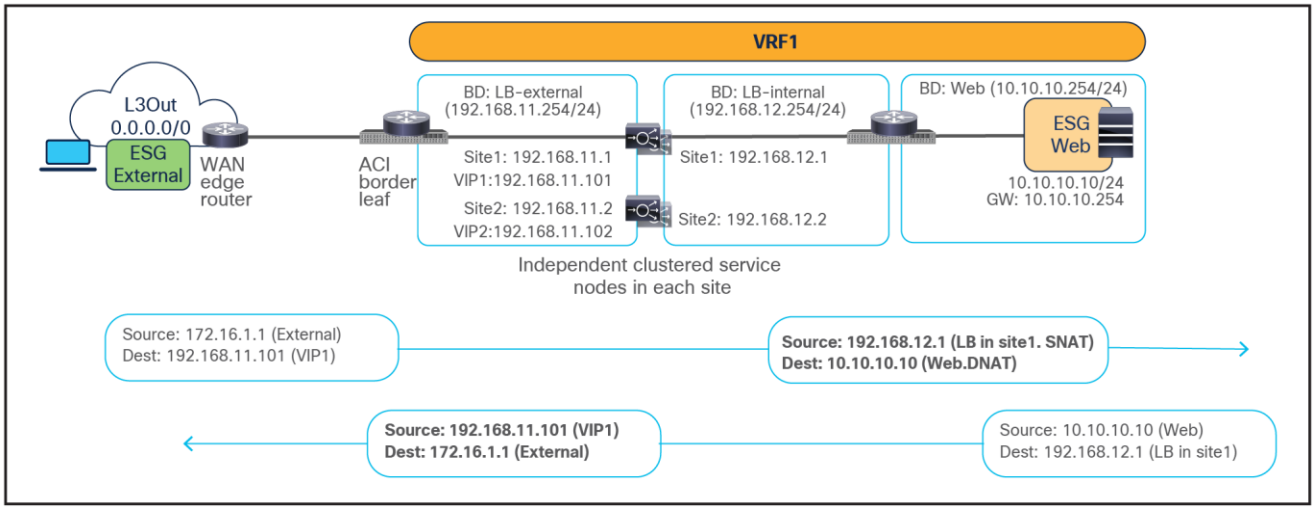


Figure 38.
Example of a north-south load balancer with a SNAT design

The figure below illustrates an example of communication between the external network and an internal Web ESG in an ACI Multi-Site deployment where we have two connections: one is between the external network and the VIP (the frontend connection), and the other one is between the load-balancer internal IP address and the real servers in the Web ESG (the backend connection). In this example, the internal Web ESG and the L3Out are defined in the same VRF.

- The incoming traffic originating from the external client is destined to the VIP, so it will be received on the L3Out connection of one of the connected sites and will then reach the load balancer without PBR as long as the VIP is reachable (this is basic forwarding behavior).
- The load balancer changes the destination IP to one of the real servers associated to the VIP. In this example, the load balancer also translates the source IP to the SNAT IP owned by the load balancer.
- After that, the traffic is forwarded to the real server.

Note: The suboptimal inbound traffic can be avoided by leveraging host-route advertising to optimize the traffic path for ingress, if the VIP of the load balancer belongs to a stretched subnet. Alternatively, it is possible to use VIP addresses in separate IP subnets for load balancers deployed in different sites. In the case of service graph, the separate IP subnets need to be configured under the same service BD because an L2-stretched service BD is required for a service graph.

© 2026 Cisco and/or its affiliates. All rights reserved.

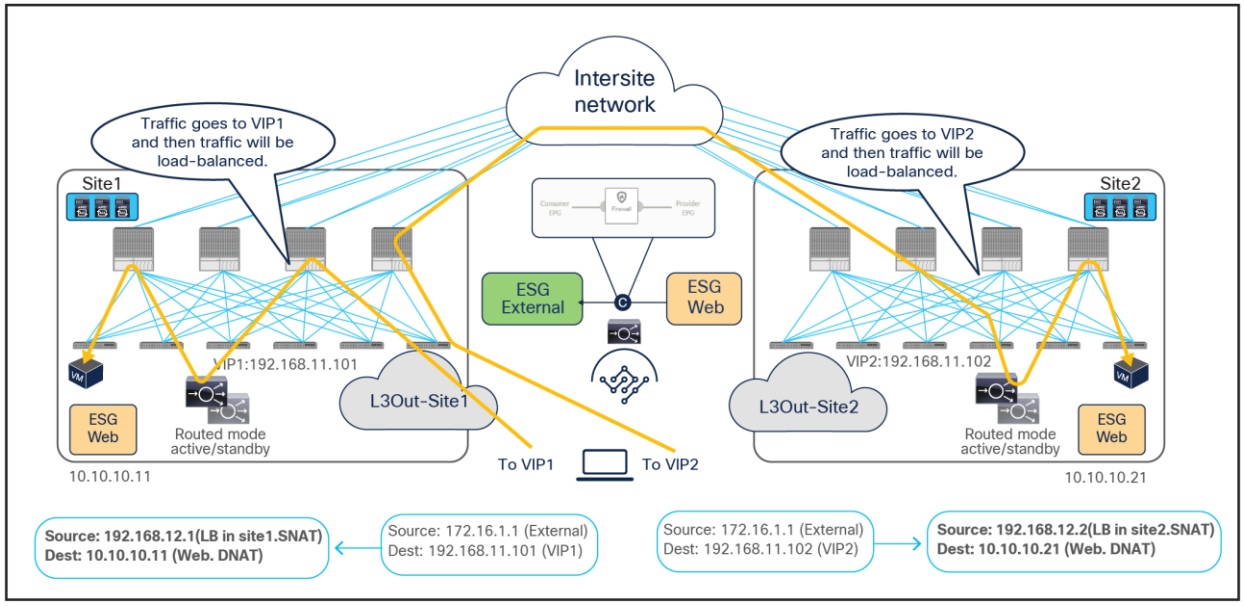


Figure 39.
Load balancer with SNAT inbound traffic flows (north-south)

Because the return traffic is destined to the SNAT IP owned by the load balancer that handled the incoming traffic flow, PBR is not required for the return traffic either.

- The load balancer receives the traffic from the Web real server endpoint and changes the source and destination IP addresses (the source becomes the VIP, the destination becomes the external client).
- The traffic is sent back to the fabric and forwarded to the external client through a local L3Out connection (unless a specific configuration is provisioned to prefer a remote L3Out connection to communicate with the external client).

Though there may be an “asymmetric” use of the L3Out connection (for example, for VIP2, inbound traffic uses L3Out-Site1, whereas outbound traffic is sent through L3Out-Site2), there is always a “fully symmetric” use of the same service node for both legs of the communication.

© 2026 Cisco and/or its affiliates. All rights reserved.

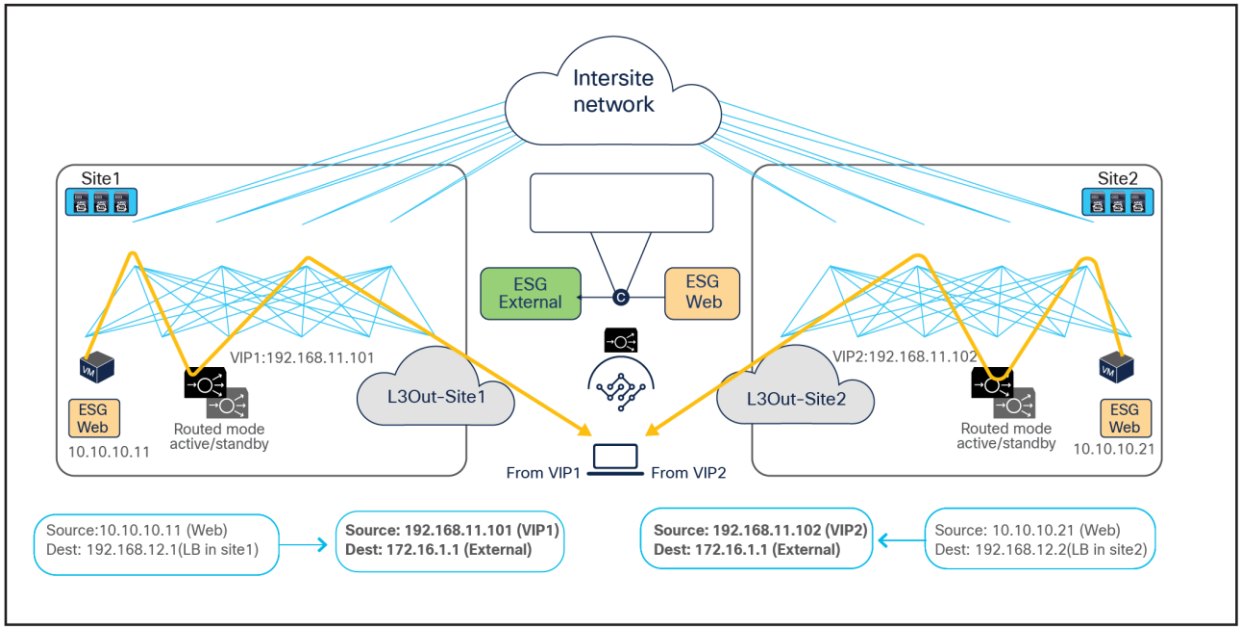


Figure 40.
Load balancer with SNAT inbound traffic flows (north-south)

The examples above show the load balancer and the real server as part of the same site, but in this use case they could also be deployed in different sites (Figure 41 and Figure 42). This is because the VIP, the SNAT IP, and the real servers' addresses are always reachable through regular forwarding from different sites. That said, the use of a local real-server farm is ideal in terms of traffic path optimization.

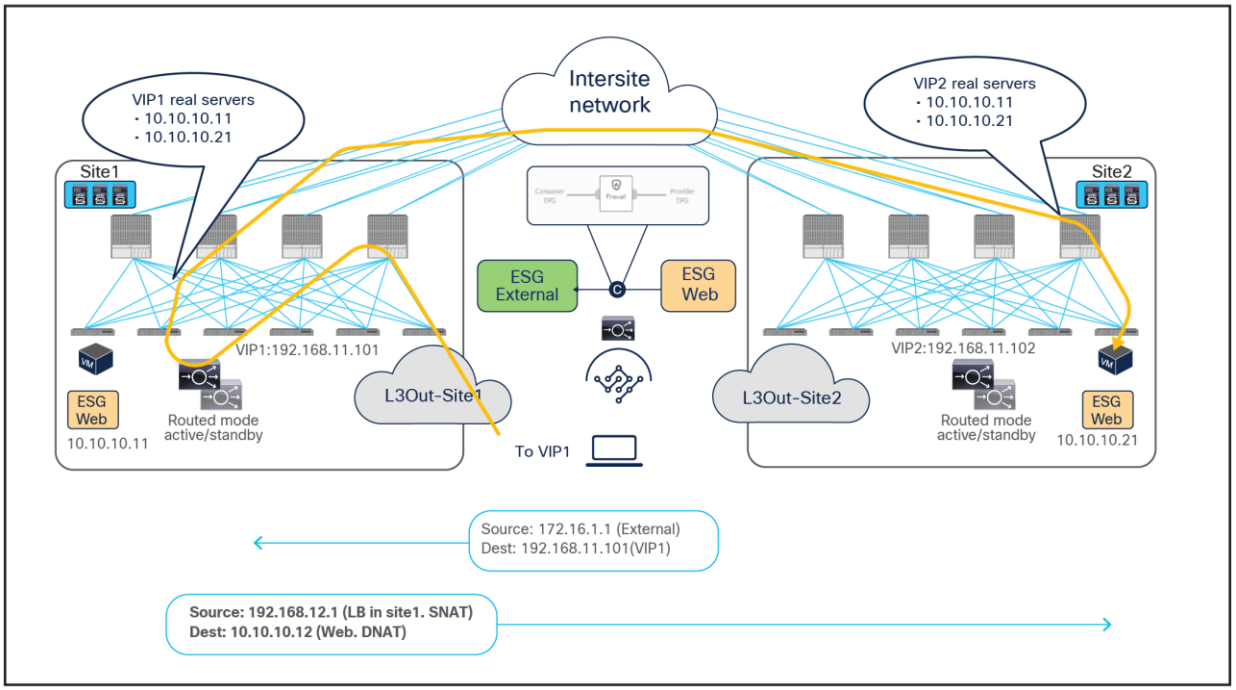


Figure 41.
Load balancer with SNAT inbound traffic flows (with VIP and real server in different sites)

© 2026 Cisco and/or its affiliates. All rights reserved.

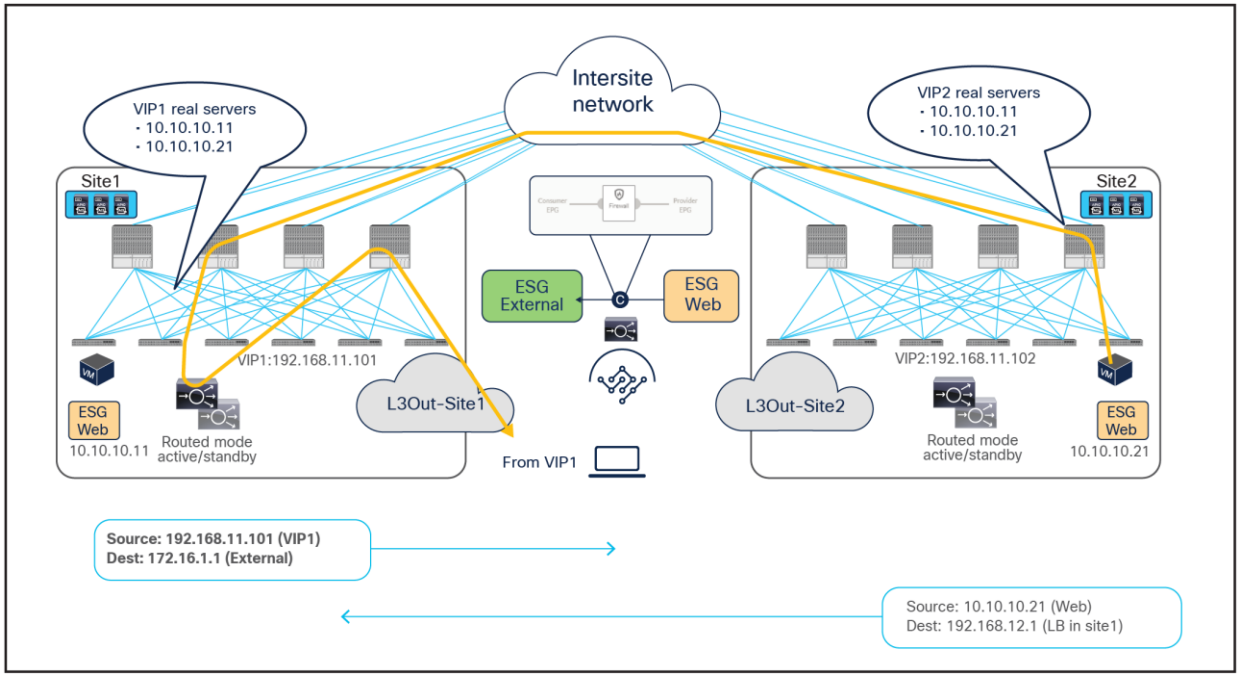


Figure 42.
Load balancer with SNAT outbound traffic flows (with VIP and real server in different sites)

East-west traffic use case (ESG-to-ESG and vzAny-to-ESG)

The figure below shows a typical Cisco ACI network design for east-west routed load-balancer insertion with SNAT. This design is similar to that for the north-south-routed load-balancer use case. In this example, the consumer Web ESG and the provider App ESG have a contract with a load-balancer service graph. Endpoints in the App ESG are real servers associated to the VIP on the load balancer.

As previously discussed for the north-south use case, the assumption is that each load balancer pair has assigned a unique VIP address that is part of the same service BD. If a service graph is not defined, each load balancer pair can use a unique VIP address in different service BD. Also, even if this example focuses on an intra-VRF contract, an inter-VRF contract for east-west communication is also supported.

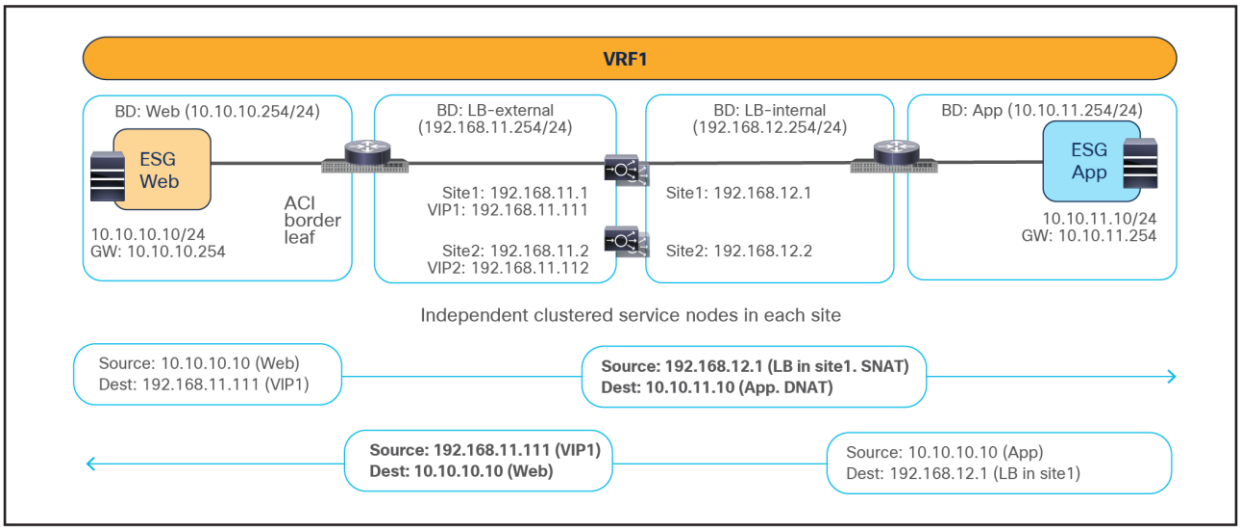


Figure 43.
Example of an east-west load balancer with a SNAT design

© 2026 Cisco and/or its affiliates. All rights reserved.

The figure below illustrates an example of east-west communication between a consumer ESG Web and a provider ESG App in a Multi-Site scenario where we have two connections: one is between the Web endpoint and the VIP (the frontend connection) and the other is between the load balancer and the real servers in the App ESG (the backend connection).

- The traffic originating from the Web endpoint is destined to the VIP, so it will reach the load balancer without requiring PBR as long as the VIP is reachable.
- The load balancer changes the destination IP to one of the real servers associated to the VIP. At the same time, the load balancer translates the source IP to the SNAT IP owned by the load balancer.
- The traffic is then sent back to the fabric and forwarded to the real server.

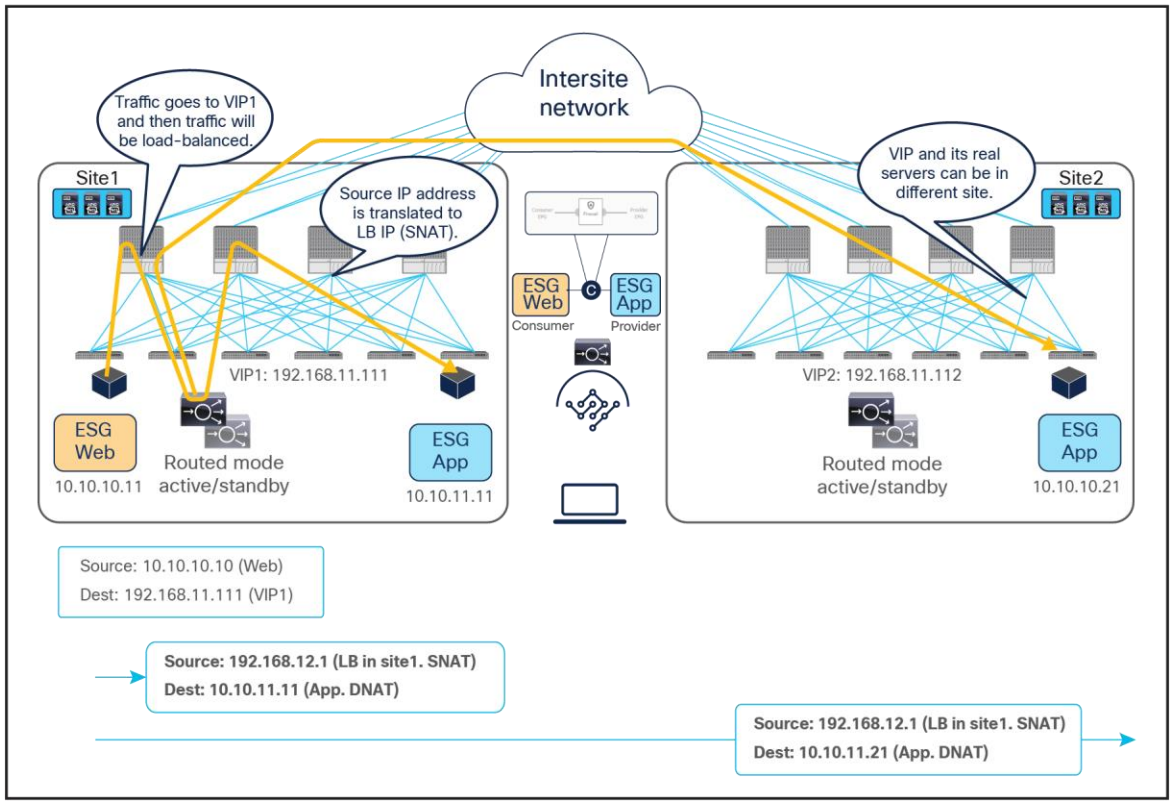


Figure 44.
Load balancer with SNAT incoming traffic flows (east-west)

For the provider-to-consumer traffic direction:

- The return traffic originated by the App real server is destined to the SNAT IP owned by the load balancer that took care of the incoming traffic; therefore, applying the PBR policy is not required for the return traffic either.
- The load balancer changes the source and destination IPs and sends the traffic back to the fabric.
- The traffic is forwarded back to the consumer endpoint.

© 2026 Cisco and/or its affiliates. All rights reserved.

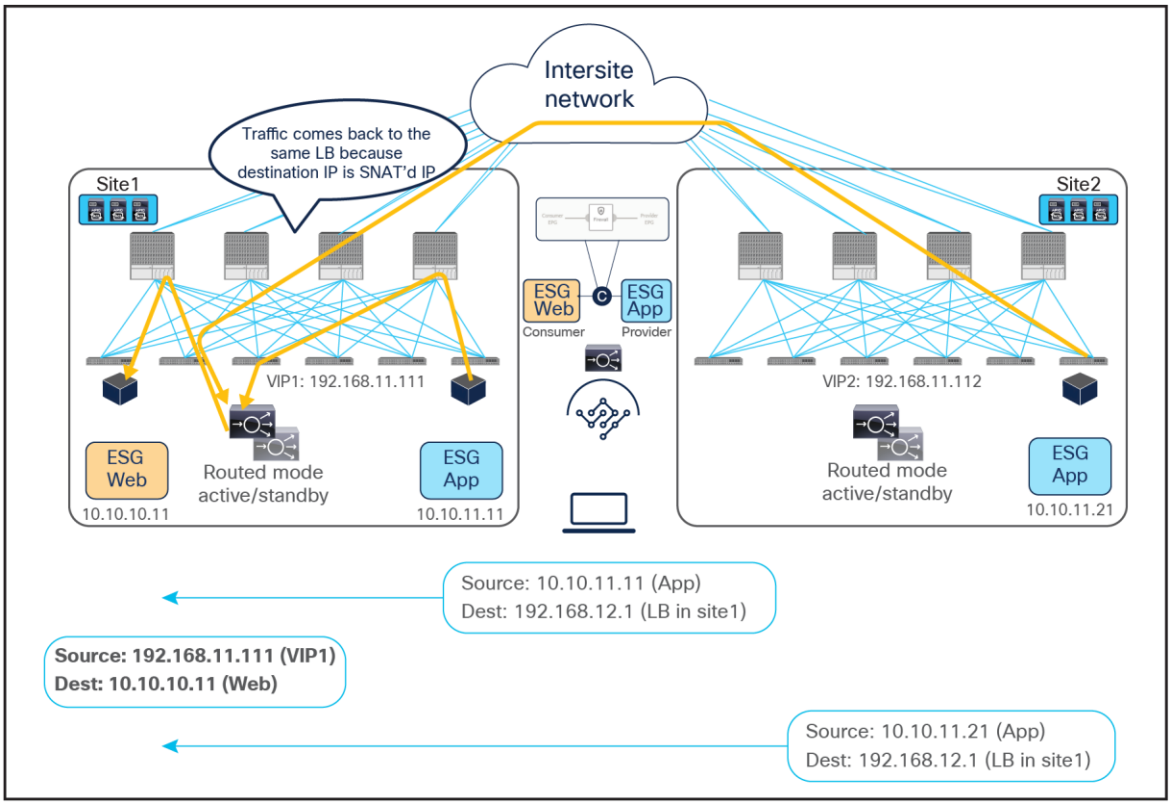


Figure 45.
Load balancer with SNAT return traffic flows (east-west)

Note: Though, in this example, the load balancer and the real server are in the same site, they can be in different sites, similar to the north-south-routed load-balancer insertion example earlier.

The use of SNAT is very handy to ensure that the return traffic goes back to the same load balancer that handled the incoming flow, therefore simplifying the design. However, a possibly undesirable consequence is that real servers lose visibility into the client's source IP address. When such visibility is a design requirement, you should avoid using SNAT on the load balancer, in order to ensure preserving the client's source IP. This mandates the introduction of PBR to properly steer the return traffic through the same load balancer that handled the first leg of the communication, as discussed in the next section.

Load balancer without SNAT (use of PBR for the return traffic)

In this deployment model, PBR is required for the return traffic between the real servers and the clients, because the load balancer does not perform SNAT for incoming traffic. The Incoming traffic flow destined to the VIP still does not require PBR and leverages basic forwarding.

There is an important considerations for deploying this design option with ACI Multi-Site:

- The load balancer and the real-server farm where traffic is load balanced must be deployed in the same site.

Although this section uses intra-VRF ESG-to-ESG contract as an example, inter-VRF contract and vzAny-to-ESG contract can be used as well because the policy enforcement behavior, redirecting traffic by the provider leaf, is same for both cases.

© 2026 Cisco and/or its affiliates. All rights reserved.

North-south traffic use case (ESG-to-ESG and vzAny-to-ESG)

The figure below shows a sample Cisco ACI network design for north-south-routed load balancer insertion without SNAT. The consumer L3Out ESG and the provider Web ESG have a contract with associated a service graph with PBR for the return traffic flow. Endpoints in the Web ESG are the real servers associated to the VIP of the load balancer. There can be multiple load balancers, which can be represented by multiple high-availability pairs deployed in separate sites.

The usual assumption here is that each load balancer gets assigned a unique VIP address that is part of the same BD and that Global Server Load Balancing (GSLB) is then used for load balancing traffic for a given application to multiple VIPs. Although the figure below illustrates an intra-VRF design, the definition of Web ESG and L3Out ESG in different VRFs is also a valid design. In this multi-VRF scenario, the service BD where the load balancer is connected must be in either the consumer or the provider VRF.

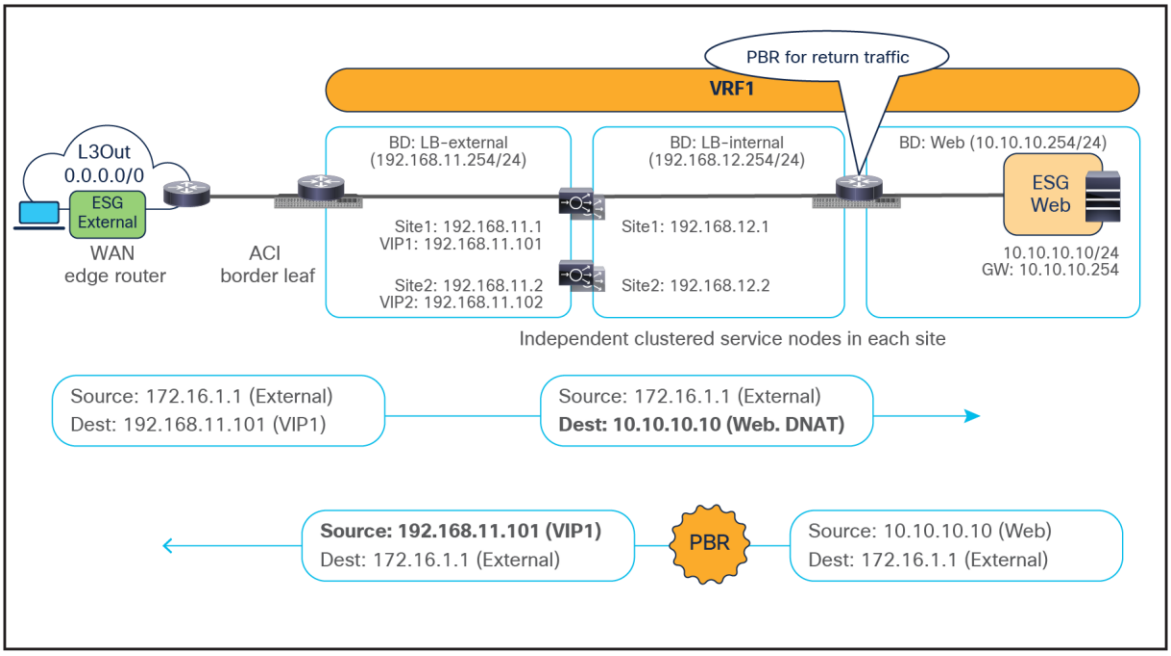


Figure 46.
Example of a north-south load-balancer design without SNAT

The figure below illustrates an example of an inbound traffic flow between the external network and an internal Web ESG in a Multi-Site deployment where we have two connections: one is between the external client and the VIP (the frontend connection) and the other is between the load balancer and the real servers that are part of the Web ESG (the backend connection).

- The incoming traffic originated from the external client and destined to the VIP is received on the L3Out connection of a given site, and reaches the load balancer without requiring PBR as long as the VIP is reachable (this is basic intra-site or inter-site forwarding).
- The load balancer changes the destination IP to one of the real servers associated to the VIP, but leaves unaltered the source IP addresses (representing the external client) and forwards the traffic back to the fabric.
- The traffic is then forwarded to the real server, which must be deployed in the local site. As clarified below, this is needed to ensure that PBR can steer the return flow to the same load balancer that handled the incoming traffic.

© 2026 Cisco and/or its affiliates. All rights reserved.

As usual, the suboptimal inbound traffic shown for communicating with the VIP2 could be avoided by leveraging host-route advertisement to optimize the traffic path for ingress communication or by taking the VIP addresses of the load balancers deployed in separate sites from different IP subnets. Note that the service BD must be L2-stretched when using a service graph. Thus, multiple IP subnets must be provisioned for the same service BD to use VIP addresses from different IP subnets.

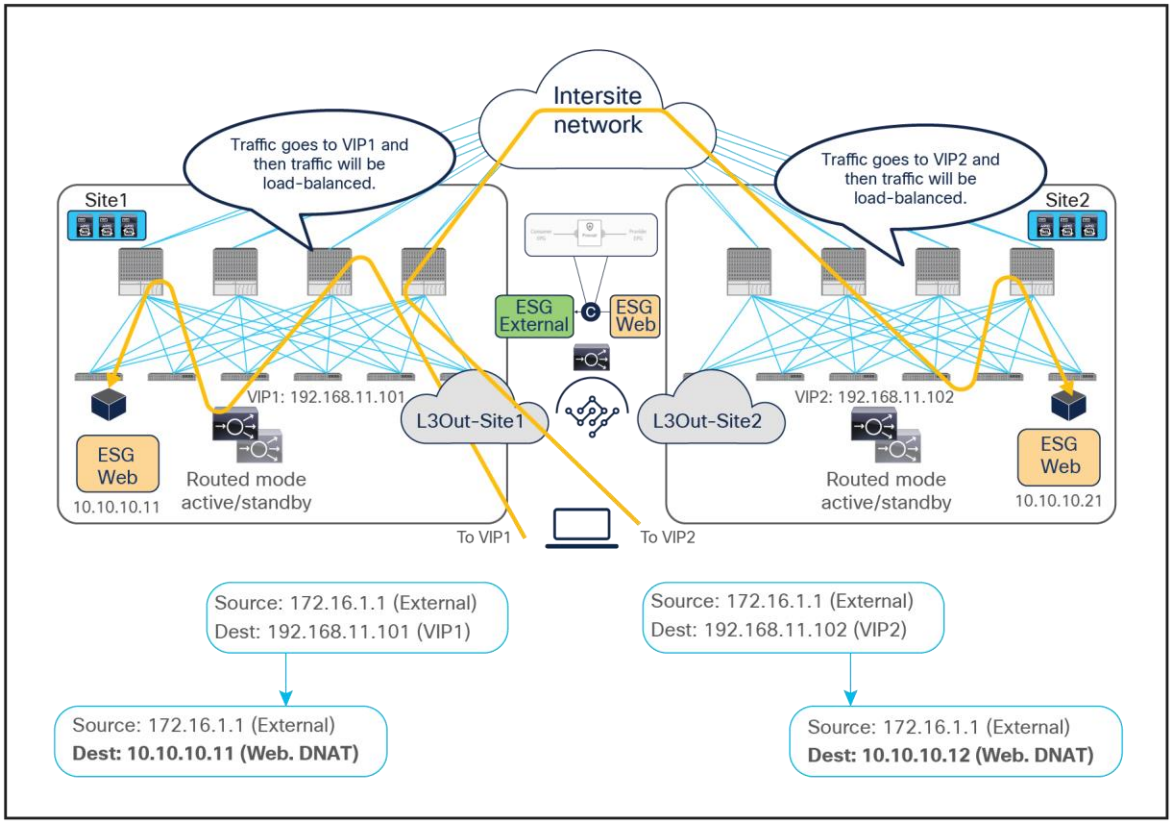


Figure 47.
Load balancer without SNAT inbound traffic flows (north-south)

For the outbound direction, the traffic is destined to the original client's IP address connection; therefore, PBR is required to steer the return traffic back to the load balancer. Otherwise, the external client would receive the traffic with the source IP being the real server's IP instead of the VIP. Such traffic will be dropped because the external client did not initiate traffic to the real server IP.

- The Web ESG sends traffic back to the external client. The PBR policy is always applied on the provider leaf where the Web endpoint is connected, so it can only redirect the traffic to a local load-balancer. This is the reason why the VIP and the real servers must be in the same site in this deployment model. The provider leaf can always resolve the consumer ESG class ID and applies the PBR policy in this case, because the IP prefix identifying the external clients and associated to the External ESG is statically configured.
- The load balancer changes only the source IP address to match the locally defined VIP and sends the traffic back to the fabric.
- The traffic is forwarded toward the external client leveraging, by default, a local L3Out connection.

© 2026 Cisco and/or its affiliates. All rights reserved.

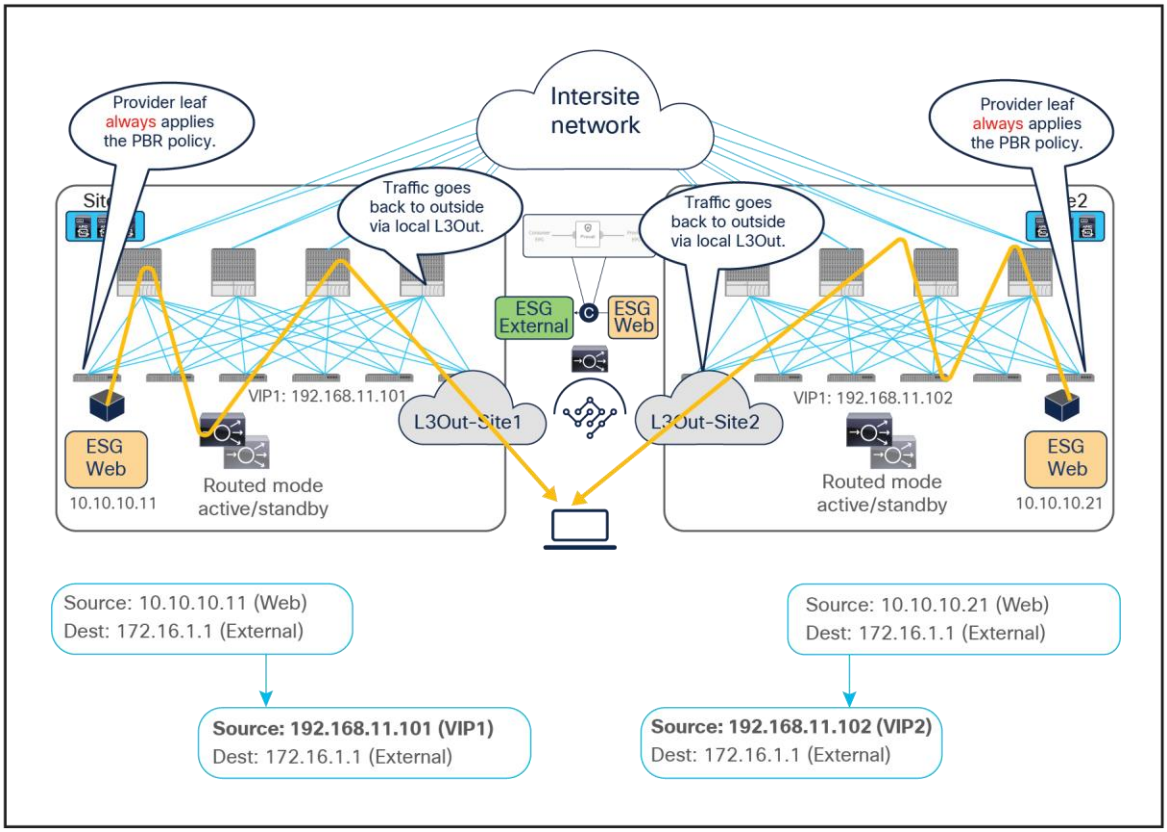


Figure 48.
Load balancer without SNAT outbound traffic flows (north-south)

Though there may be an “asymmetric” use of the L3Out connections (that is, for VIP2, inbound traffic uses L3Out in site1, whereas outbound traffic is sent through L3Out in site2), there is always a “fully symmetric” use of the same service node for both legs of the communication as long as the load balancer and the real servers are deployed in the same site. Otherwise, the return traffic would be redirected to the load balancer in a different site and lose traffic symmetricity.

Figure 49 and Figure 50 illustrate an example of this problem: the load balancer in site1 has both local site endpoint 10.10.10.11 and remote site endpoint 10.10.10.21 as real servers associated to VIP1. If the incoming traffic to VIP1 is load balanced to 10.10.10.21 in site2, the PBR policy for the return traffic enforced on the provider leaf in site2 would redirect the traffic to the local load balancer, creating traffic asymmetry.

© 2026 Cisco and/or its affiliates. All rights reserved.

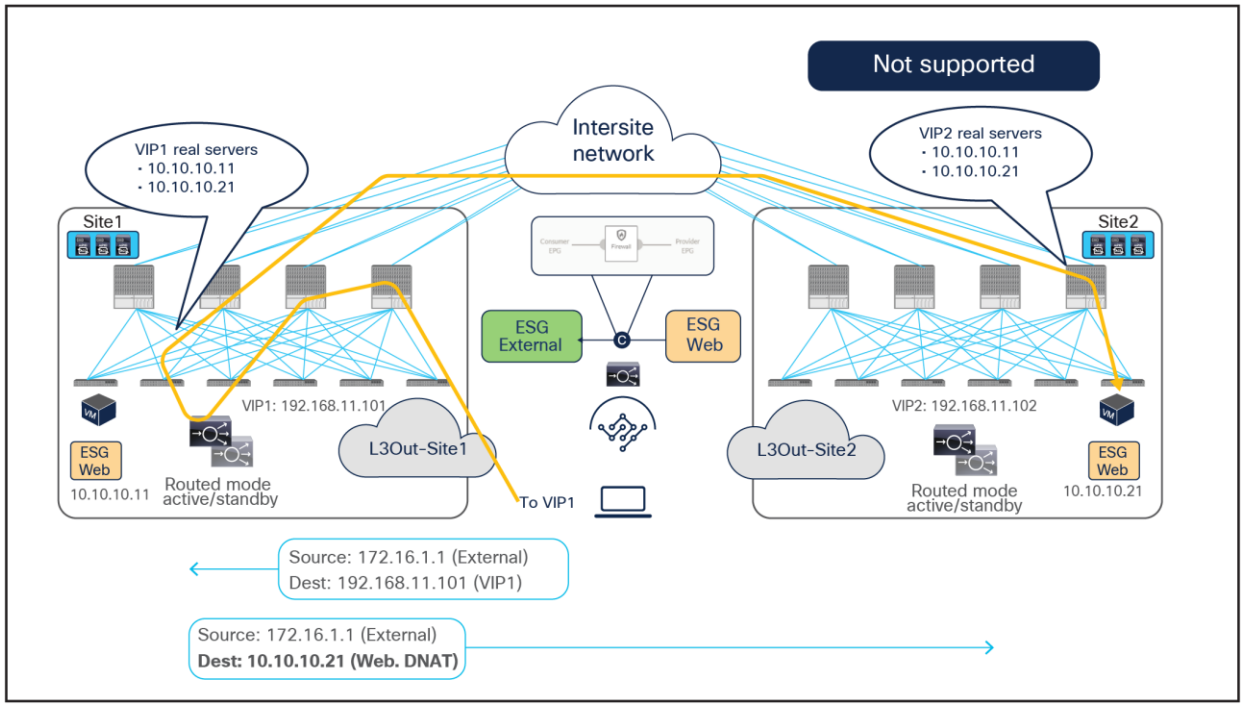


Figure 49.
Load balancer without SNAT inbound traffic flows (Having the VIP and real server in different sites is not supported).

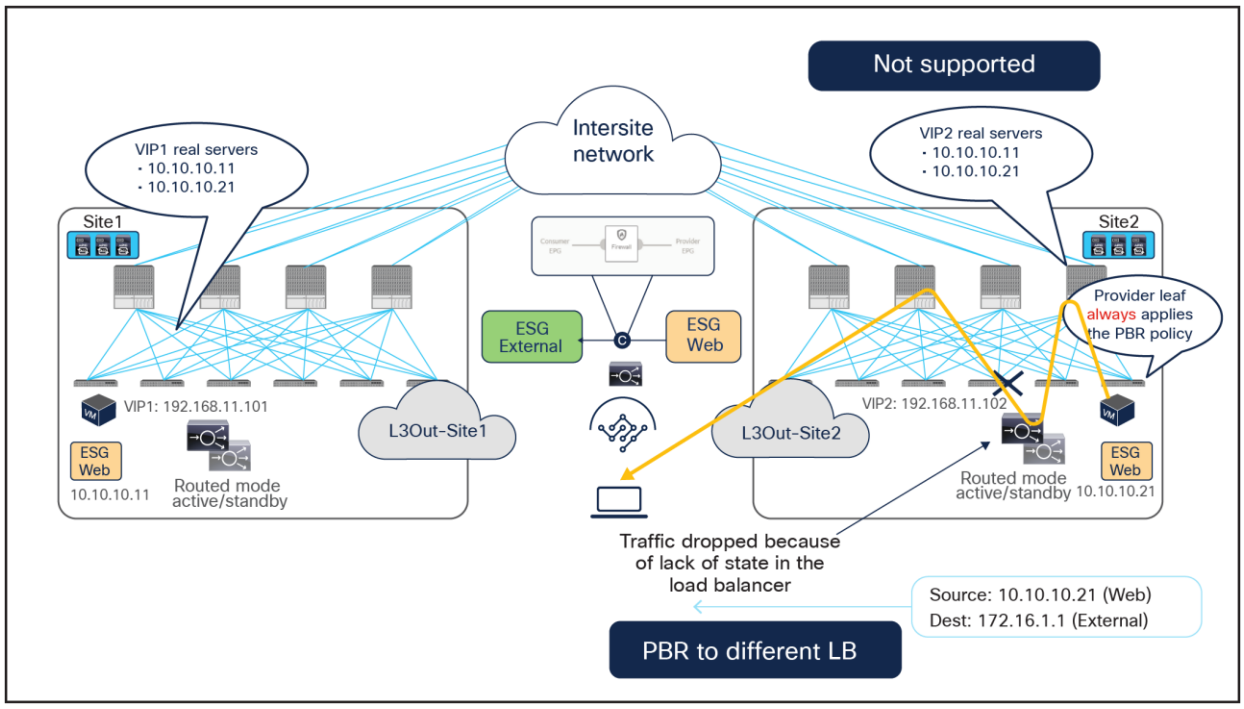


Figure 50.
Load balancer without SNAT inbound traffic flows (Having the VIP and real server in different sites is not supported).

© 2026 Cisco and/or its affiliates. All rights reserved.

East-west traffic use case (ESG-to-ESG and vzAny-to-ESG)

The figure below shows a typical Cisco ACI network design for east-west-routed load-balancer insertion without SNAT. This design is similar to that for the north-south load-balancer use case previously discussed. The consumer Web ESG and the provider App ESG have a contract with a load-balancer service graph. The endpoints in App ESG are real servers associated to the VIP on the load balancer and must be connected in the same site where the VIP is active.

The assumption here is that the VIP is in the same BD, each load balancer pair has a unique VIP address, and Global Server Load Balancing (GSLB) is used for load balancing to multiple VIPs.

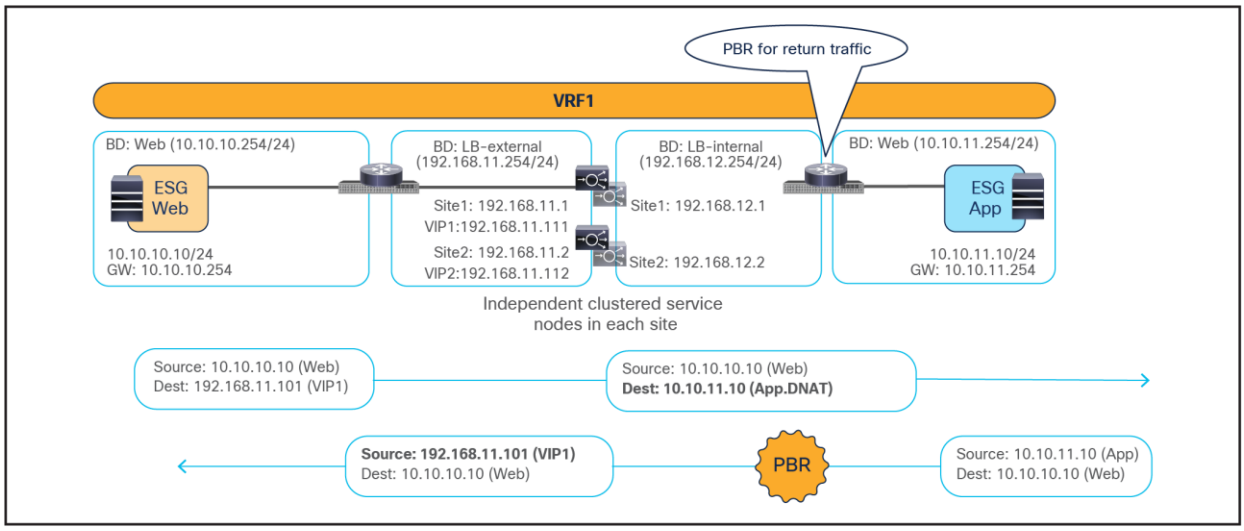


Figure 51.
Example of east-west load balancer design without a SNAT

The figure below illustrates an example of east-west communication between a consumer ESG Web and a provider ESG App in an ACI Multi-Site architecture. Here we have two connections: one is between a Web endpoint and the VIP (the frontend connection), and the other is between the load balancer and the real servers in the App ESG (the backend connection).

- The consumer-to-provider traffic is destined to the VIP, so the traffic reaches the load balancer without the need of PBR as long as the VIP is reachable. Notice how the VIP could be locally deployed or available in a remote site.
- The load balancer changes the destination IP to one of the real servers associated to the VIP, but it does not alter the source IP (since SNAT is not enabled). The traffic is then sent back to the fabric.
- The traffic is forwarded to the real server, which must be connected in the same site.

In the example below, the Web endpoint accesses the VIP addresses of both of the load balancers deployed in the local and remote sites, which then redirect traffic to local server farms.

© 2026 Cisco and/or its affiliates. All rights reserved.

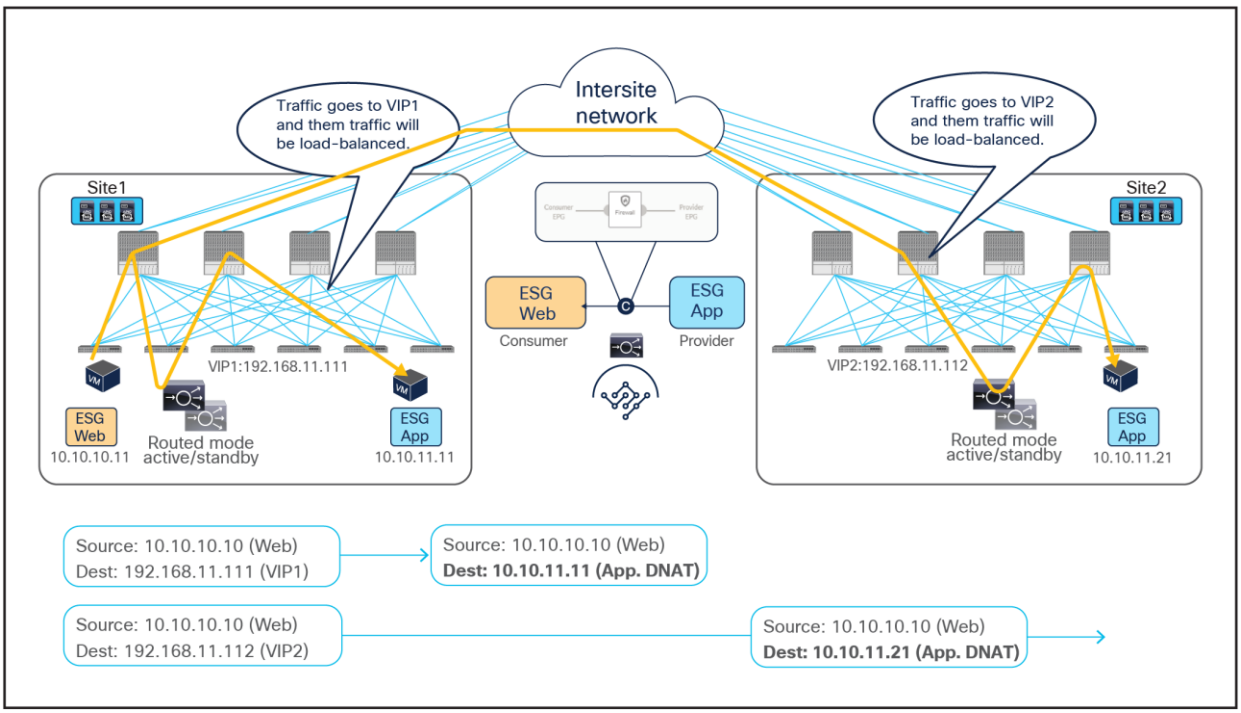


Figure 52.
Load balancer without SNAT incoming traffic flows (east-west)

Because the return traffic is destined to the original source IP of the Web endpoint, PBR is required to force the return traffic through the same load balancer used for the consumer-to-provider direction.

- The App endpoint sends the traffic back to the Web endpoint PBR, and the policy is applied on the provider leaf node where the App endpoint is connected. This ensures that the return traffic is steered toward the same load balancer because the load balancer and the real server must always be in the same site in this deployment model. Otherwise the return traffic would be redirected to a different load balancer from the one used for the first leg of the communication, thus causing loss of traffic symmetricity (similar to what was shown previously, in Figure 49 and Figure 50). The provider leaf cannot resolve the destination class ID if the consumer Web ESG is not based on IP address and the consumer endpoint IP is not yet learned. Hence, the conversational learning is required, similarly to the ESG-to-ESG PBR use case for firewall insertion previously discussed.
- The load balancer changes only the source IP address to match the locally defined VIP and sends the traffic back to the fabric.

The traffic is forwarded toward the Web endpoint that could be locally connected or deployed in a remote site.

© 2026 Cisco and/or its affiliates. All rights reserved.

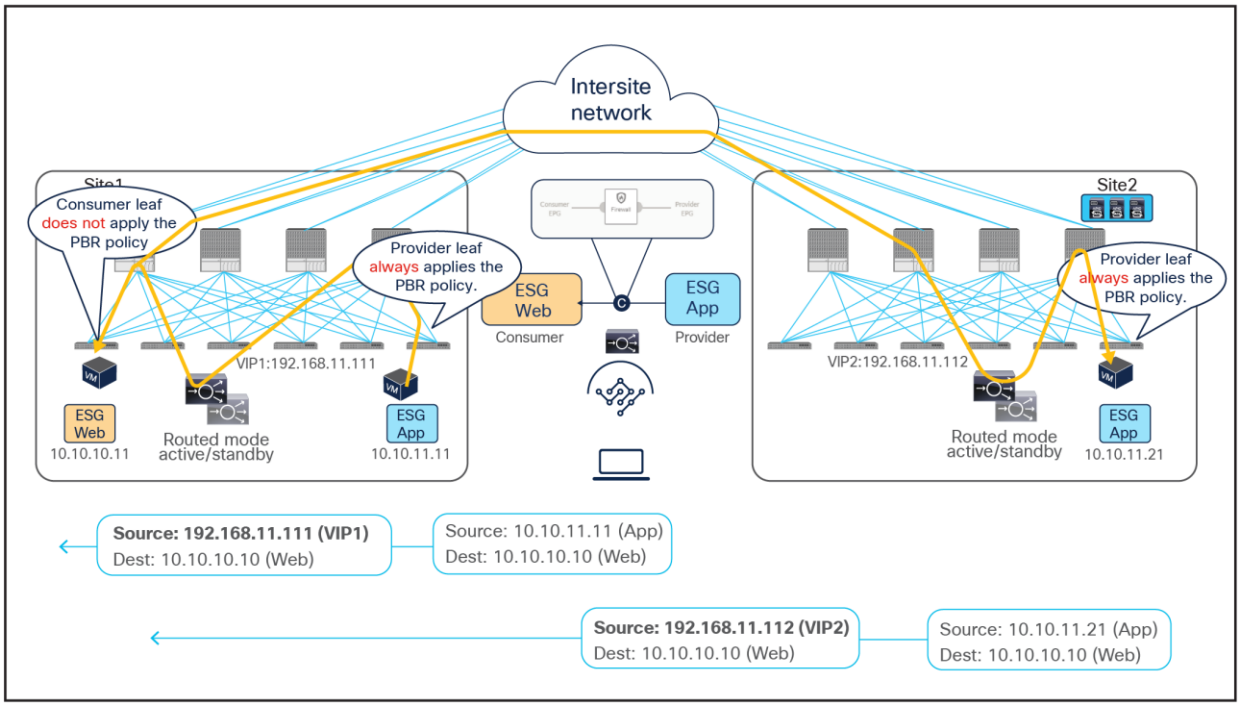


Figure 53.
Load balancer without SNAT return traffic flows (east-west)

PBR to a firewall and a load balancer without SNAT (two nodes service graph)

This section covers a two-node firewall and load-balancer insertion use case for north-south and east-west communication. PBR is enabled for both directions (provider-to-consumer traffic and vice versa) to redirect traffic to the firewall, but PBR for the load balancer is only needed for the provider-to-consumer direction (since SNAT is not configured).

The same specific design considerations mentioned in the previous section for the load balancer only scenario are still valid here; thus, it is mandatory for the load balancer and the real servers to reside in the same site.

Though the example we present in Figure 54, below, has the firewall as the first service function and the load balancer without SNAT as the second, other service-function combinations or sequences are also possible, as, for example, those in the bulleted list below:

- The first service function is the firewall; the second is the load balancer with SNAT.
- The first service function is the load balancer with SNAT; the second is the firewall.
- The first service function is the load balancer without SNAT; the second is the firewall.
- The first service function is the firewall; the second is the IPS.

Note: As of Cisco ACI Release 6.1(4), a Multi-Site service graph can contain up to two service functions when defined in a Multi-Site template (that is, a template that can also be used to provision objects stretched across sites) and up to five service functions in case of autonomous templates. Also, a service graph with two or more nodes is not supported for vzAny-to-vzAny PBR use case.

© 2026 Cisco and/or its affiliates. All rights reserved.

North-south traffic use case

The figure below shows a sample Cisco ACI network design for a two-node PBR service chain (a firewall and a load balancer without SNAT) applied to north-south-routed communication. A contract with an associated service graph with PBR is applied between the consumer External ESG and the provider Web ESG. The endpoints in the Web ESG are real servers associated to the VIP on the load balancer. As always, the firewall and load balancer services are deployed as a distributed set of highly available service nodes deployed in each site.

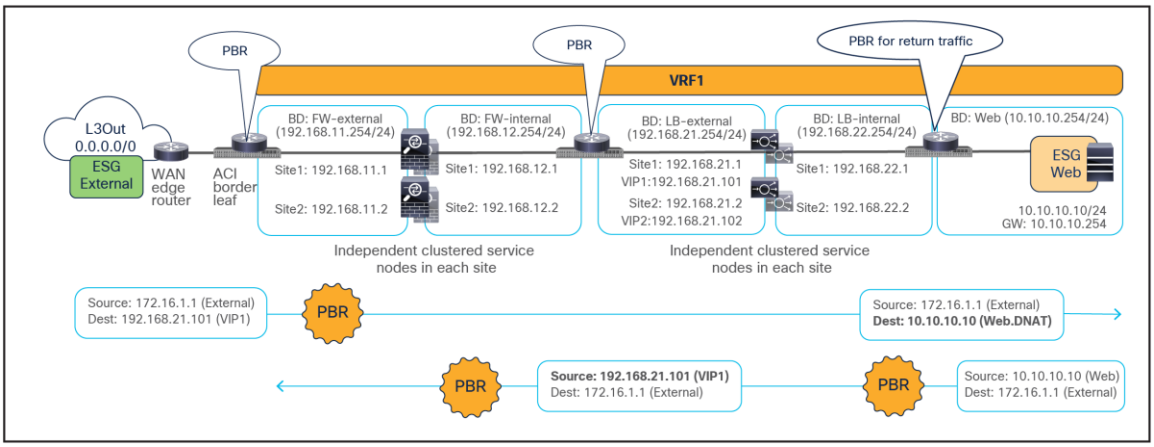


Figure 54.
Sample design of a north-south firewall with PBR and a load balancer without SNAT

Figure 55 and Figure 56 illustrate an example of an inbound traffic flow between the external network and an internal Web ESG in an ACI Multi-Site deployment, where we have two connections: one is between the external network and the VIP (the frontend connection), and the other is between the load balancer and the real servers part of the Web ESG (the backend connection).

- Traffic originating from a client in the external network is destined to the VIPs (one in each site), thus the traffic is forwarded to the leaf nodes where the load balancers are connected as long as the VIPs are reachable. Notice how the VIPs could be locally deployed or reachable in a remote site.
- The PBR policy is applied on the load-balancer leaf nodes (since they represent the provider leaf nodes for the north-south traffic from the external network to the VIPs) and redirects the traffic to the first service function, which is the firewall. As previously discussed, the consumer leaf nodes don't apply the policy because of the "redirect override" flag. Since the PBR policy is applied on the load-balancer leaf nodes, the consequence is that all inbound north-south flows will always be redirected to the firewall in the site where the load balancer with the destination VIP is located, independently from the specific L3Out where the inbound flow is received (as Figure 55 clearly highlights).
- The traffic is inspected by the firewalls; if allowed, it is then sent back to the fabric and reaches the VIPs residing on the load balancers.
- The load balancers change the destination IP to one of the real servers associated to the VIP and send the traffic back to the fabric (the load balancers do not change the source IP address since SNAT is not enabled).
- The traffic is forwarded to the real server destination, which must be deployed locally in each site.

© 2026 Cisco and/or its affiliates. All rights reserved.

Notice how redirection for inbound flows happens only to one service function (the firewall), since, from the firewall to the real server, all that is needed is just regular traffic forwarding. Also, the suboptimal inbound traffic shown in Figure 55 can be avoided by leveraging host-route advertisement to optimize the traffic path for ingress communication when the VIPs of the load balancer belong to a stretched BD.

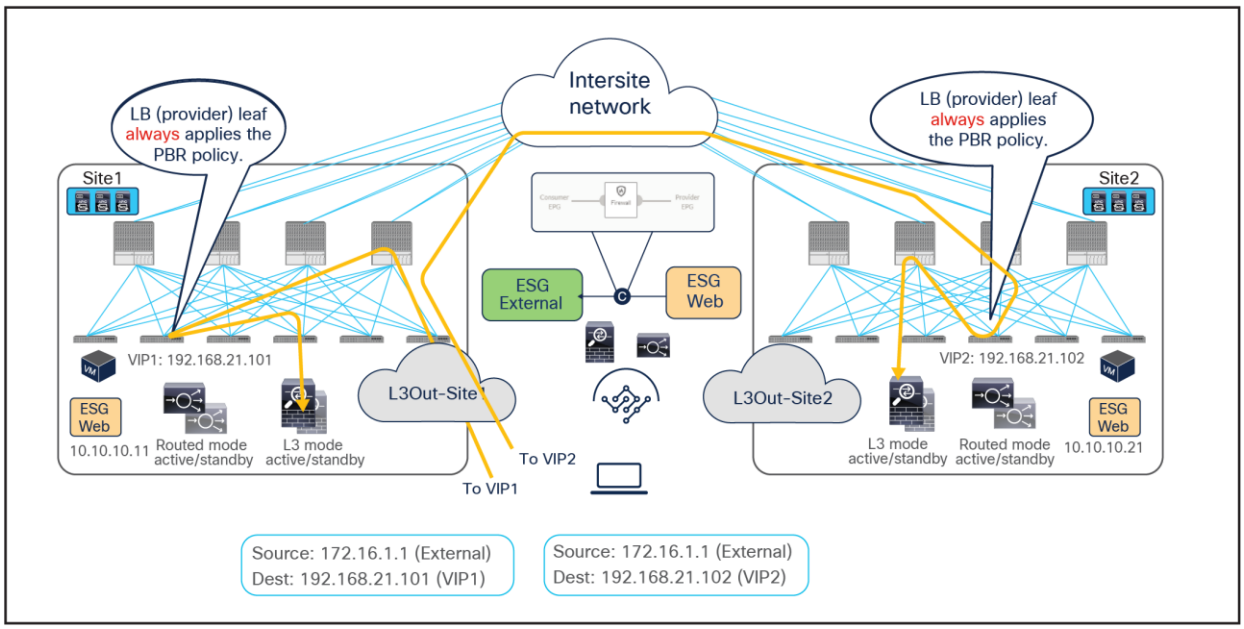


Figure 55.
Firewall with PBR and load balancer without SNAT inbound traffic flows (external-client to VIP)

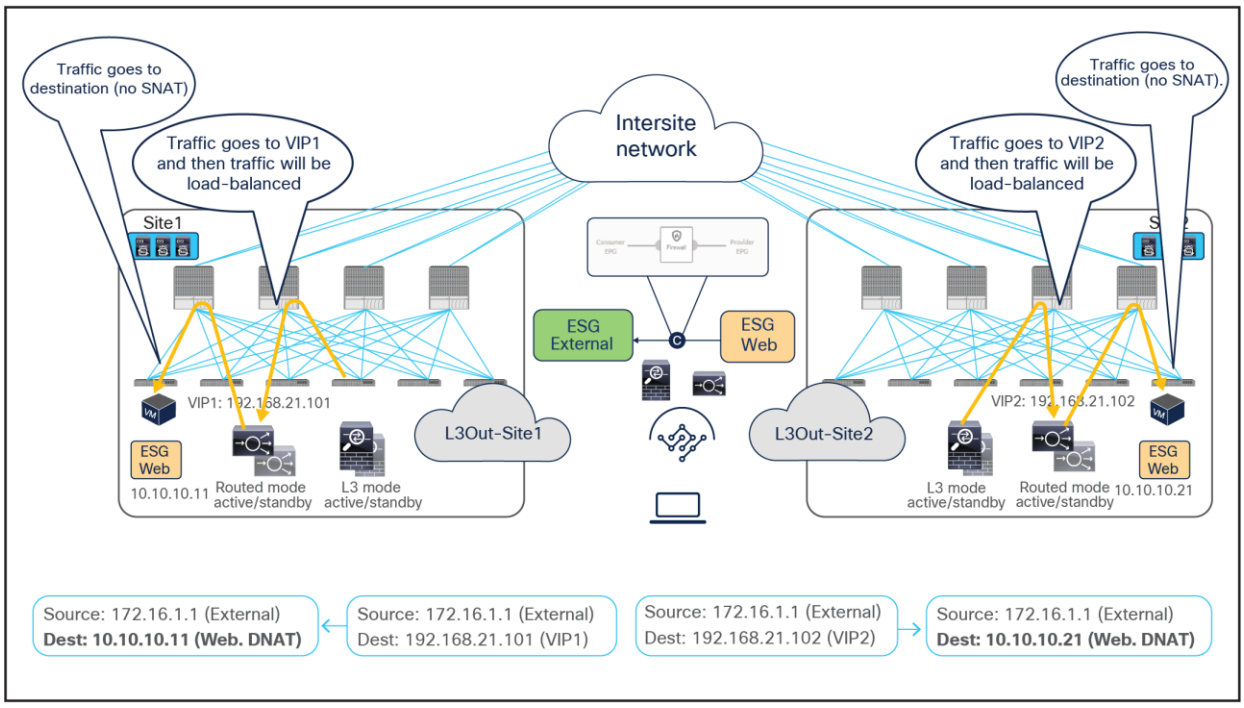


Figure 56.
Firewall with PBR and load balancer without SNAT inbound traffic flows (VIP to real-servers)

© 2026 Cisco and/or its affiliates. All rights reserved.

As previously discussed, since the return traffic flows are destined to the original IP address of the external client, PBR is required to steer the flow through the load balancers.

- The PBR policy is associated to a contract between the Web ESG and the External ESG, therefore it is applied on the provider leaf nodes where the Web endpoints are connected and redirects the traffic to the local load balancer. Once again, this is the reason the VIP and the server farm (Web ESG) must be deployed in the same fabric.
- Once the load balancers receive the traffic, they change the source IP to the VIP and forward the traffic back to the ACI fabric.
- At this point the second redirection associated to the contract between the Web ESG and the External ESG kicks in on the load balancer leaf nodes where the load balancers' external interfaces are connected, and the traffic is forwarded to the firewalls.
- After the firewalls have applied their locally configured security policies, the traffic is forwarded back to the fabric so it can reach the external client through the local L3Out connections.

Therefore, in the outbound direction redirection must happen twice, first to the load-balancer service and then to the firewall service.

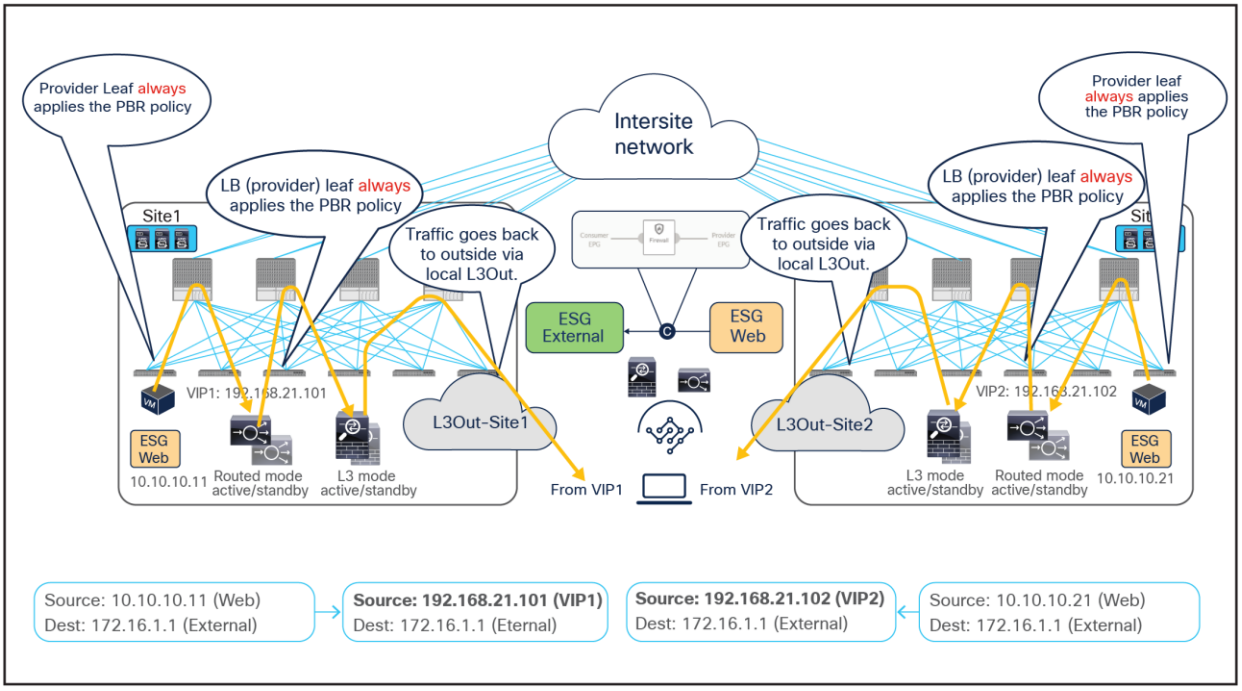


Figure 57.
Firewall with PBR and load balancer without SNAT outbound traffic flows (north-south)

© 2026 Cisco and/or its affiliates. All rights reserved.

East-west traffic use case

The figure below shows a sample Cisco ACI network design for a two-node service chain, this time applied to the east-west traffic use case. This design is similar to the one for the north-south use case previously discussed. The contract with an associated two-nodes (firewall and load balancer) service graph with PBR is now applied between a consumer Web ESG and a provider App ESG. The endpoints in the App ESG are real servers associated to the VIP of the load balancer.

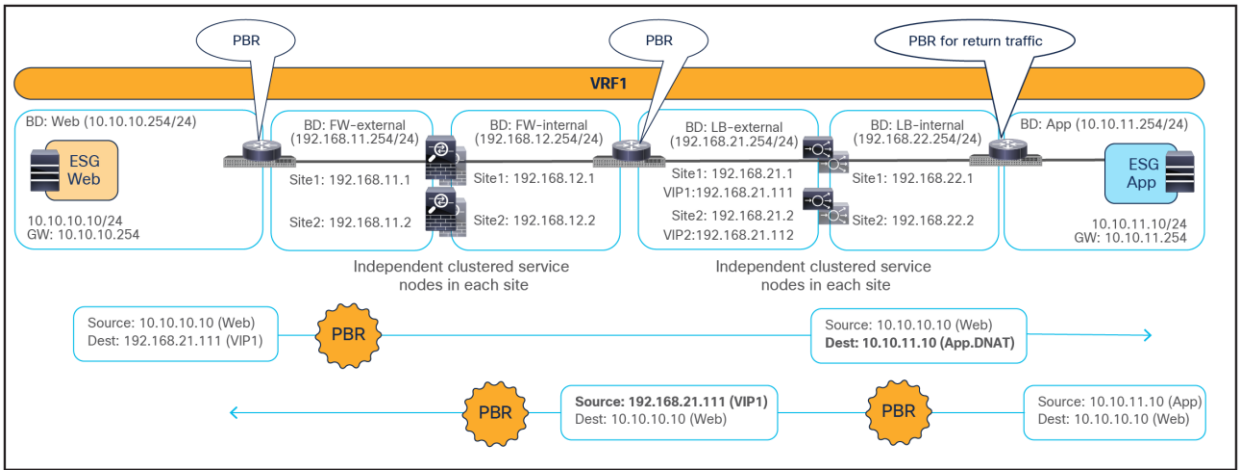


Figure 58.
Sample design of an east-west firewall with PBR and a load balancer without SNAT

Figure 59 and Figure 60 illustrate an example of east-west communication between the consumer ESG Web and the provider ESG App in an ACI Multi-Site deployment where we have two connections: one is between a Web endpoint and the VIP (the frontend connection), and the other is between the load balancer and the real servers in the App ESG (the backend connection). As usual for all the scenarios not leveraging SNAT, the load balancer and the real servers must be deployed in the same site.

- The traffic originating from the consumer (Web) endpoint is destined to the VIP, so it reaches the leaf nodes where the load balancers are connected as long as the VIPs are reachable. In the example below, the same consumer ESG accesses two different VIPs (VIP1 and VIP2) defined in site1 and site2.
- The PBR policy is applied on the load-balancer leaf nodes and redirects traffic to the local firewalls (because they represent the provider leaf nodes for the traffic from the consumer to the VIP). As previously explained, this mandates the definition of an IP prefix under the consumer EPG identifying the endpoints part of that security group.
- The firewalls apply their security policies and then send the traffic back to the fabric toward the VIPs.
- The load balancers change the destination IPs to one of the real servers associated to the VIPs and forward the traffic to the destinations. In this example, the load balancers do not perform SNAT and hence do not alter the source IP address.

© 2026 Cisco and/or its affiliates. All rights reserved.

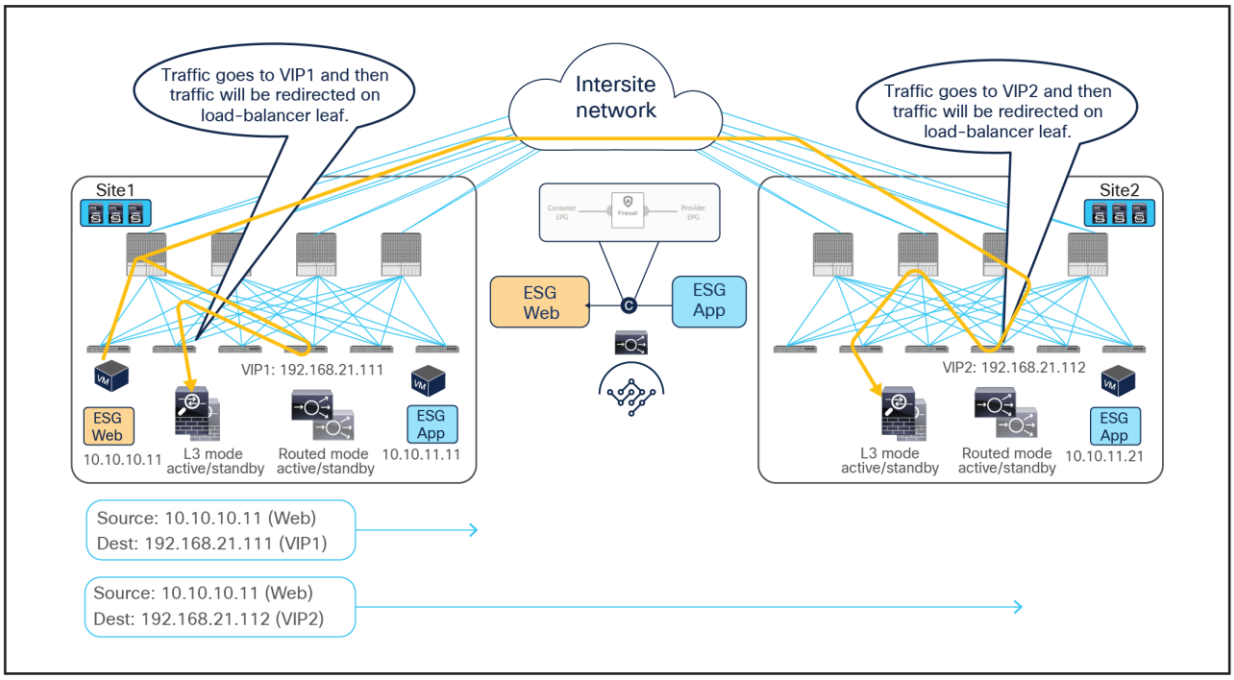


Figure 59.
Firewall with PBR and load balancer without SNAT east-west traffic flows (client to VIPs)

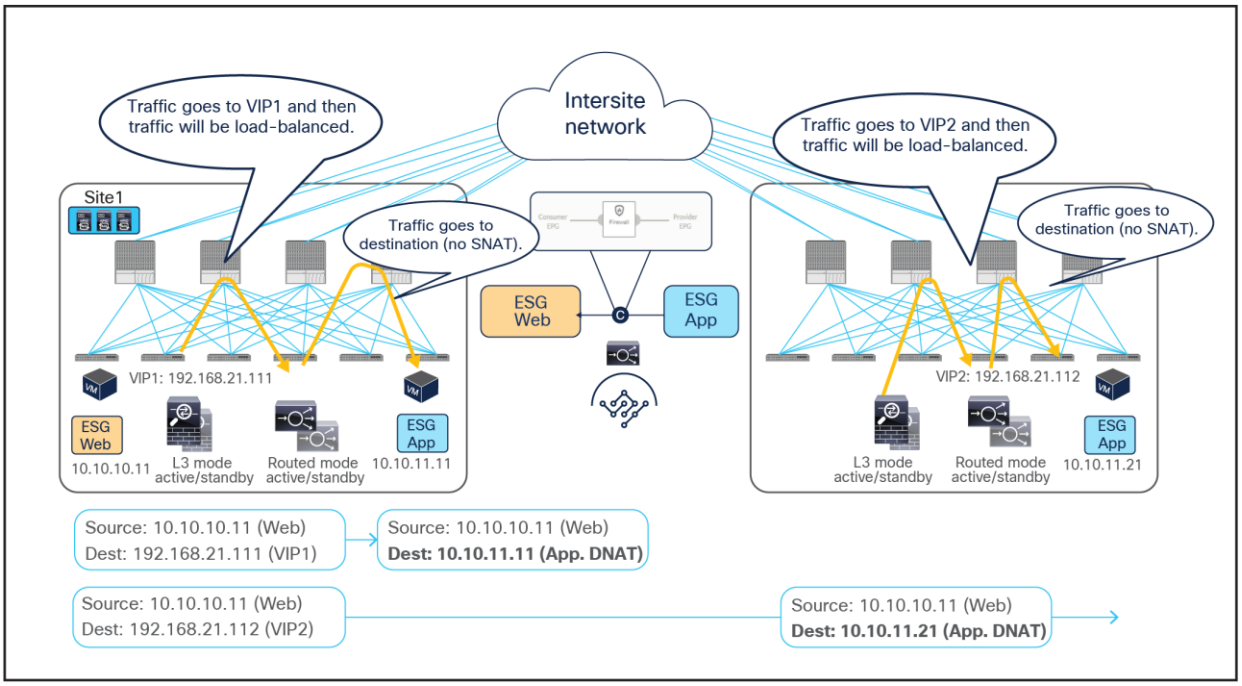


Figure 60.
Firewall with PBR and load balancer without SNAT east-west traffic flows (VIP to real-server)

© 2026 Cisco and/or its affiliates. All rights reserved.

Because the return traffic is destined to the original source IP (the Web endpoint), PBR is required to steer the return traffic to the load balancers.

- The provider endpoints (real servers) originate traffic destined to the consumer endpoint. The PBR policy gets applied on the leaf nodes where the real servers are connected, because they represent the provider leaf nodes for the east-west communication between the App (provider) endpoint and the Web (consumer) endpoint. The traffic gets steered to the load balancers, which must be located in the same site.
- The load balancers change the source IP to match the VIP address and send the traffic back to the ACI fabric.
- Another PBR policy is then applied on the load-balancer leaf nodes to redirect the traffic to the firewalls.
- The firewalls perform their security policy enforcement and send the traffic back to the fabric.
- The traffic is forwarded to the consumer (Web) endpoint, which can be located in the local site or in a remote site (as highlighted in Figure 62).

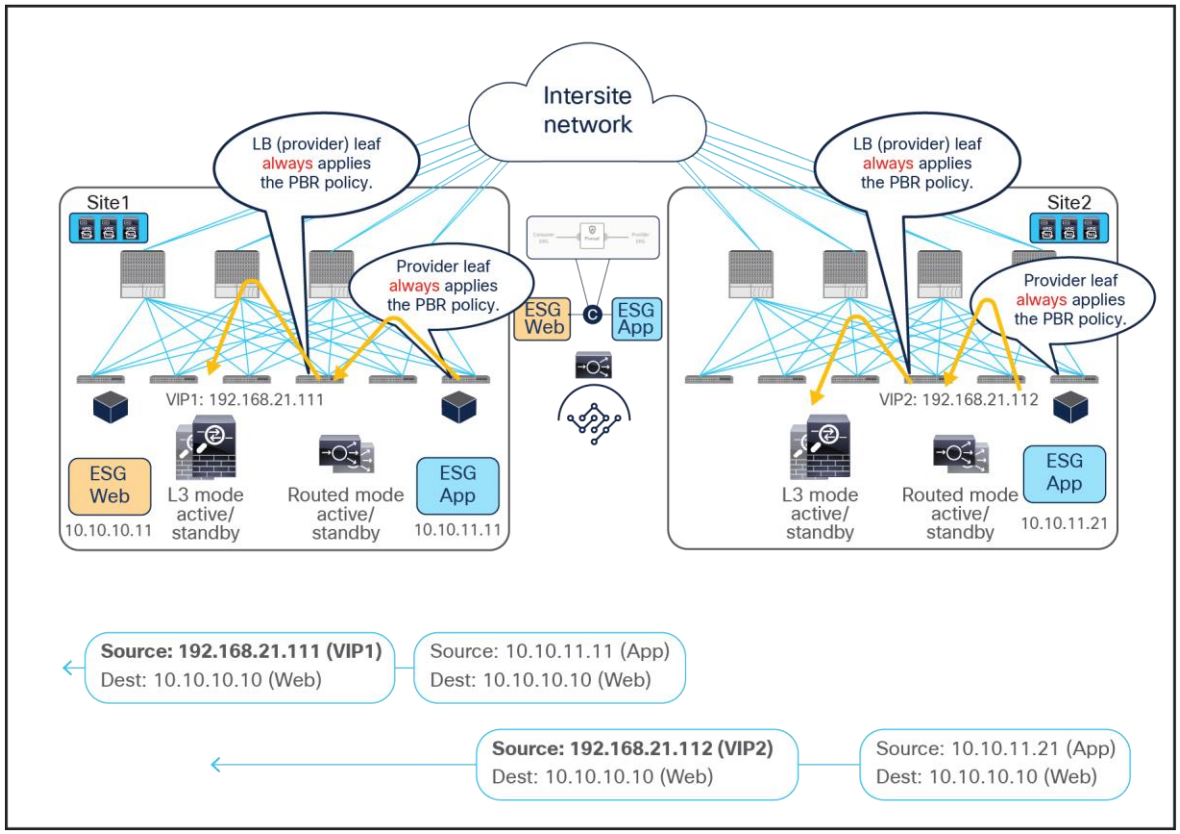


Figure 61.
Firewall with PBR and load balancer without SNAT east-west traffic flows (real servers to firewalls)

© 2026 Cisco and/or its affiliates. All rights reserved.

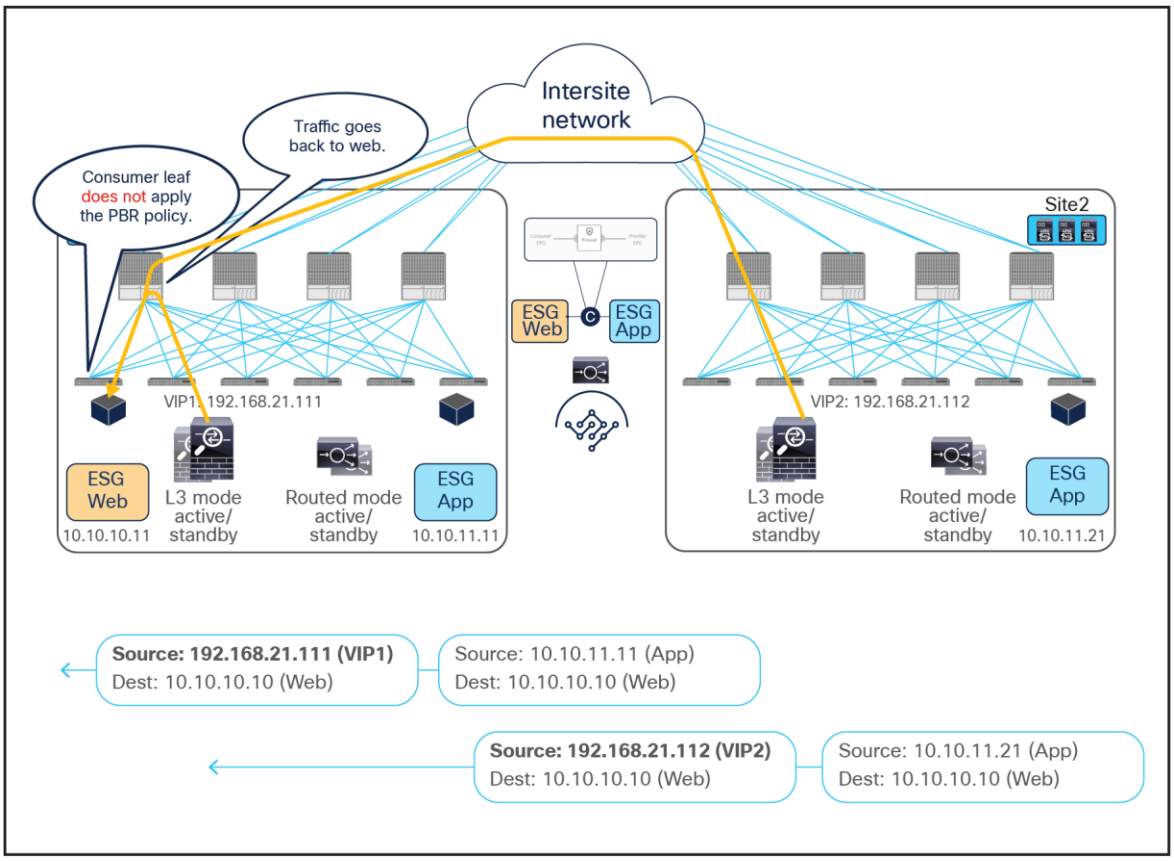


Figure 62.
Firewall with PBR and load balancer without SNAT east-west traffic flows (firewalls to client)

Advanced design examples

This section covers several examples of designs that leverage combinations of the supported use cases explained in the previous sections.

Example 1: Insert firewall for most (but not all) inter-ESG traffic in a VRF.

The figure below illustrates a sample design that has the following requirements:

- Application-centric (Multiple ESGs are configured as part of the same BD and IP subnet.)
- All of inter-ESG traffic needs to be inspected by a firewall except specific ESG-to-ESG combinations (in this example, App-to-DB traffic needs to be permitted at the fabric level without being redirected to the firewall).
- Intra-ESG communication should happen freely.

© 2026 Cisco and/or its affiliates. All rights reserved.

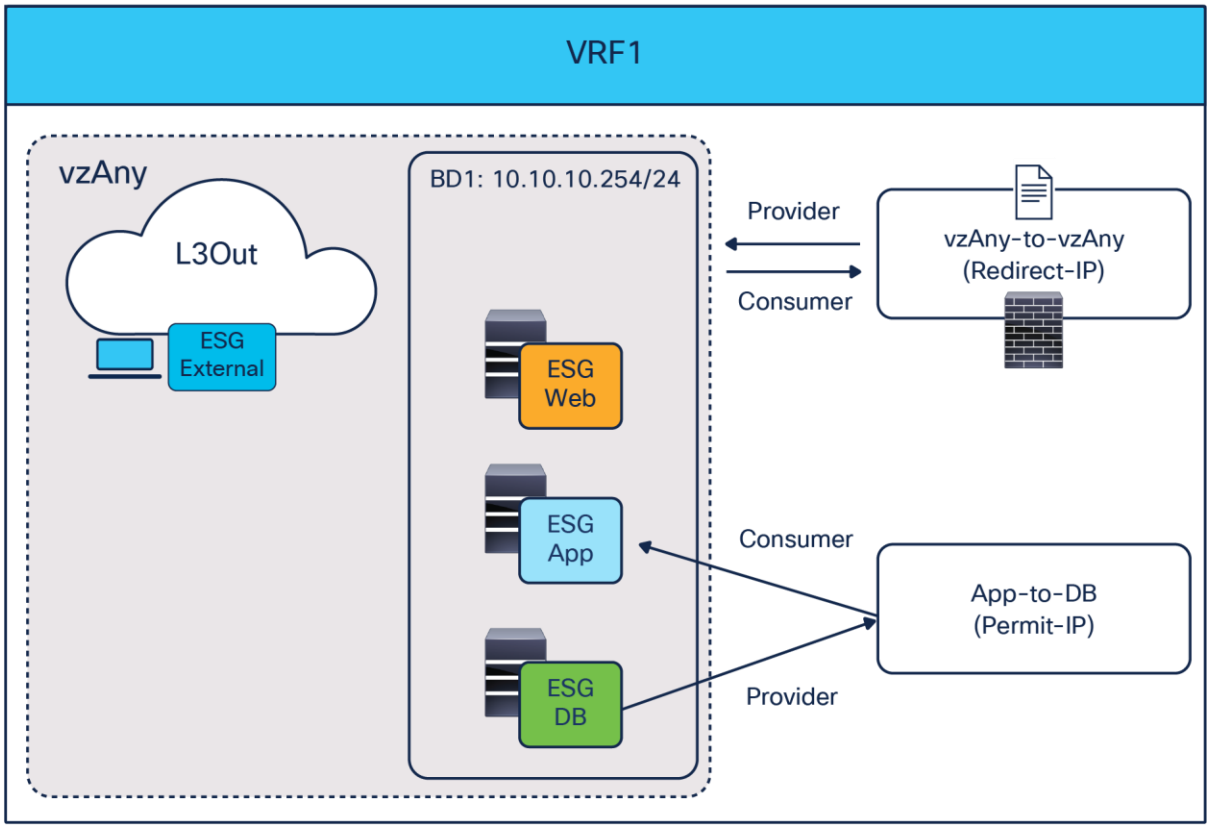


Figure 63.
Sample design 1: vzAny-to-vzAny PBR with specific ESG-to-ESG permit contract to bypass firewall

By using a vzAny-to-vzAny contract (with a “permit IP” filter) associated to a service graph with PBR, all inter-ESG traffic in the VRF is redirected to the firewall. Intra-ESG traffic is always permitted because the intra-ESG implicit permit rule (priority 3) wins over the vzAny-to-vzAny redirect rules (priority 17).

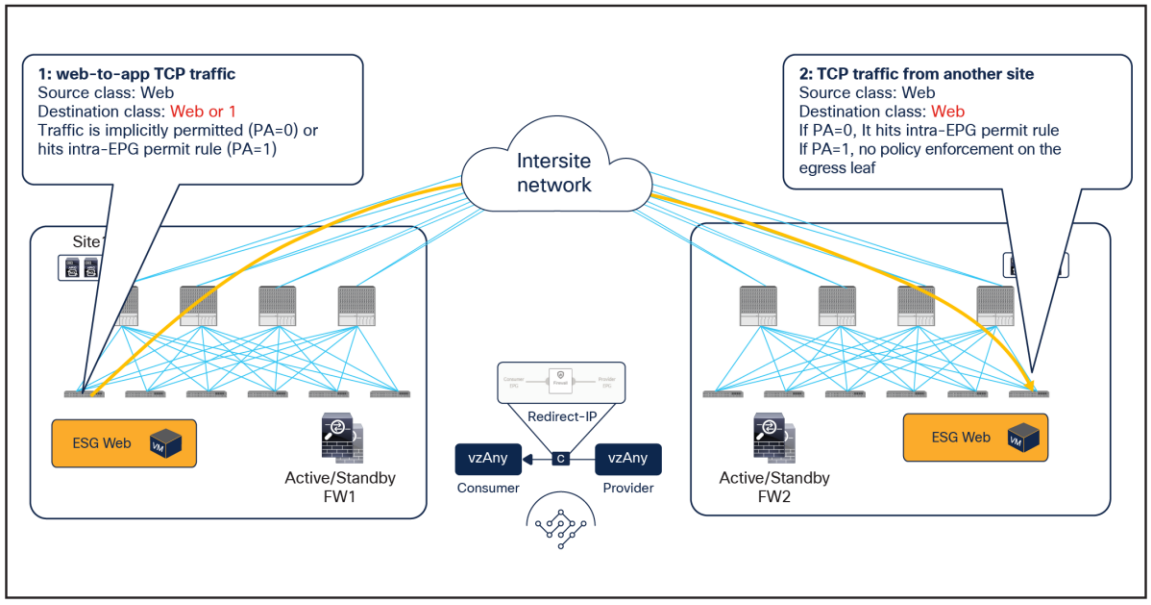


Figure 64.
Intra-ESG permit rule wins over vzAny-to-vzAny rule

© 2026 Cisco and/or its affiliates. All rights reserved.

By adding a specific ESG-to-ESG contract with a permit action, the redirection to the firewall can be bypassed for this communication because the specific ESG-to-ESG contract rules (priority 7 or 9) wins over the vzAny-to-vzAny contract redirect rules (priority 17).

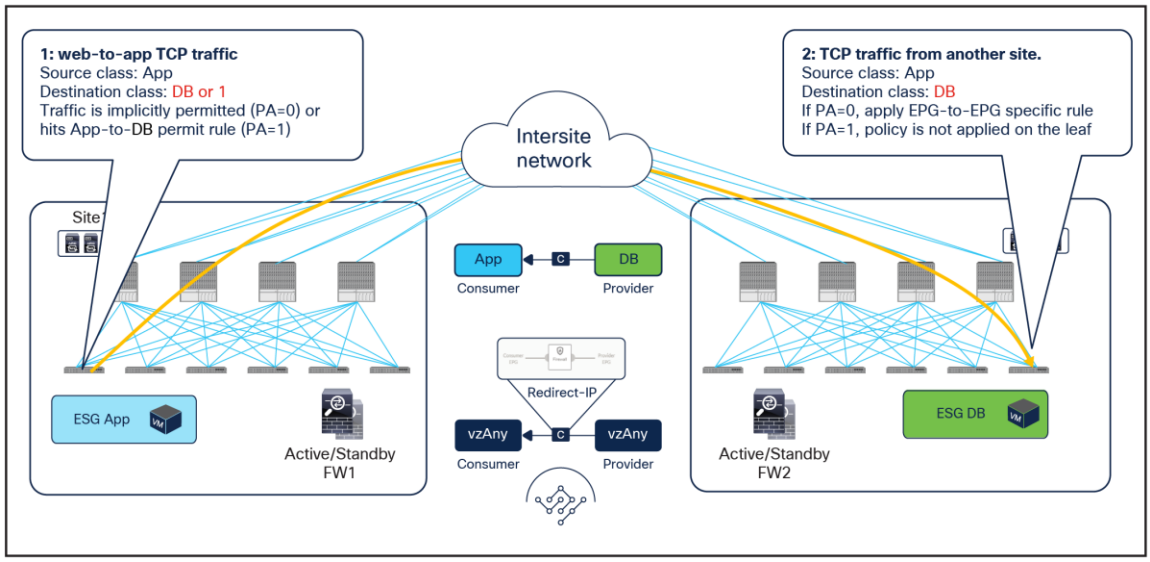


Figure 65.
Specific ESG-to-ESG rule wins over vzAny-to-vzAny rule

In addition to this, by using specific filters for the vzAny-to-vzAny contract with PBR, other inter-ESG traffic can be denied or just permitted in the VRF. For example, if vzAny-to-vzAny contract with PBR uses a more specific “permit TCP” filter instead of a “permit IP” filter, UDP traffic between ESG Web and ESG App will be denied because there is no permit or redirect rule applicable to that type of traffic (Figure 66).

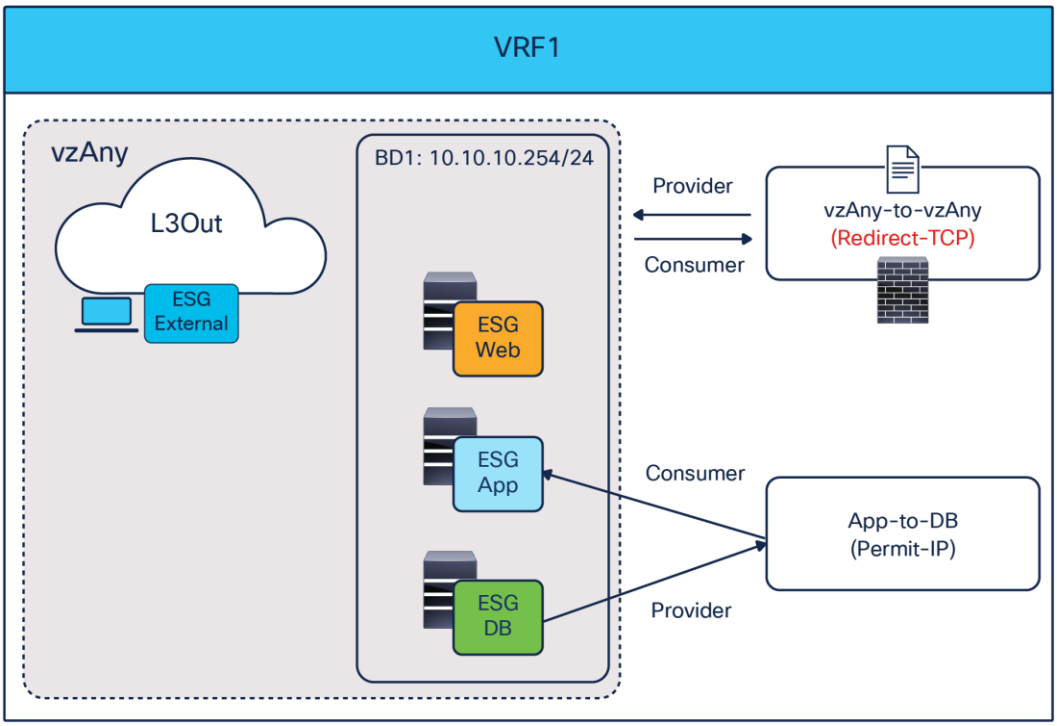


Figure 66.
Use specific filter to redirect specific vzAny-to-vzAny traffic

© 2026 Cisco and/or its affiliates. All rights reserved.

Example 2: Use different firewalls for north-south and east-west traffic

Example 1, above, uses the same one-arm firewall for both north-south and east-west traffic. If the requirement was, instead, to use a different two-arm firewall for north-south traffic (often this need is driven by security reasons), it is possible to introduce a separate vzAny-to-ESG contract with PBR that can coexist with the vzAny-to-vzAny contract with PBR used in Example 1 (Figure 67).

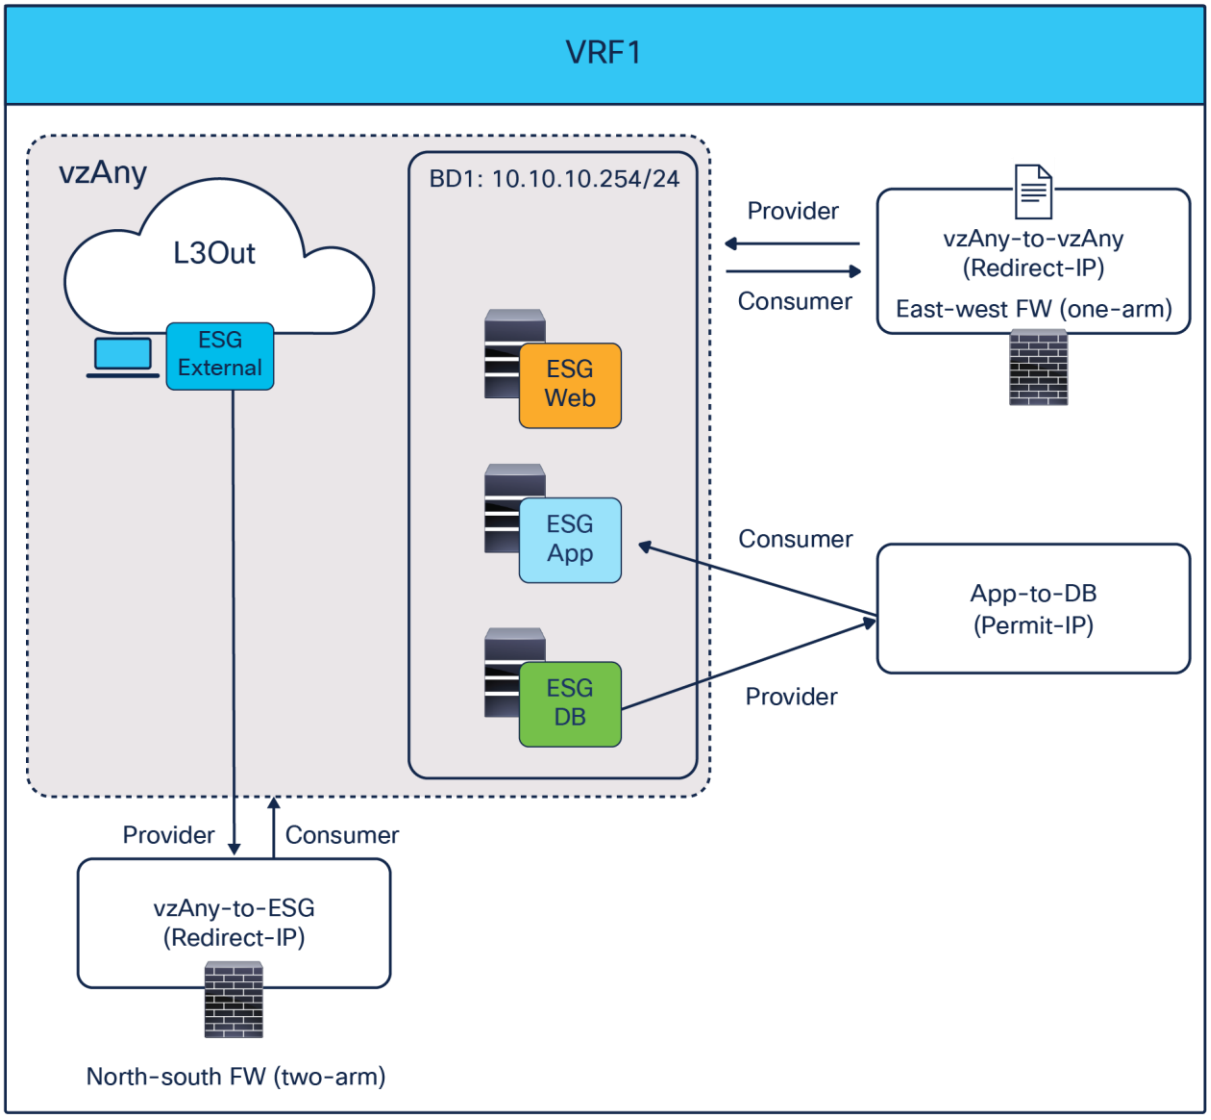


Figure 67.
Use a different two-arm firewall for north-south traffic

Since the vzAnu-to-ESG contract's rule (priority 10) wins over the vzAny-to-vAny rule (priority 17), north-south traffic will always be redirected to the two-arm firewall, whereas east-west traffic will still use the other one-arm firewall (with the exception of the communication between App and DB permitted by a specific contract).

© 2026 Cisco and/or its affiliates. All rights reserved.

Configuration examples

Overview

This section describes the general service chaining configuration in Cisco Nexus Dashboard (ND).

Note: This document shows GUI screenshots taken from Cisco ACI Release 6.2(1g) and ND release 4.1(1g). Thus, the GUI "look and feel" in this document might be slightly different from your specific ACI or NDO GUI.

Some objects must be created on each ACI domain and on the ND before going into the service-chaining configuration. This section does not cover how to configure interface policies, domains, tenants, VRFs, BDs, EPGs, L3Out, ESGs and contracts. The assumption is that these items are already configured.

For more information on the use of ND to deploy configurations and/or objects stretched across sites, please refer to the configuration guide listed at https://www.cisco.com/c/en/us/support/data-center-analytics/nexus-dashboard/series.html.

For the deployment of a service chaining specific to an ACI Multi-Site architecture, there are three configuration steps that must be performed on ND. The sections that follow provide more detailed information on each of the configuration steps listed below:

- Create a tenant policy template for defining an IP-SLA monitoring policy
- Create a service device template for defining the service device(s) to which redirect the traffic
- Configure a service chaining in a contract

Note: if a contract has a stretched consumer and provider objects, such as vzAny of a stretched VRF across sites or an ESG stretched across site, PBR configuration must be also done through ND.

Create a tenant policy template for IP-SLA monitoring policy

This step is to create an IP-SLA monitoring policy in a tenant policy template associated to a specific tenant and mapped to all the sites where the tenant is deployed. IP-SLA tracking is used to check availability information of each PBR destination and to detect each PBR destination MAC dynamically. It is recommended to enable IP-SLA tracking for faster failure detection. In this example, we are going to create an IP-SLA policy using Internet Control Message Protocol (ICMP) to monitor the status of the service device(s).

The first step is to create a tenant policy template. The location is at Manage > Orchestration > Tenant Templates > Tenant Policies > Actions > Create Tenant Policy Template.

© 2026 Cisco and/or its affiliates. All rights reserved.

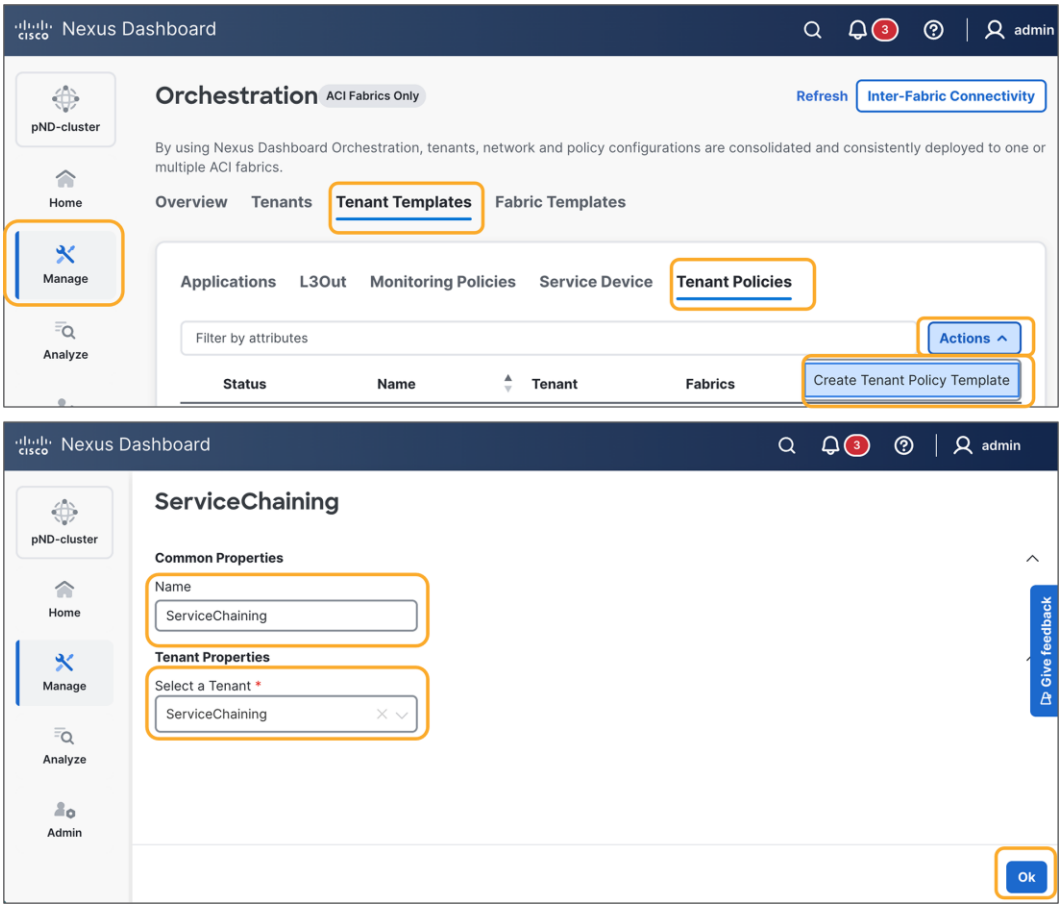


Figure 68.
Create tenant policy template

In the created tenant policy template, associate the template to the sites and save the template.

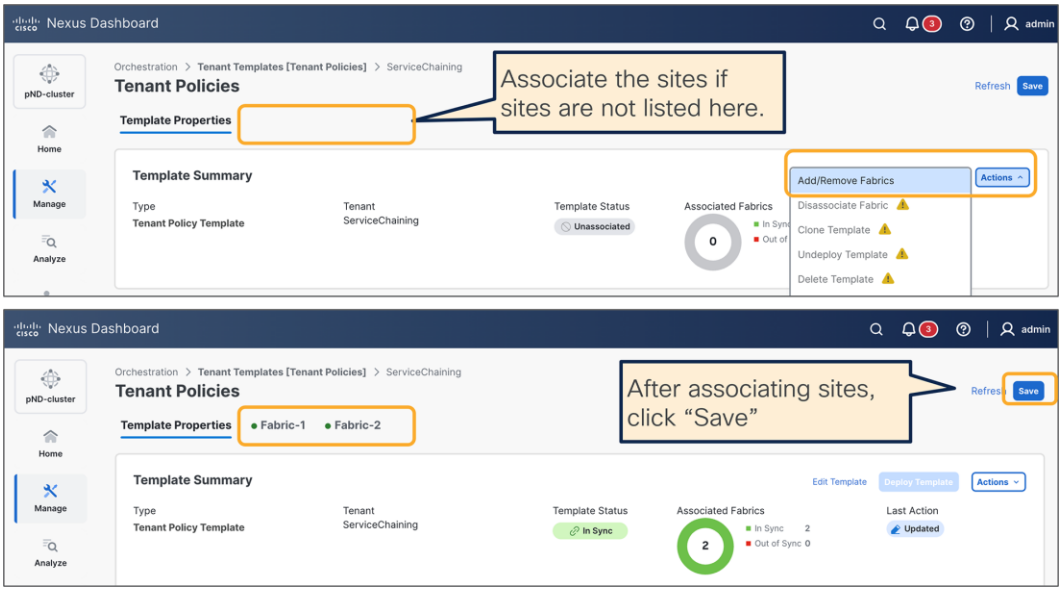


Figure 69.
Associate the sites

© 2026 Cisco and/or its affiliates. All rights reserved.

Then, Create Object > IPSLA Monitoring Policy, and specify SLA Type, SLA Frequency, Detect Multiplier, and other options if needed.

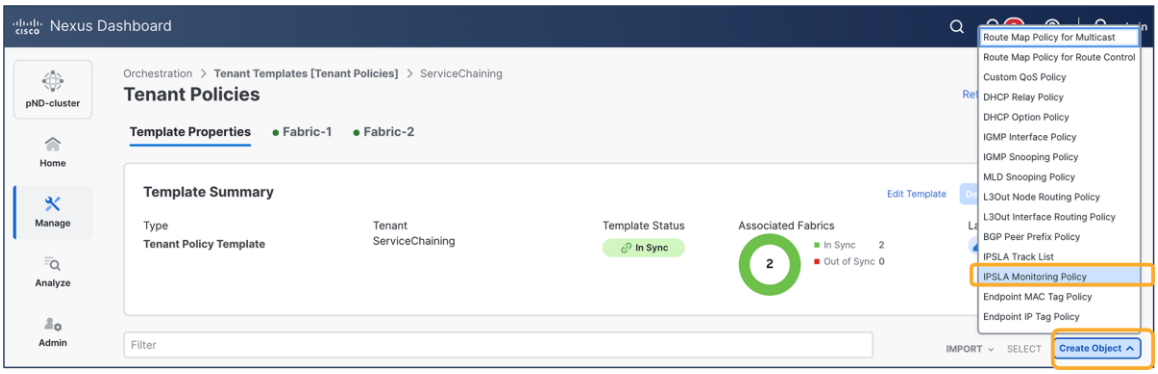


Figure 70.
Select “IPSLA Monitoring Policy”

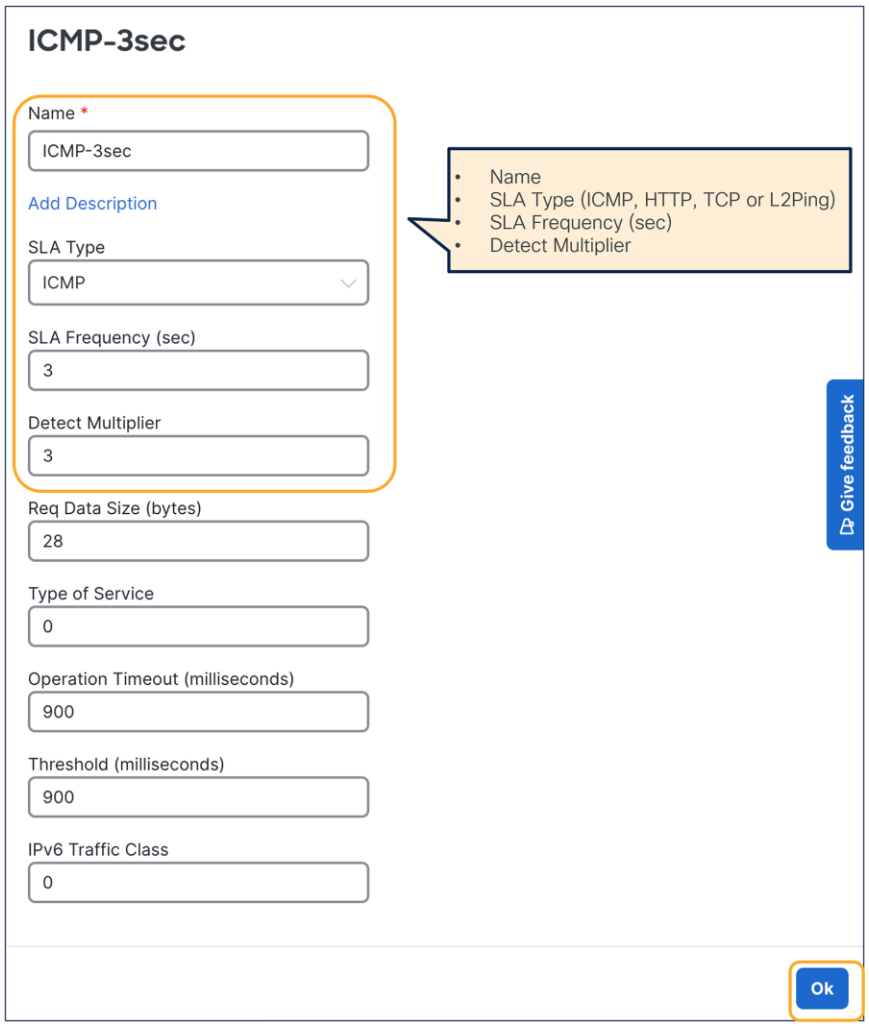


Figure 71.
Create IPSLA monitoring policy

© 2026 Cisco and/or its affiliates. All rights reserved.

Finally, deploy the tenant policy template to the sites.

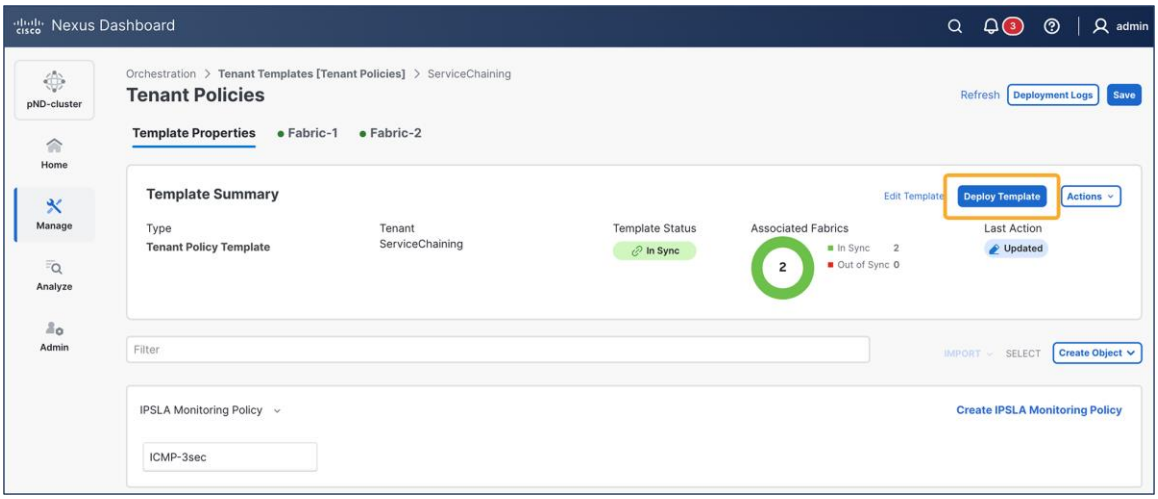


Figure 72.
Deploy the template to the sites

If the deployment is successfully done, the IP-SLA policy is created in the tenant on each APIC domain. The location on APIC is at Tenant > Policies > Protocol > IP SLA > IP SLA Monitoring Policies.

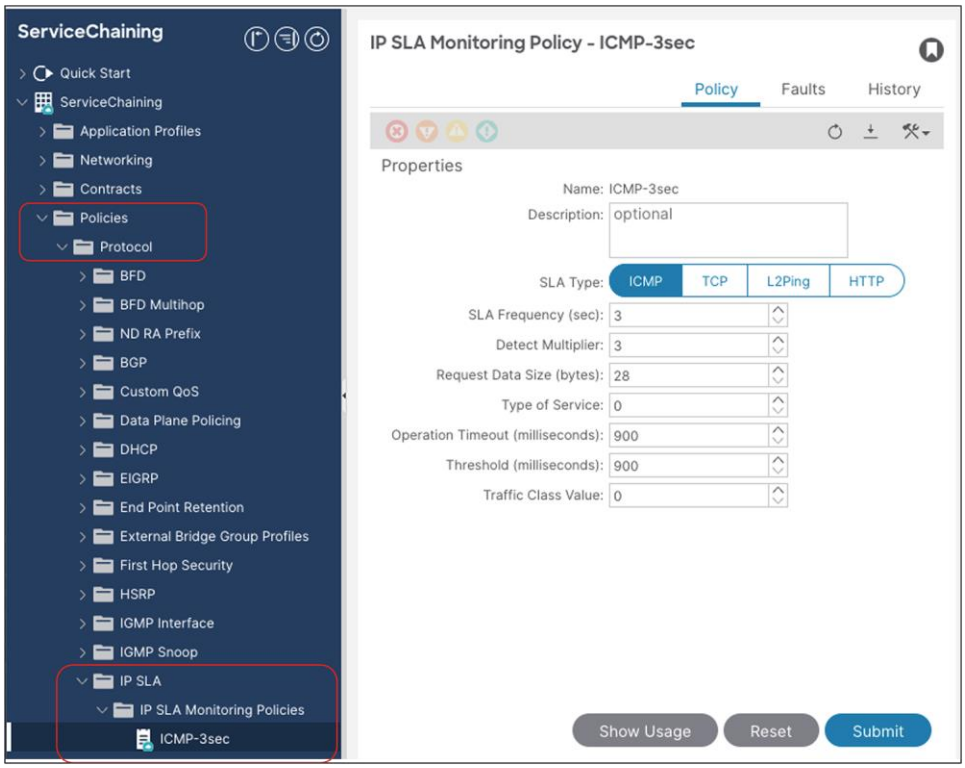


Figure 73.
Verify the configuration on APIC (IP-SLA monitoring policy)

© 2026 Cisco and/or its affiliates. All rights reserved.

Create a service device template

Service device templates allow you to define the service nodes (L4-L7 devices and PBR policies) associated to a given tenant in one or more sites. A service device template has a template-level configuration and a site-level configuration as follows:

- Template-level configuration:
  - Device type: FW/LB/Others
  - Device mode: L3/L2/L1 (routed/transparent/inline)
  - Number of interfaces of the service device: one-arm, two-arm, or more
    - BD or L3Out for each service-device interface
    - Redirection (PBR) enabled or disabled for each service-device interface
- Site-level configuration:
  - Domain type: physical or VMM domain
  - Path (physical domain) or VM (VMM domain) information

Note: PBR to a destination in another site because of local PBR destination failure is NOT supported. Thus, it is strongly recommended to deploy the local service nodes in a highly available way.

The first step is to create a service device template. The location is Manage > Orchestration > Tenant Templates > Service Device > Create Service Device Template.

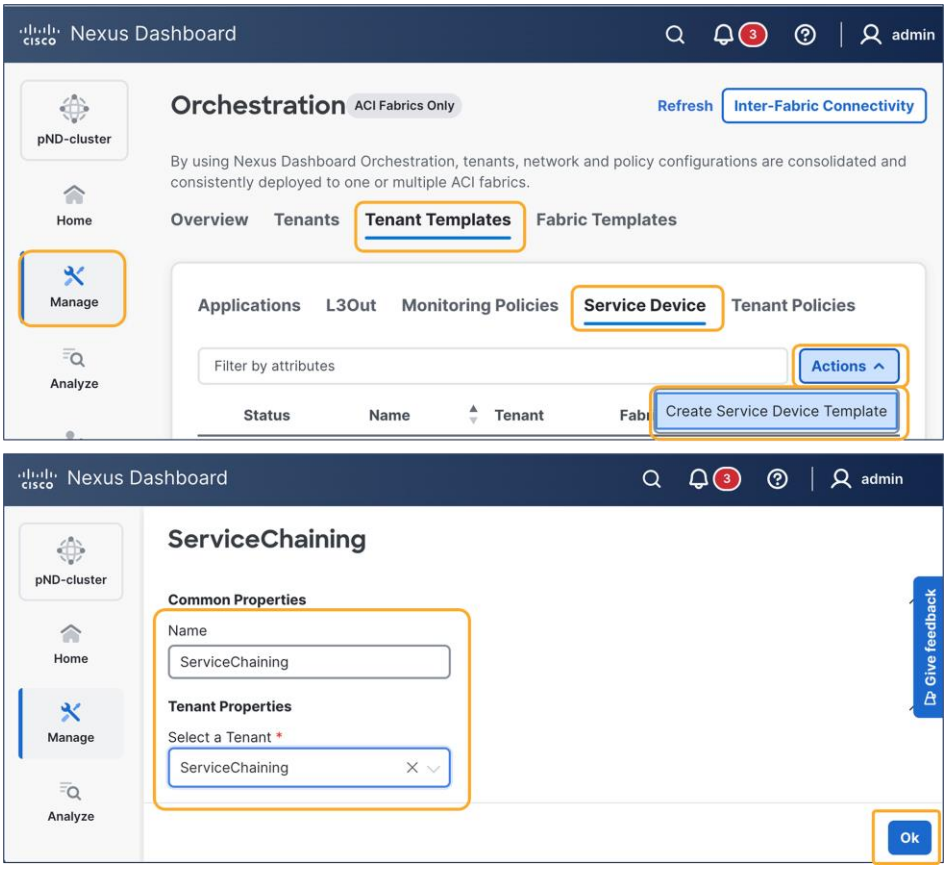


Figure 74.
Create a service device template

© 2026 Cisco and/or its affiliates. All rights reserved.

In the created service device template, select “Add/Remove Sites” to associate the template to the sites, and save the template.

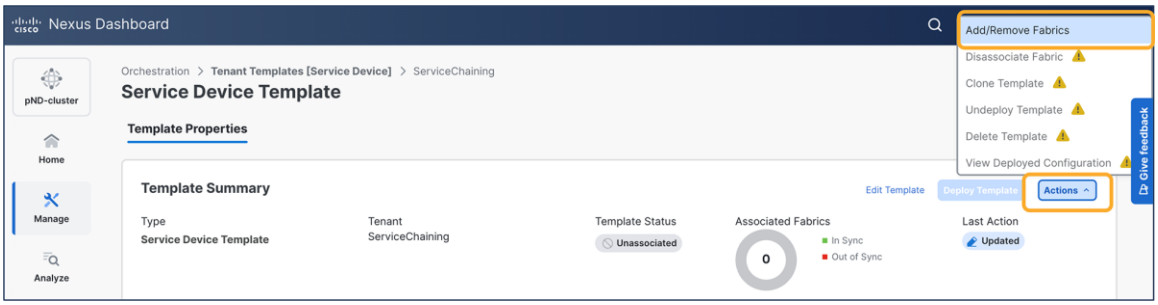


Figure 75.
Associate the template to the sites

Next, you are going to create a service device in the template. The location is at Create Object > Service Device Cluster. After you complete template-level configurations for the service device, you are going to configure site-level configurations.

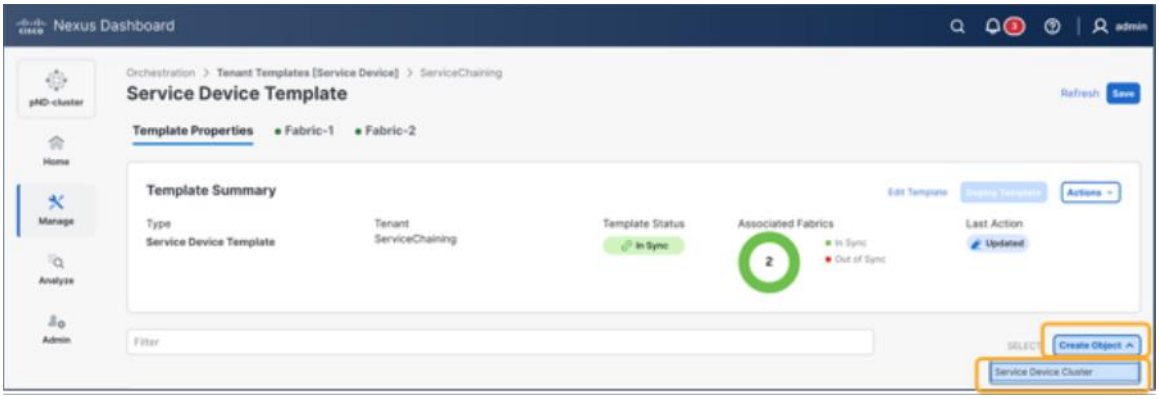


Figure 76.
Create a Service Device Cluster

In this document, we are going to configure the following service devices illustrated in the figure below:

- L3 firewall with two interfaces
- L3 load balancer with two interfaces

Depending on the service-device insertion use case, you are going to redirect traffic to only one or to both interfaces. Although some use cases only support one-arm mode insertion, the service device can have more than one interface.

Note that the BDs where the service-devices’ interfaces are connected need to be selected as part of the service device template configuration. Thus, please ensure you have already created those BDs in the tenant by using specific application templates.

© 2026 Cisco and/or its affiliates. All rights reserved.

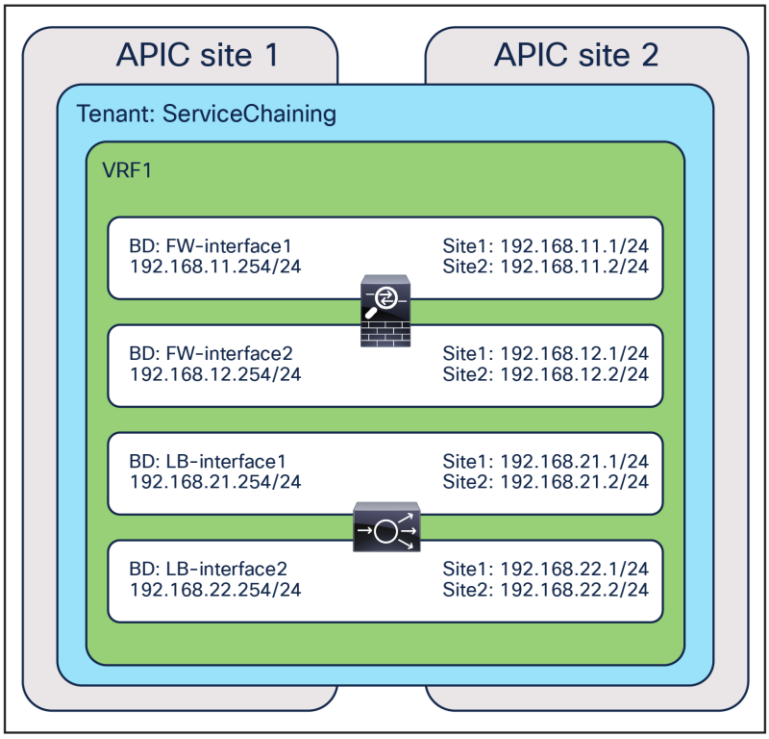


Figure 77.
Example of a service-device configuration

L3 firewall configuration example

At the template level for a service device template configuration, you first need to select a device type and device mode. Some options are grayed out if the options are not applicable based on your selection.

The next step is to select the connectivity mode. If one-arm is selected, you are going to select a bridge domain or an L3Out for one interface. If redirect is enabled, the IP-SLA policy configuration option is shown in the UI. Though not mandatory, it is generally recommended to specify an IP-SLA policy that enables IP-SLA tracking for the PBR destinations for the interface.

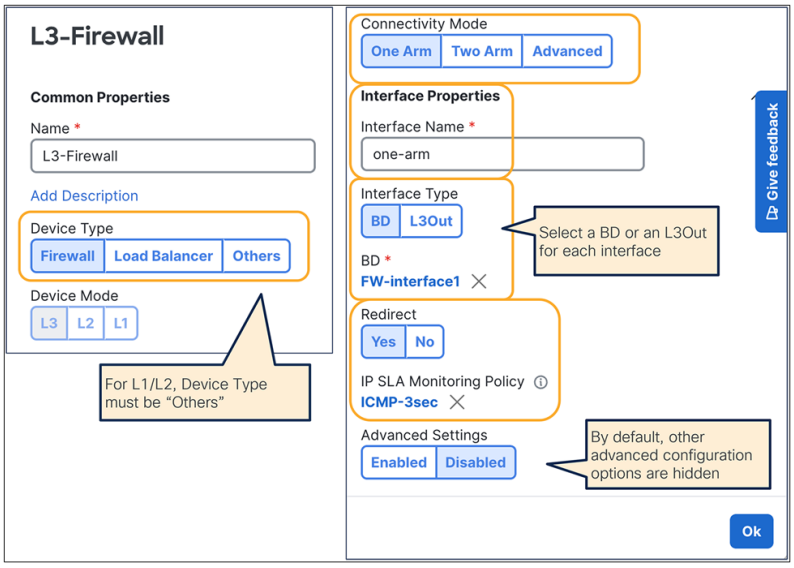


Figure 78.
Example of a Service Device Cluster configuration (one-arm)

© 2026 Cisco and/or its affiliates. All rights reserved.

If you select two-arm, the UI shows the table that has the list of interfaces. You need to repeat the configuration for each interface by clicking the pencil icon for each interface. If you have more than two interfaces, select Advanced and continue configuring the third interface, and so on.

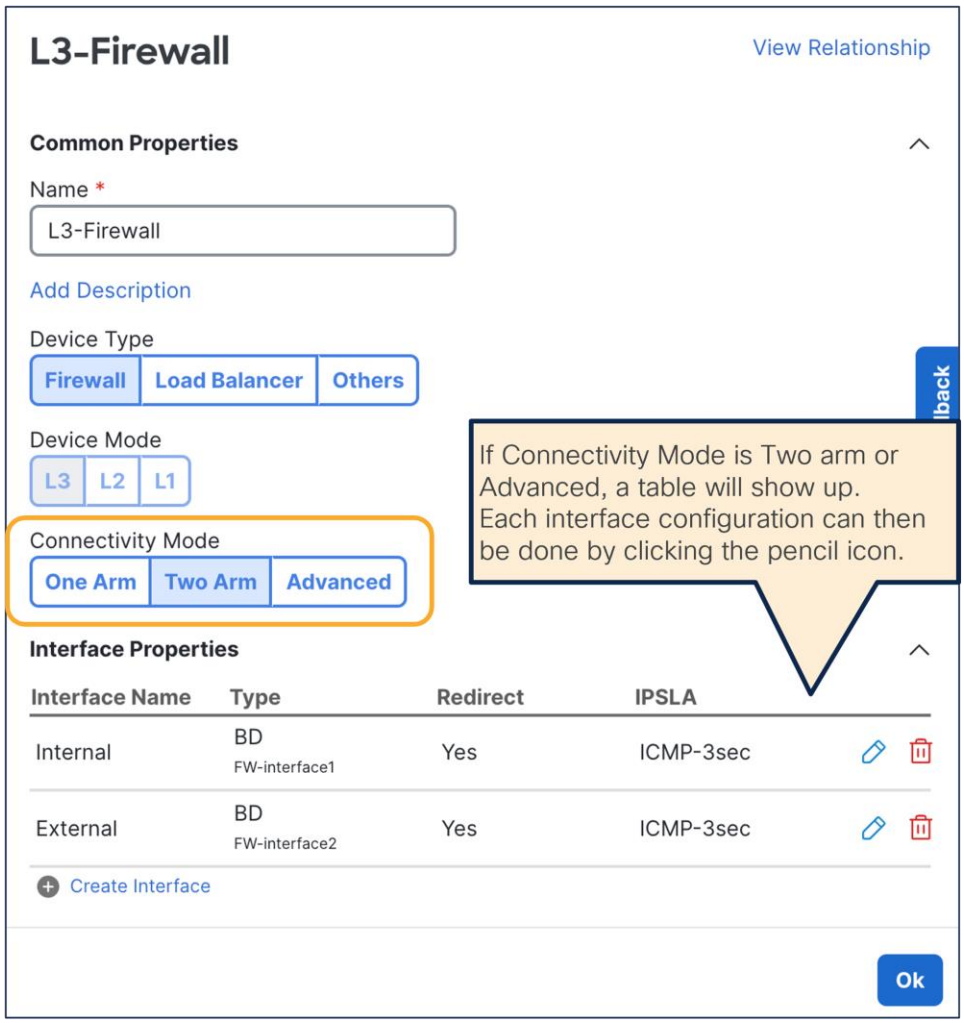


Figure 79.
Example of a Service Device Cluster configuration (two-arm)

© 2026 Cisco and/or its affiliates. All rights reserved.

After the template-level configuration is completed, the next step is to configure site-level configurations. Select one of the sites and click the service device that you just created.

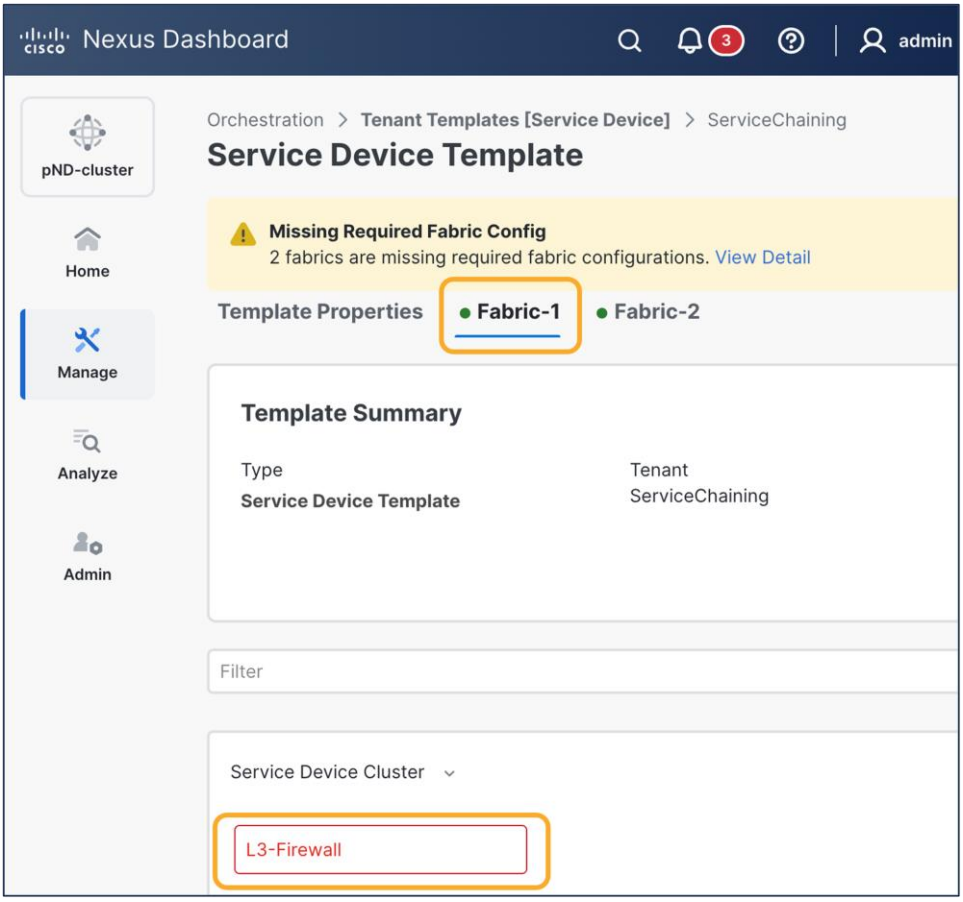


Figure 80.
Site-level configuration

Select a physical domain or a VMM domain. Depending on the domain, the required configuration options will be different.

In the case of VMM domain, by default a VLAN will be allocated dynamically from the VLAN pool used in the VMM domain. Thus, a VLAN ID is not a mandatory configuration, but you can specify a VLAN if the VLAN pool has a static VLAN range. If Link Aggregation Control Protocol (LACP) is used for virtual switch to upstream switch connectivity, select Enhanced LAG Option. Then select a VM and its virtual Network Interface Card (NIC).

If you have more than one interface, repeat the step for other interfaces. The figure below shows a two-arm-mode firewall example.

© 2026 Cisco and/or its affiliates. All rights reserved.

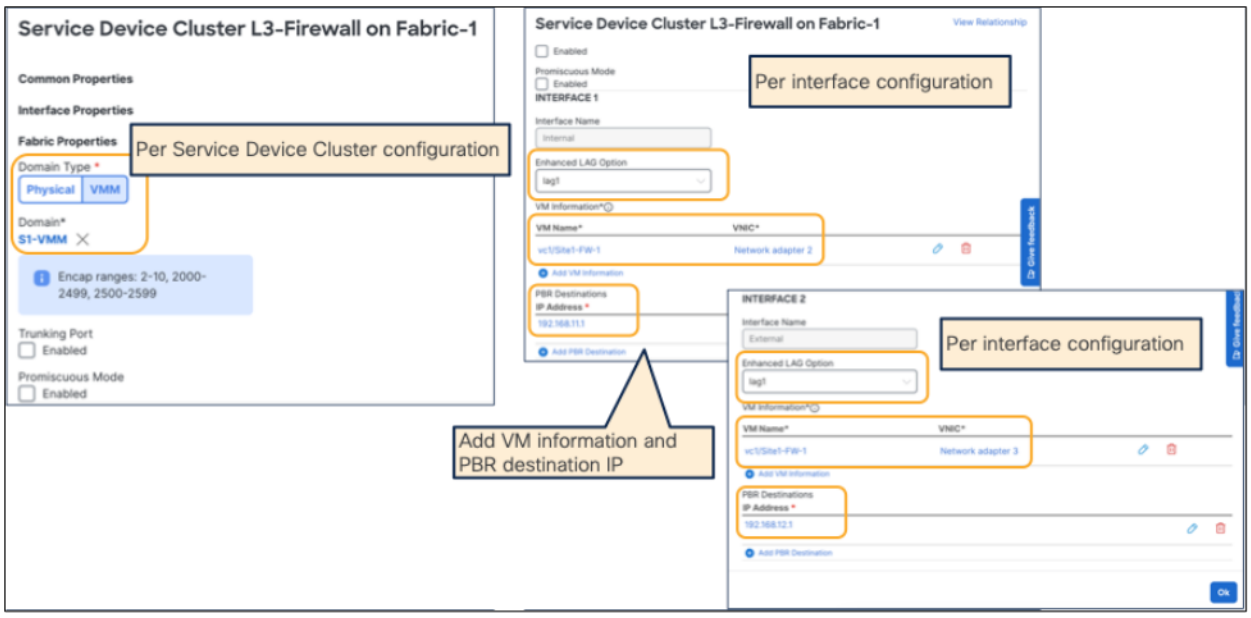


Figure 81.
Example of a site-level configuration (VMM domain)

If you have more than one VM, repeat this step for other VMs.

- Figure 82 shows an example of an active-standby HA pair: two VMs with one PBR destination IP per PBR policy (two-arm deployment). A single IP identifies the HA pair on each of the defined interfaces, and, as a consequence, a health group is automatically created.
- Figure 83 shows, instead, the deployment of independent service nodes: two VMs with two PBR destination IPs per PBR policy (in this example, also a two-arm deployment). In this case, a specific TAG is required for health-group configuration and must be associated to each PBR destination to group the PBR destination IP for consumer-to-provider direction and the PBR destination IP for provider-to-consumer direction from multiple PBR destination IP addresses.

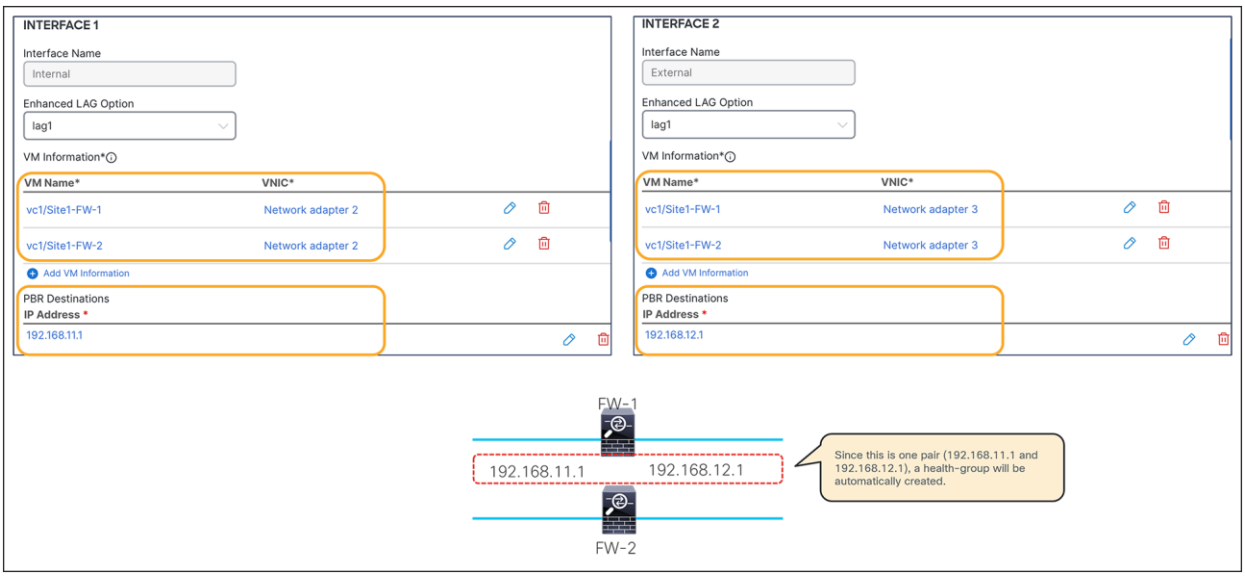


Figure 82.
Example of a site-level configuration for an active/standby HA pair that is part of a VMM domain

© 2026 Cisco and/or its affiliates. All rights reserved.

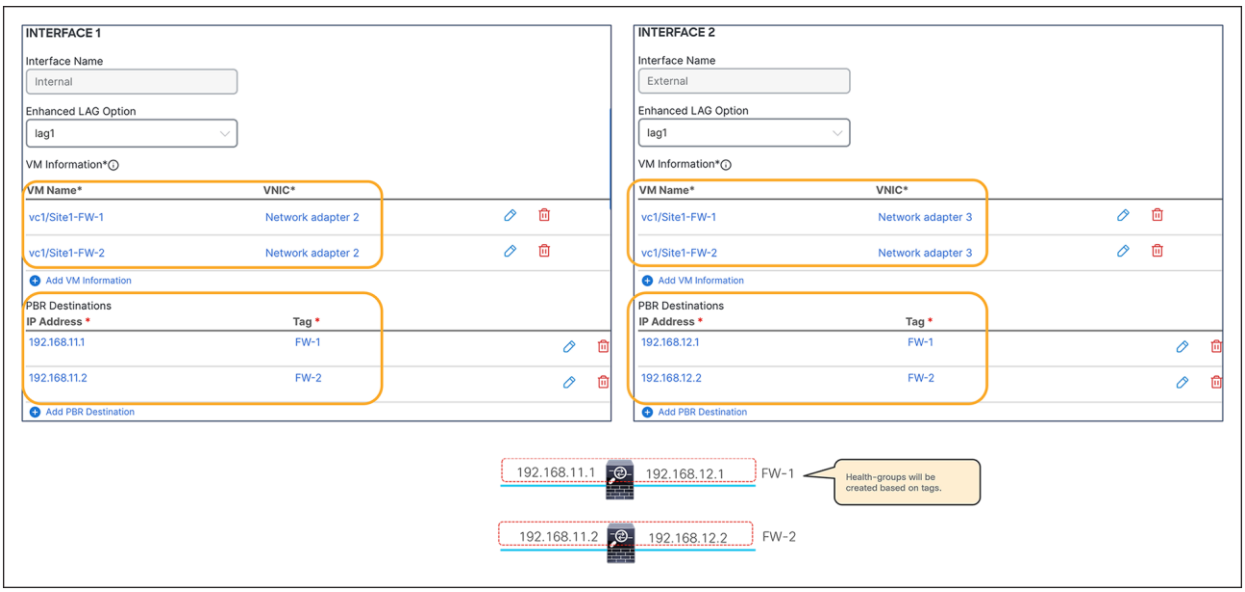


Figure 83.
Example of a site-level configuration for independent service nodes that are part of a VMM domain

In the case of a physical domain, select a path and a VLAN ID for each interface similar to an EPG with static path bindings. Figure 84 is an example of a two-arm active/standby HA example: two VMs with one PBR destination IP per PBR policy.

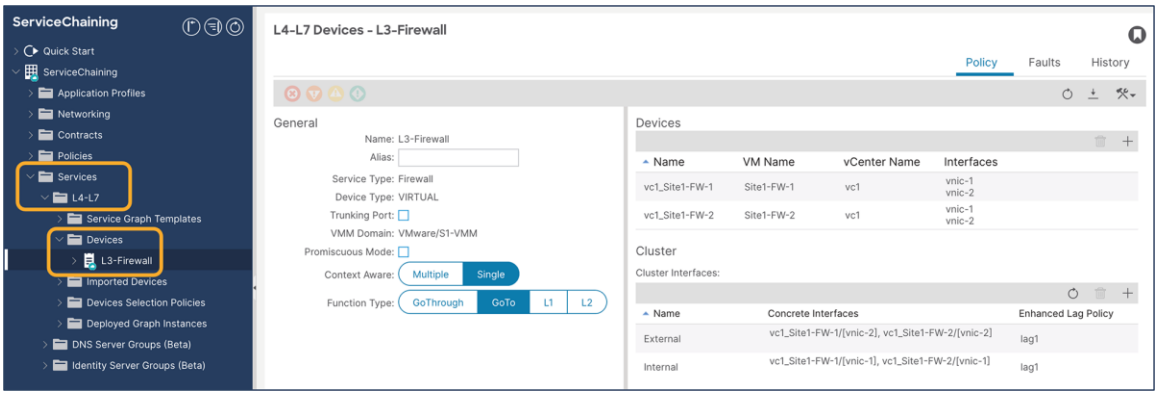


Figure 84.
Example of a site-level configuration (physical domain)

© 2026 Cisco and/or its affiliates. All rights reserved.

After the site-level configuration is done for a site, repeat it for the remaining site(s) and then deploy the template.

If deployment is successfully done, the L4-L7 Device and the PBR policy are created in the tenant in each APIC domain. The locations on APIC are at Tenant > Services > L4-L7 > Devices for the L4-L7 Devices and at Tenant > Policies > Protocol > L4-L7 Policy-Based Redirect for the PBR policy.



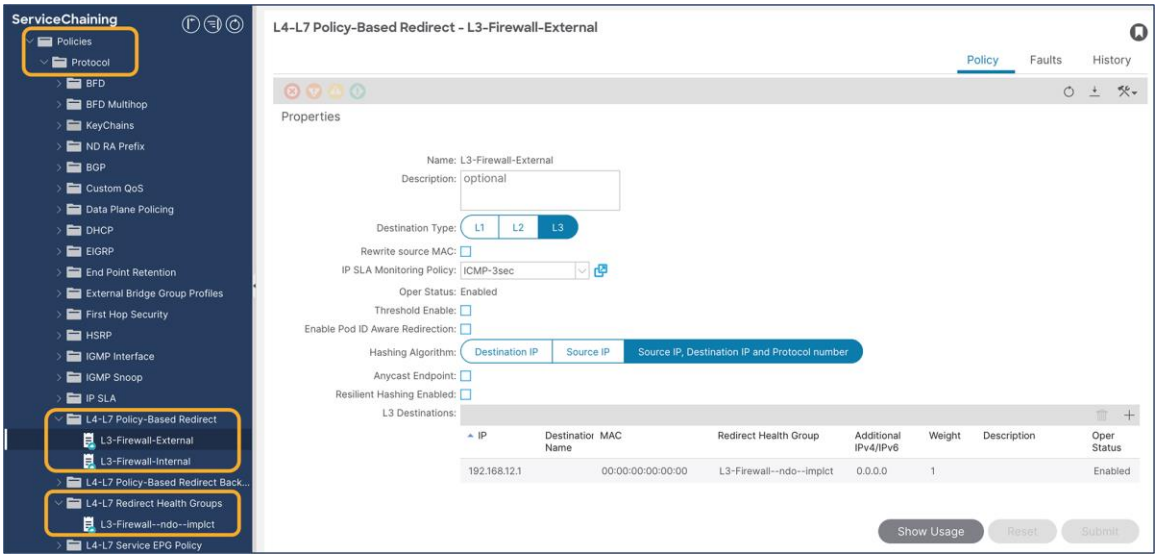


Figure 85.
Verify the configurations on APIC (L4-L7 Device and PBR policy

© 2026 Cisco and/or its affiliates. All rights reserved.

Advanced options

If “Advanced Option” is enabled, the UI shows additional configuration items, such as load-balancing hashing options, etc. This approach is taken to show the minimum configuration options (usually the most commonly used), unless advanced configurations are required.

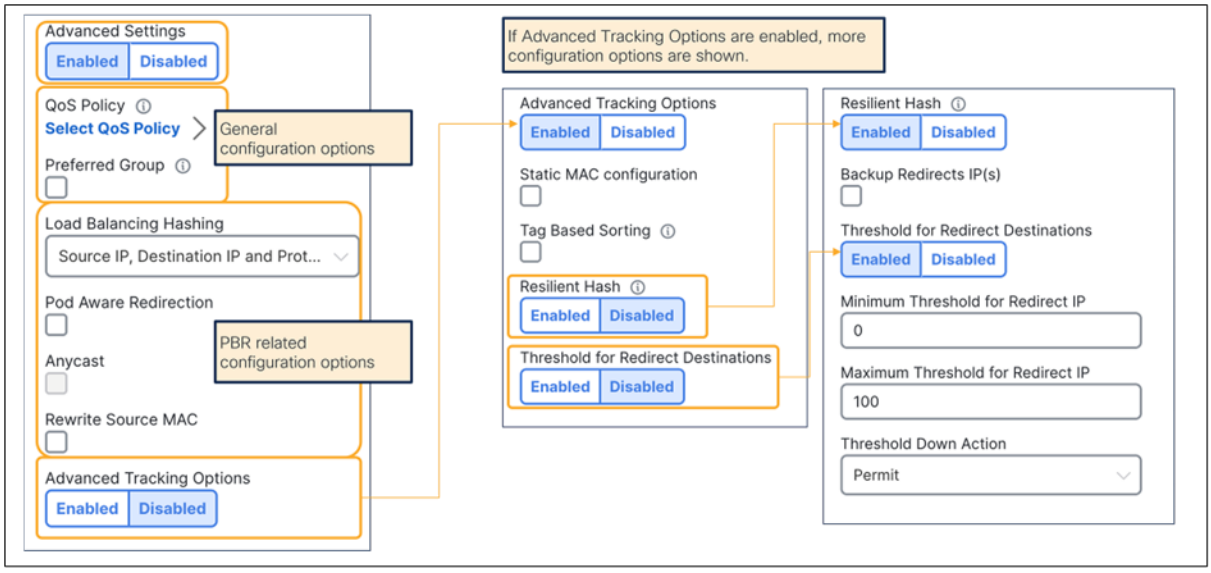


Figure 86.
Advanced configuration options (template-level)

Depending on the “Advanced” settings in the template-level configuration, you might require additional configurations in the site-level configuration. For example, if “Pod Aware Redirection” is enabled, you need to configure a pod ID for each PBR destination IP.

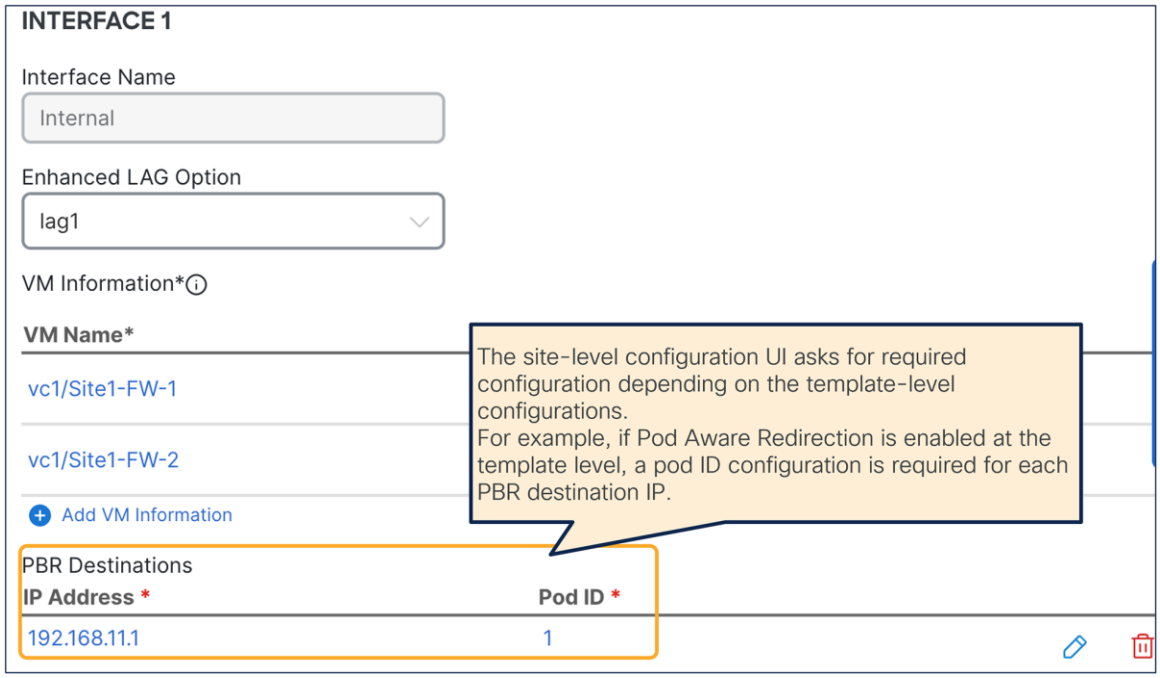


Figure 87.
Advanced configuration options (site-level)

© 2026 Cisco and/or its affiliates. All rights reserved.

L3 load-balancer configuration example

Because the configuration steps and options required for the load balancer are almost identical to those for the firewall example presented in the previous subsection, this subsection shows only the configurations required for a two-arm load balancer deployed as a virtual machine that is part of a VMM domain.

At the template-level of the service device template configuration, you first need to select a device type and device mode. Some options are grayed out if the options are not applicable based on your selection. The figure below shows an example of a two-arm load balancer where PBR is enabled on the internal interface only. If redirect is not enabled, there is no need to select an IP-SLA policy.

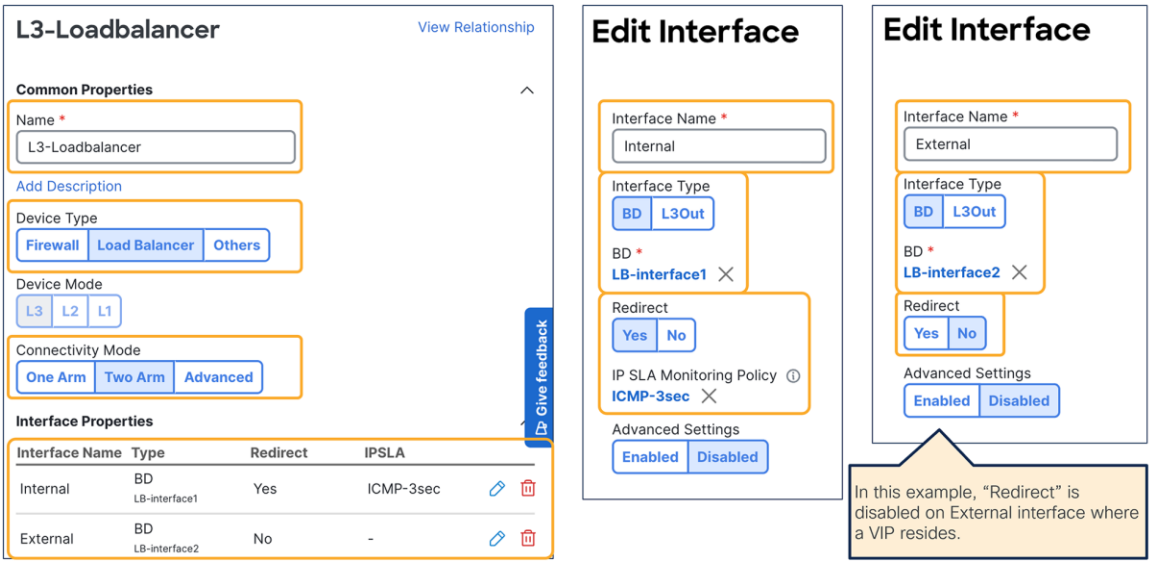


Figure 88.
Example of a load-balancer configuration (two-arm, template-level)

At the site level configuration, select a physical domain or a VMM domain. Depending on the selected domain, the required configuration options will be different. If PBR is not enabled on the interface at the template level, there is no need to configure the PBR destination's IP information for that interface.

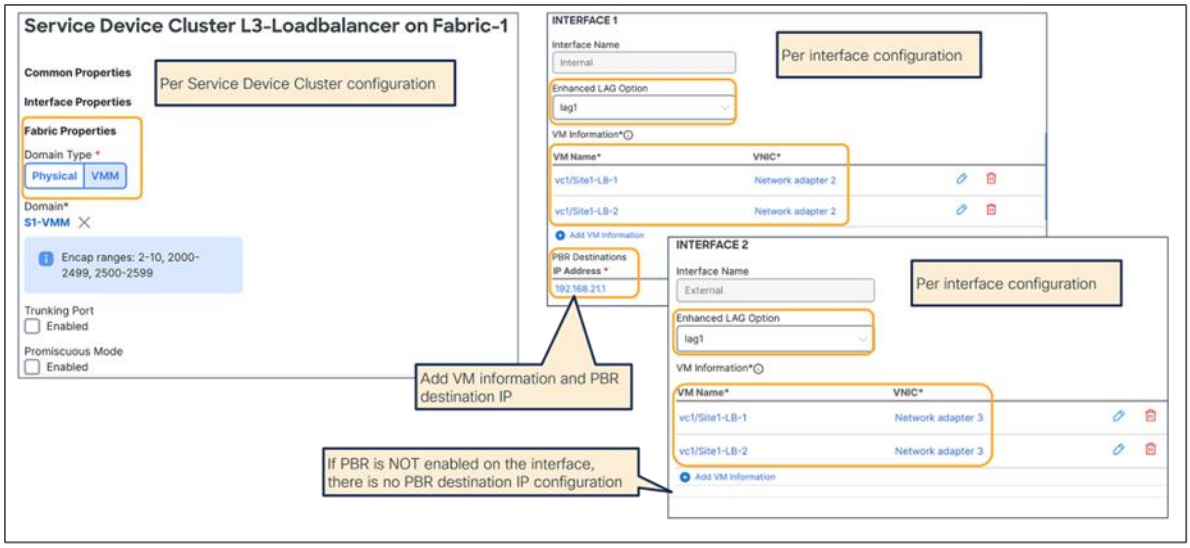


Figure 89.
Example of a load-balancer configuration (two-arm, site-level with VMM domain)

© 2026 Cisco and/or its affiliates. All rights reserved.

After the site-level configuration is done for a site, repeat it for the remaining site(s) and then deploy the template.

If the deployment is successfully done, the L4-L7 Device and the PBR policy are created in each APIC domain for the specified tenant. The locations on APIC where to find the provisioned objects are at Tenant > Services > L4-L7 > Devices for the L4-L7 Device and at Tenant > Policies > Protocol > L4-L7 Policy-Based Redirect for the PBR policy.

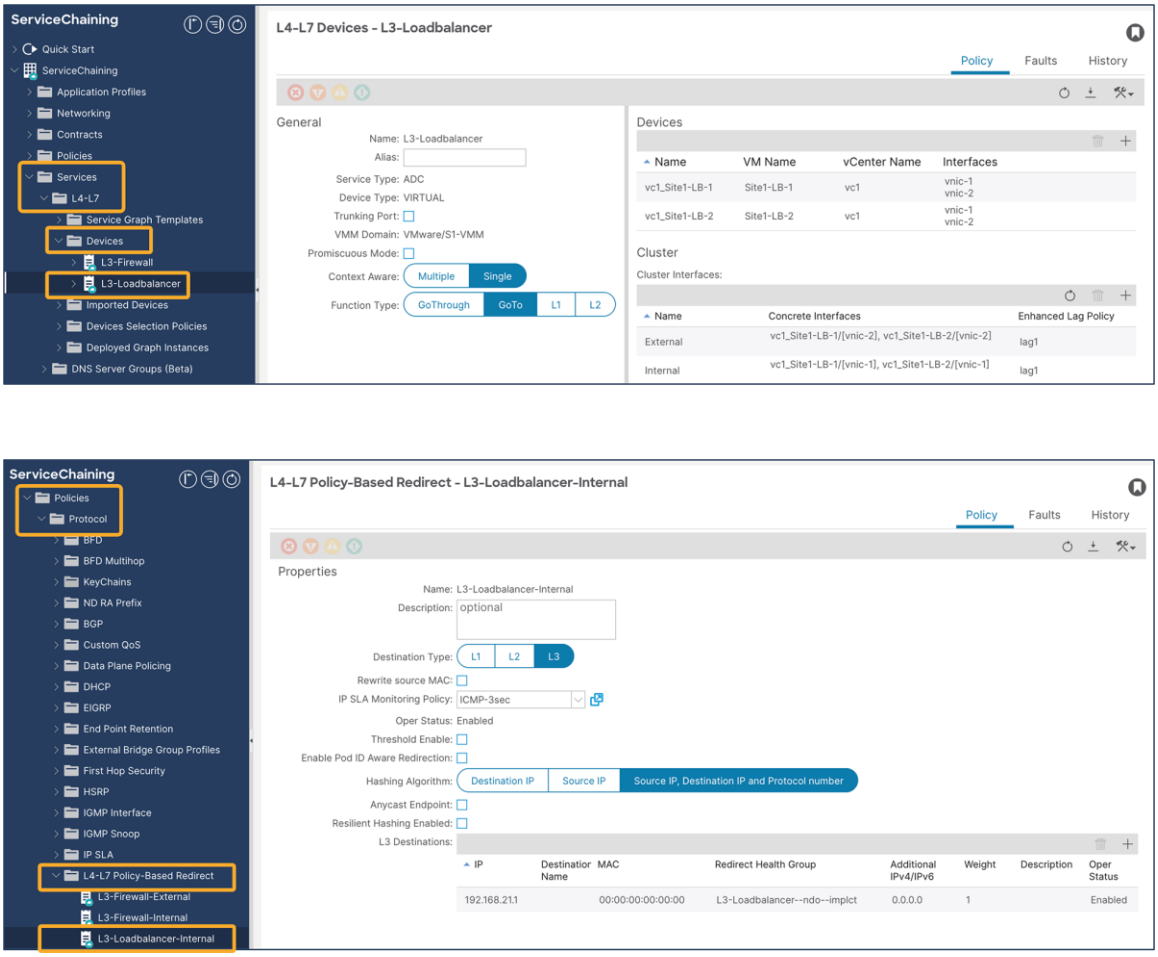


Figure 90.
Verify the configuration on APIC (L4-L7 Device and PBR policy)

© 2026 Cisco and/or its affiliates. All rights reserved.

Application template

This step slightly differs depending on use cases. This section will cover the following use cases:

- ESG-to-ESG contract with PBR
- vzAny-to-vzAny contract with PBR
- vzAny-to-ESG contract with PBR (firewall and load balancer)

Prerequisites

"L3 Multicast" and "Fabric-aware Policy Enforcement Mode" knobs must be enabled (Figure 91). For L3 Multicast, the provisioning of a Rendezvous Point (RP) is not required. The reason "L3 Multicast" is required is explained at Why do "Fabric-aware Policy Enforcement Mode" and "L3 Multicast" need to be enabled on the VRF for vzAny PBR? in FAQ section.

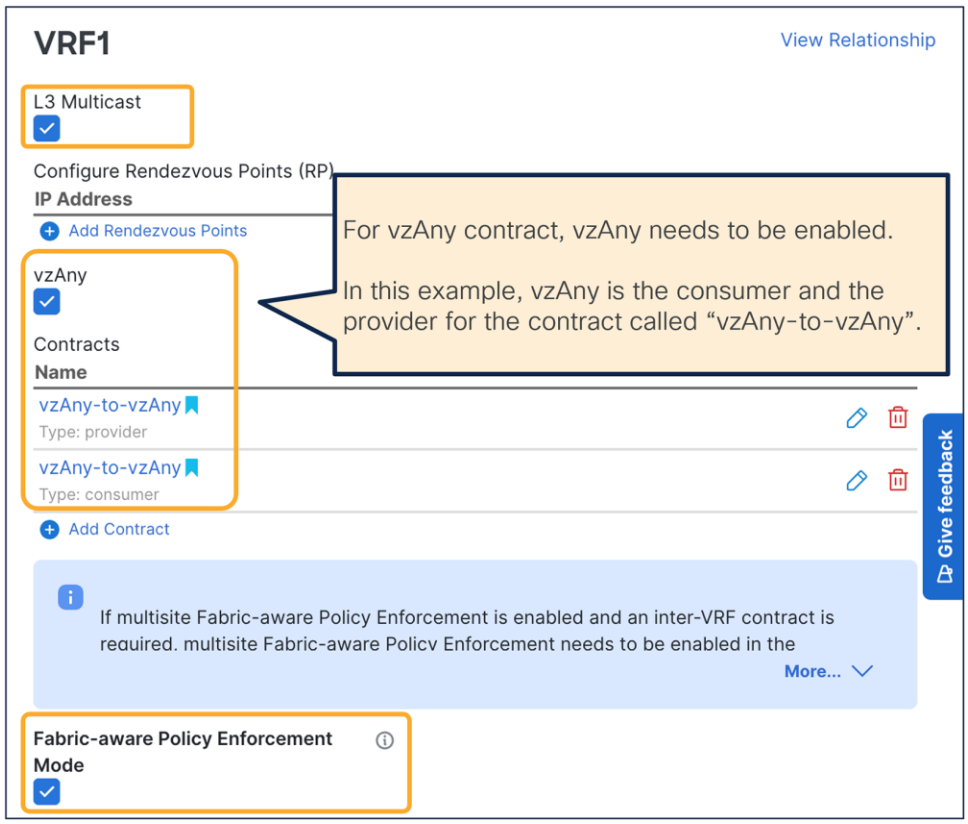


Figure 91.
"L3 Multicast" and "Fabric-aware Policy Enforcement Mode" on VRF

© 2026 Cisco and/or its affiliates. All rights reserved.

The consumer, provider, and service BDs must be set to “Hardware Proxy” mode (Figure 92).

This is because when a BD is in “Flood” mode if the destination MAC is unknown, packets will hit an implicit permit rule (any-to-BD_class_ID) and won’t be redirected. Please see Cisco ACI contract guide that has implicit rule list.

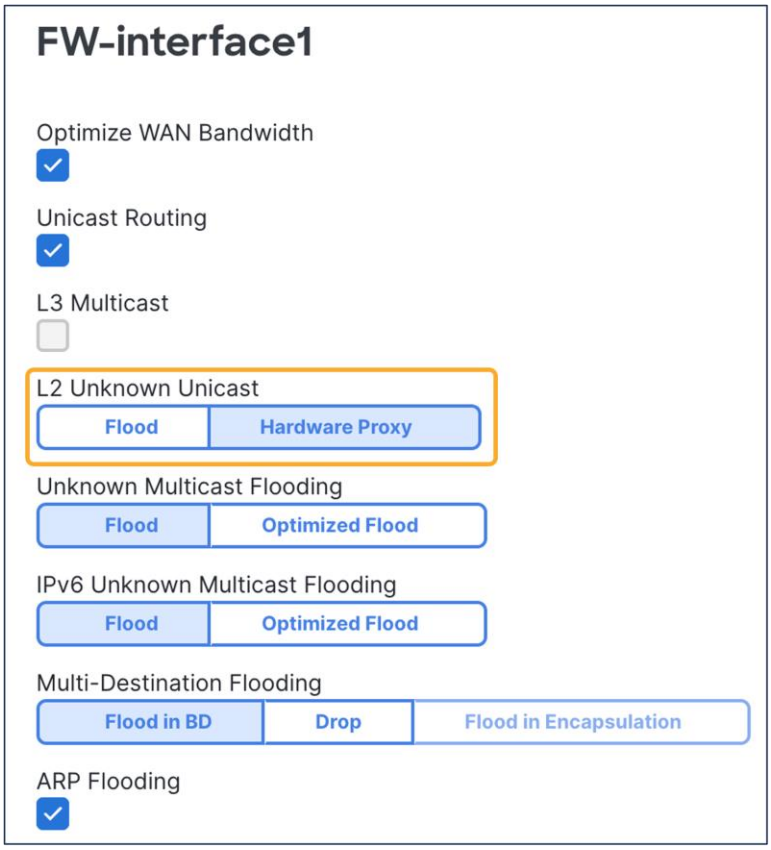


Figure 92.
BDs must be set to “Hardware Proxy Mode.”

Example 1: ESG-to-ESG and vzAny-to-ESG contracts with PBR

This subsection covers a configuration step for ESG-to-ESG and vzAny-to-ESG contracts with PBR for firewall and load-balancer insertion. One node and multinode service chaining, such as inserting a firewall and then a load balancer, are supported for ESG-to-ESG and vzAny-to-ESG contracts with PBR. PBR can be enabled on both directions or either one of the directions. Though the figure below illustrates multiple intra-VRF examples, inter-VRF is also supported. For inter-VRF, the service BDs must be in either the consumer or the provider VRF.

© 2026 Cisco and/or its affiliates. All rights reserved.

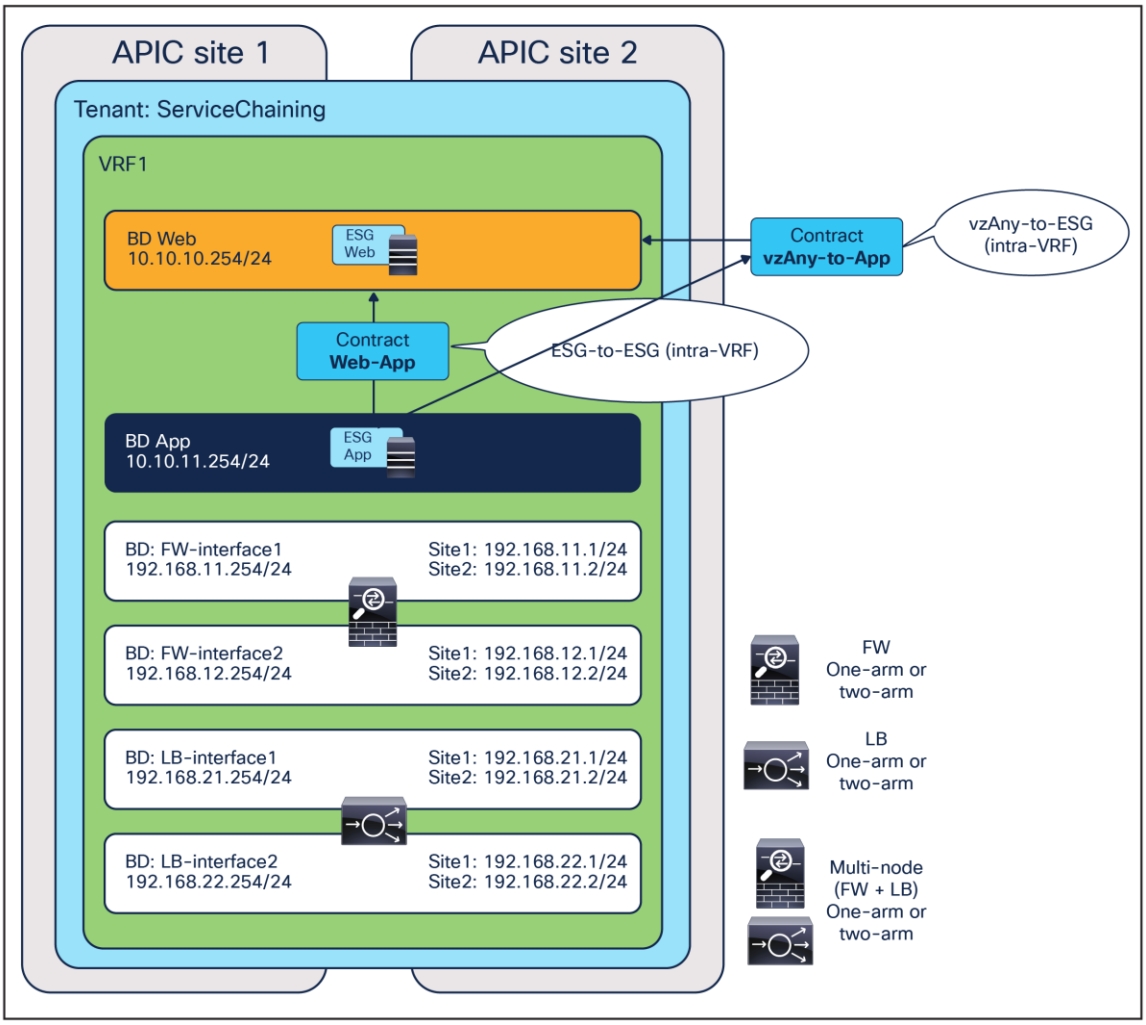


Figure 93.
ESG-to-ESG and vzAny-to-ESG contracts with PBR (intra-VRF)

This subsection shows the following configuration examples:

- Bidirectional PBR for firewall insertion
- Uni-directional PBR for load-balancer insertion (PBR is enabled for the provider-to-consumer direction.)

Configure service chaining

Service insertion is configured by associating one or more service devices with a contract. This is provisioned from Configure > Tenant Template > Applications. Select your template and then the contract.

At the bottom of the contract's configuration, ensure service chaining is selected. It shows the list of the consumers and the providers of the contract (if already configured). By clicking the "+" icon, you can add one or more service devices to create a service chaining between the consumers and the providers.

© 2026 Cisco and/or its affiliates. All rights reserved.

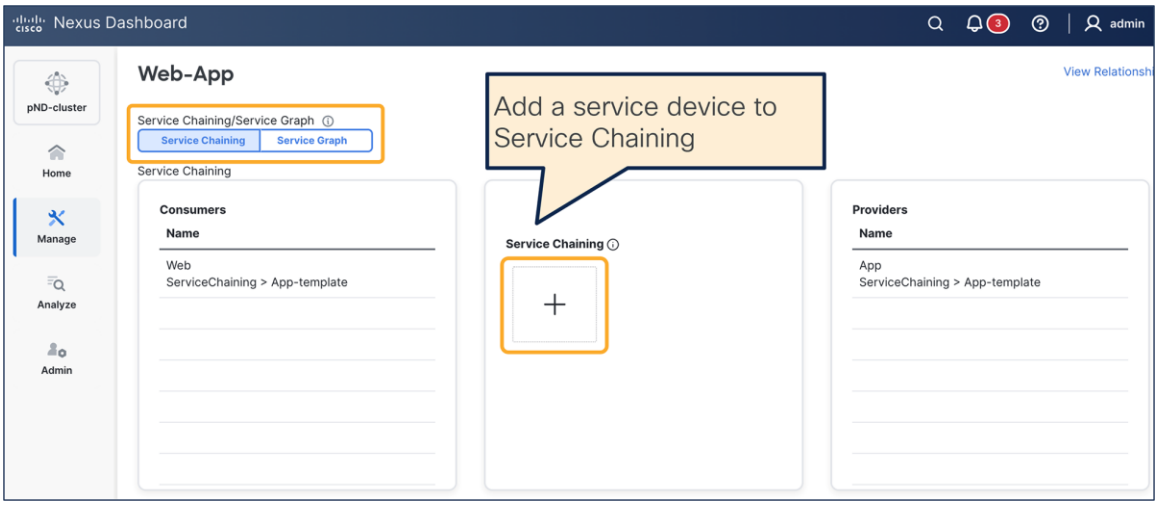


Figure 94.
Configure service chaining

Select device type, device, and interfaces. The figures below show the device settings for the following:

- One-arm mode firewall with PBR for both directions
- Two-arm mode firewall with PBR for both directions
- One-arm mode load balancer with PBR for the provider-to-consumer direction

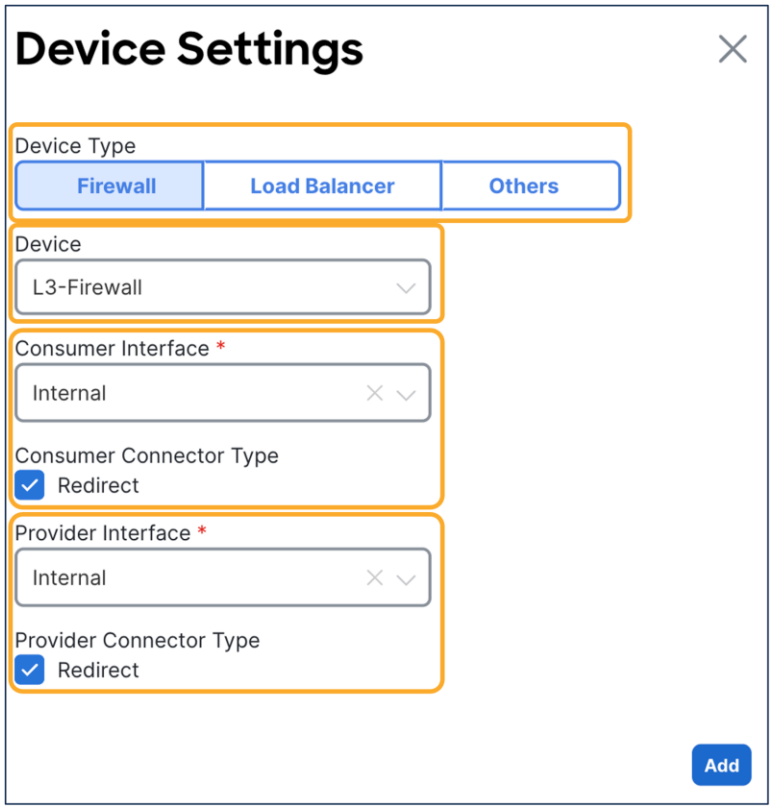


Figure 95.
One-arm mode firewall with PBR for both directions

© 2026 Cisco and/or its affiliates. All rights reserved.

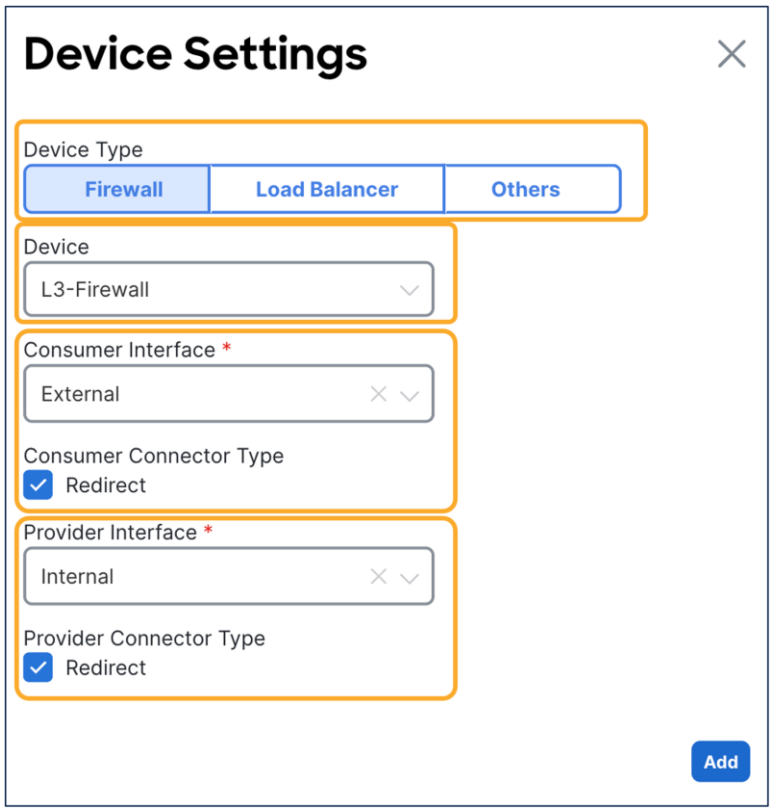


Figure 96.
Two-arm mode firewall with PBR for both directions

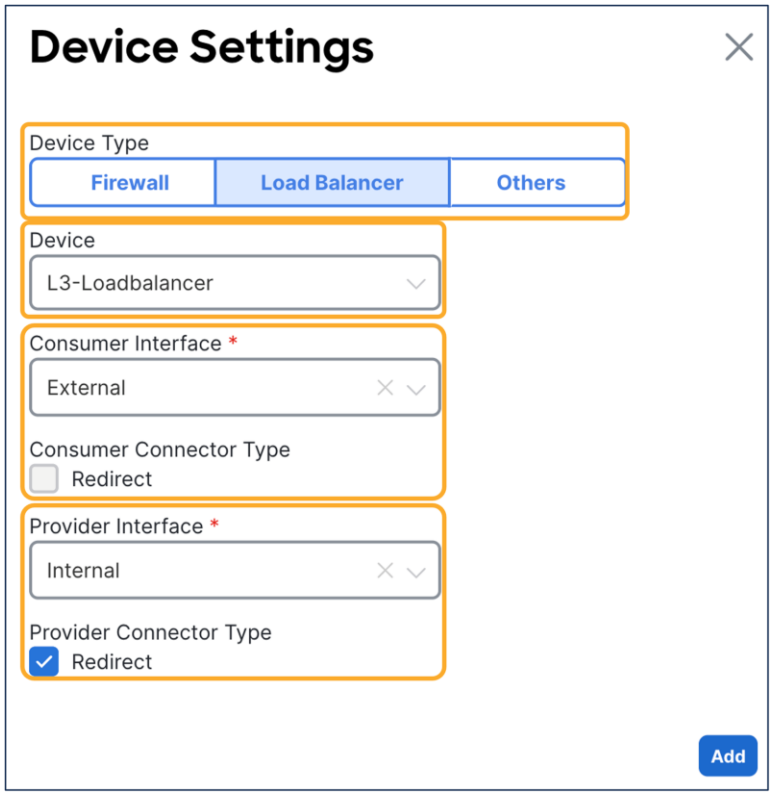


Figure 97.
Two-arm mode load balancer with PBR for the provider-to-consumer direction

© 2026 Cisco and/or its affiliates. All rights reserved.

If you have multiple nodes, add another service device and repeat the same steps. The figure below shows an example of a two-node service-chaining, in this case for firewall and load balancer.

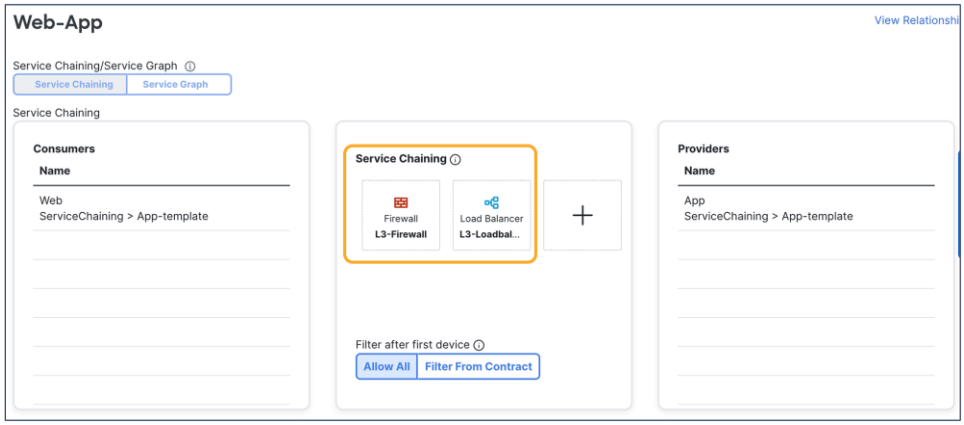


Figure 98.
Two node service chaining example

If the deployment is successfully done, the service graph, device-selection policy, and deployed graph Instance are created for the tenant on each APIC domain. This can be verified on APIC at Tenant > Services > L4-L7 > Device Selection Policies and Deployed Graph Instances.

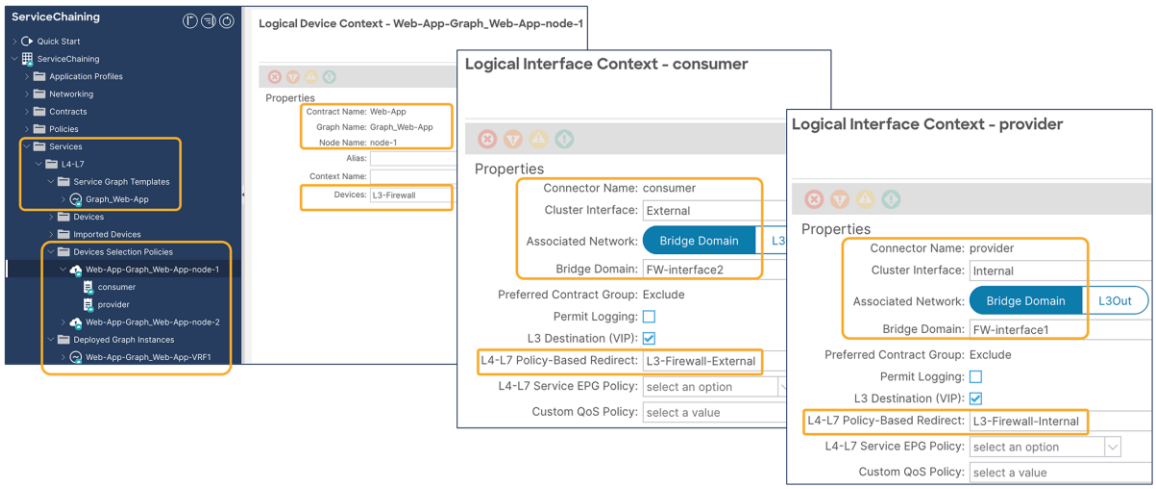


Figure 99.
Verify the configuration on APIC (Device Selection Policy and Deployed Graph Instance)

© 2026 Cisco and/or its affiliates. All rights reserved.

Example 2: vzAny-to-vzAny contract with PBR

This subsection covers a configuration step for vzAny-to-vzAny contract with PBR. vzAny-to-vzAny contract with PBR supports intra-VRF contracts only and it must be one-arm one-node bidirectional PBR.

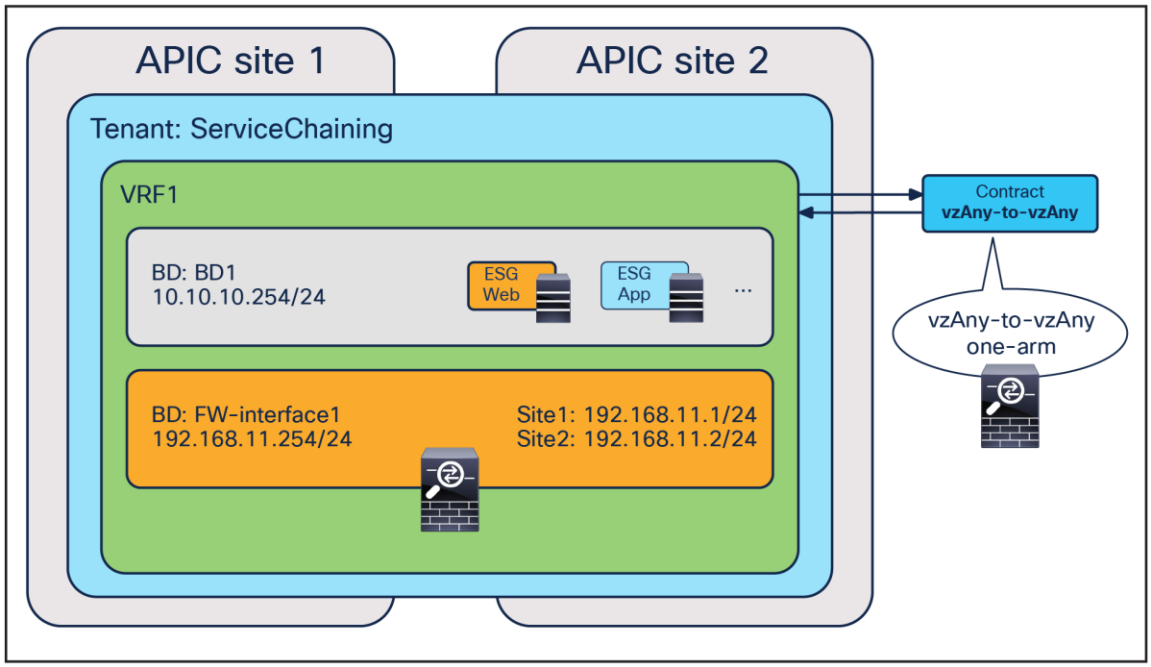


Figure 100.
vzAny-to-vzAny contract with PBR

Configure service chaining

Service insertion is configured by associating one or more service devices with a contract. This is provisioned from Configure > Tenant Template > Applications. Select your template and then the contract.

At the bottom of the contract's configuration, ensure that "Service Chaining" is selected. It shows the list of the consumers and the providers of the contract (if already configured). By clicking the "+" icon, you can add one or more service devices to create a service chaining between the consumers and the providers.

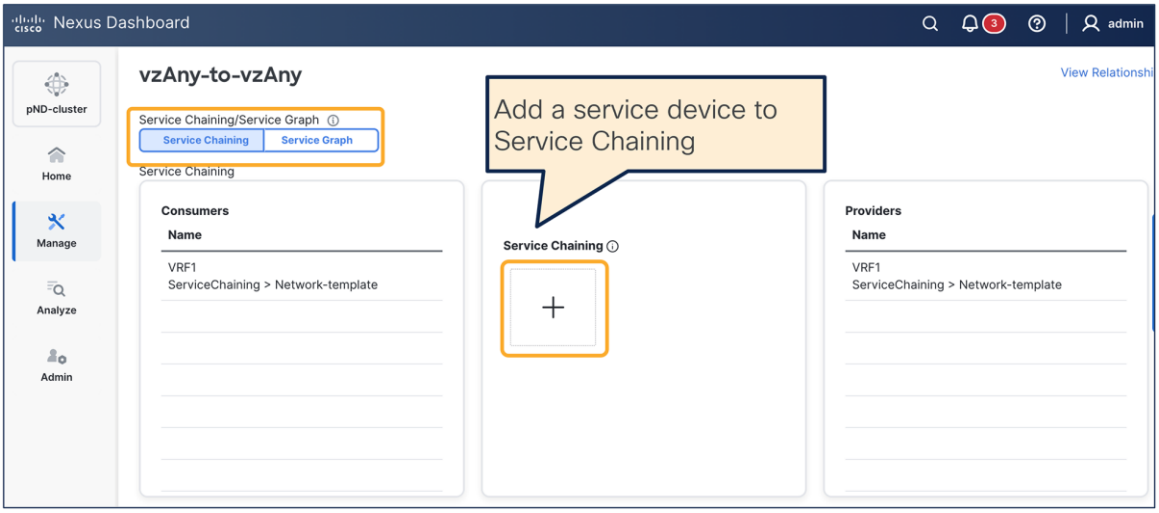


Figure 101.
Configure service chaining

© 2026 Cisco and/or its affiliates. All rights reserved.

Select device type, device, and interfaces. In the case of vzAny-to-vzAny contract with PBR, you must select the same interface and enable the redirect flag for both the consumer and the provider interfaces, because the use case is supported only with one-arm and one-node bidirectional PBR. If the service device was defined with only one interface, it would be automatically selected.

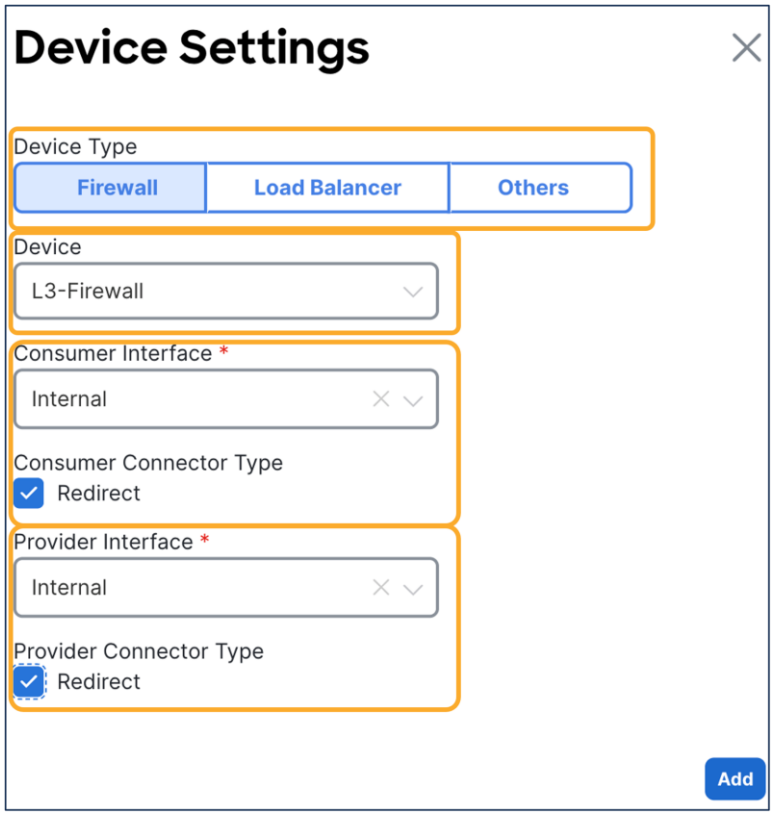


Figure 102.
Select device and interfaces

Confirm that the service device is added to the service chaining in the contract, and then deploy the template to the sites.

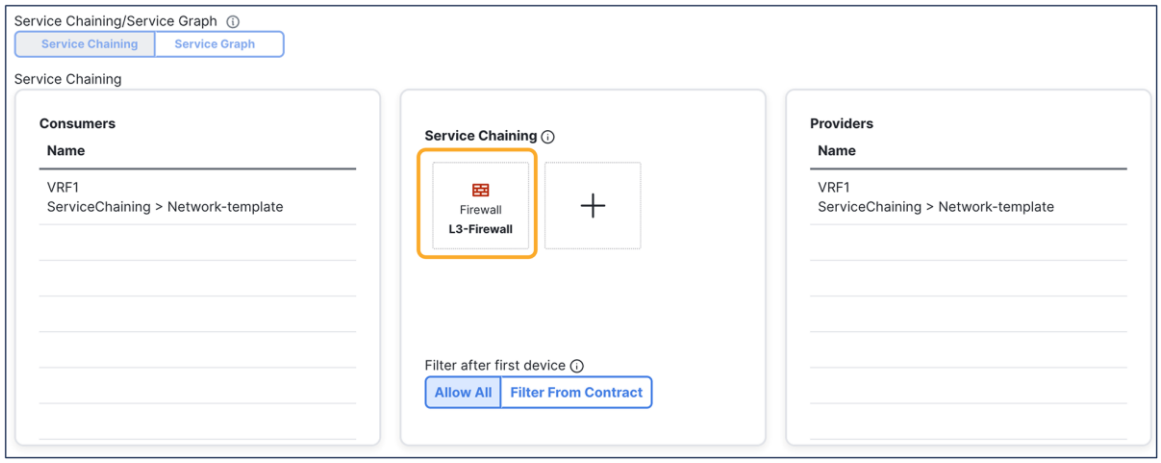


Figure 103.
Confirm cervices chaining configuration

© 2026 Cisco and/or its affiliates. All rights reserved.

If the deployment is successfully done, the service graph, device-selection policy, and deployed graph Instance are created for the tenant on each APIC domain. This can be verified on APIC at Tenant > Services > L4-L7 > Device Selection Policies and Deployed Graph Instances.

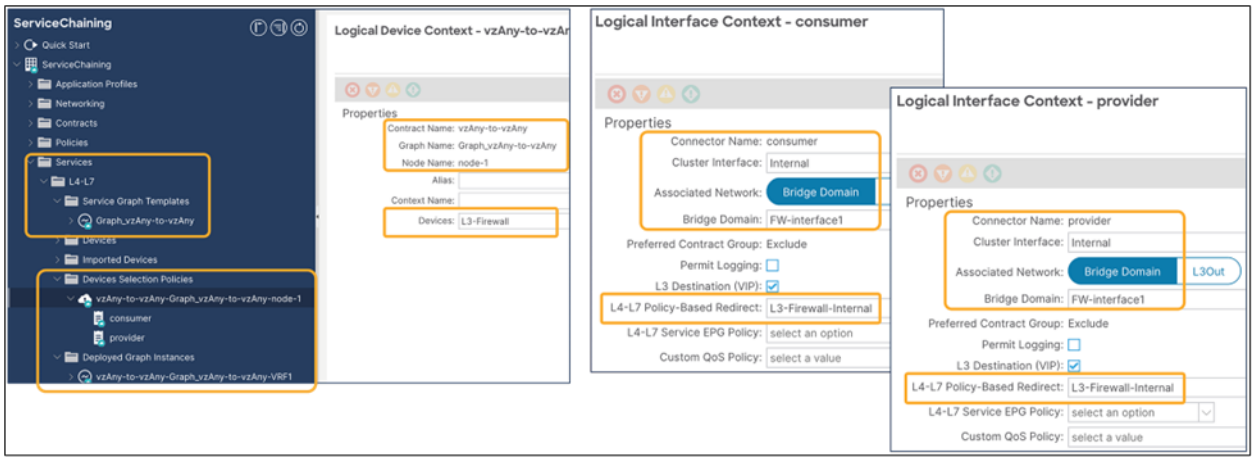


Figure 104.
Verify the configuration on APIC (Device Selection Policy and Deployed Graph Instance)

GUI and CLI output example for verification

Overview

The following steps are typical for troubleshooting. This section explains how to verify steps 2 and 3, which are specific to a service graph. This document does not cover general Cisco ACI endpoint learning or forwarding troubleshooting steps. For more information about Cisco ACI troubleshooting, refer to the following link: https://www.cisco.com/c/dam/en/us/td/docs/switches/datacenter/aci/apic/sw/4-x/troubleshooting/Cisco_TroubleshootingApplicationCentricInfrastructureSecondEdition.pdf.

1. Check if the communication between ESGs can be established without attaching the service graph with PBR to the contract:
   - Consumer and provider endpoints are learned.
   - Consumer and provider endpoints can communicate within the same site and across sites.
2. Verify the service-graph deployment (on each APIC):
   - Deployed graph Instances have no faults.
   - VLANs and class IDs for service nodes are deployed.
   - Service -node endpoints are learned.
3. Check that the traffic is successfully redirected:
   - Capture the traffic on the service node.
   - Check that the policy is properly programmed on the leaf nodes.
4. Check that the incoming traffic arrives on the consumer and provider endpoints.

© 2026 Cisco and/or its affiliates. All rights reserved.

Check that a service graph is deployed

Deployed graph instances

After a service graph is successfully applied, you can see the deployed graph instance for each contract with a service graph (Figure 105). If a service graph instantiation fails, you will see faults in the deployed graph instance.

The location is Tenant > Services > L4-L7 > Deployed Graph instances.

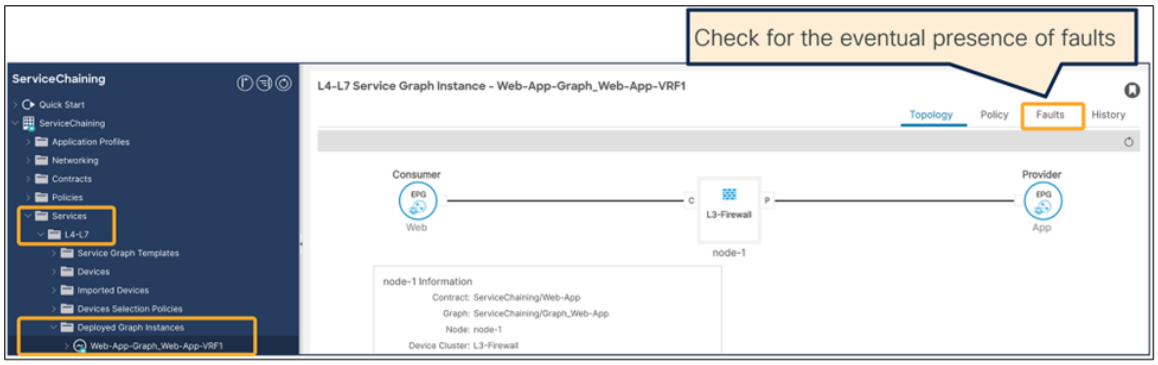


Figure 105.
Check that the graph instance is deployed on APIC

VLANs and class IDs for service node

If you see a fault, it's most likely because there is something wrong with the APIC configuration. For example, the encap VLAN is not available in the domain used for the L4–L7 device.

Once the service graph is successfully deployed without any fault in the deployed graph instances, EPGs and BDs for service node get created. The figure below shows where to find the class IDs for the service-node interfaces (Service EPGs). In this example, FW-external (consumer connector) class ID is 26 and FW-internal (provider connector) class ID is 10931.

The location is Tenant > Services > L4-L7 > Deployed Graph instances > Function Nodes.

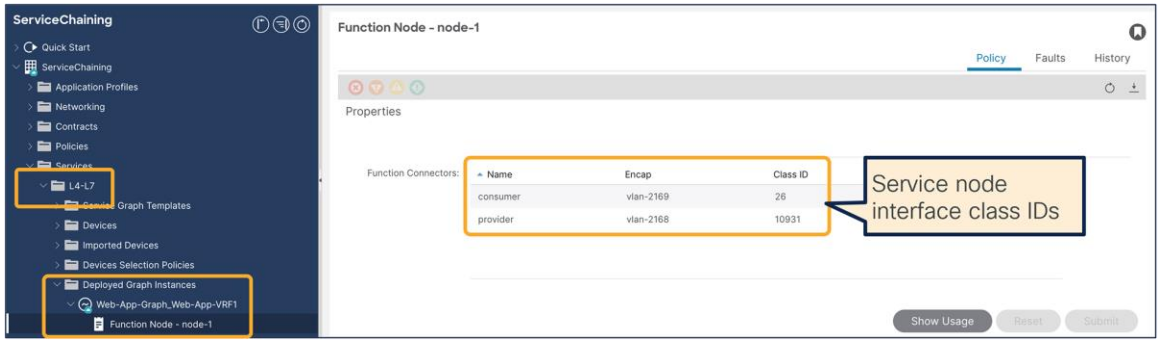


Figure 106.
Service-node interface class ID

© 2026 Cisco and/or its affiliates. All rights reserved.

These VLANs are deployed on the service leaf node where the service nodes are connected. VLAN and endpoint learning status can be checked by using “show VxLAN extended” and ‘show endpoint” on the service leaf node CLI. If you don’t see the IPs of service nodes learned as endpoints in the ACI fabric, most likely it’s a problem of connectivity or a configuration issue between the service leaf and the service node. Please check the following statuses that might have something wrong:

- Interface status on leaf interfaces connected to the service node.
- The leaf interface path and VLAN encap.
- The service node VLAN and IP address.
- The intermediate switch VLAN configuration if you have it between the service leaf node and the service node.

Check if the traffic is redirected

Capture the traffic on the service node

If end-to-end traffic stops working once you enable PBR, even though the service-node endpoints are learned in the ACI fabric, the next troubleshooting step is to check if traffic is redirected and where the traffic is dropped.

To verify whether traffic is redirected to the service node, you can enable capture on the PBR destination. The figure below shows an example of where you should see the traffic redirected. In this example, ESG Web to ESG App (ESG-to-ESG) contract with PBR is configured and 10.10.11.11 in site1 Web ESG tries to access 10.10.12.12 in site2 App ESG. Because the endpoint 10.10.12.21 is in the provider ESG, the PBR policy is applied in site2, thus traffic should be seen on the PBR destination in site2.

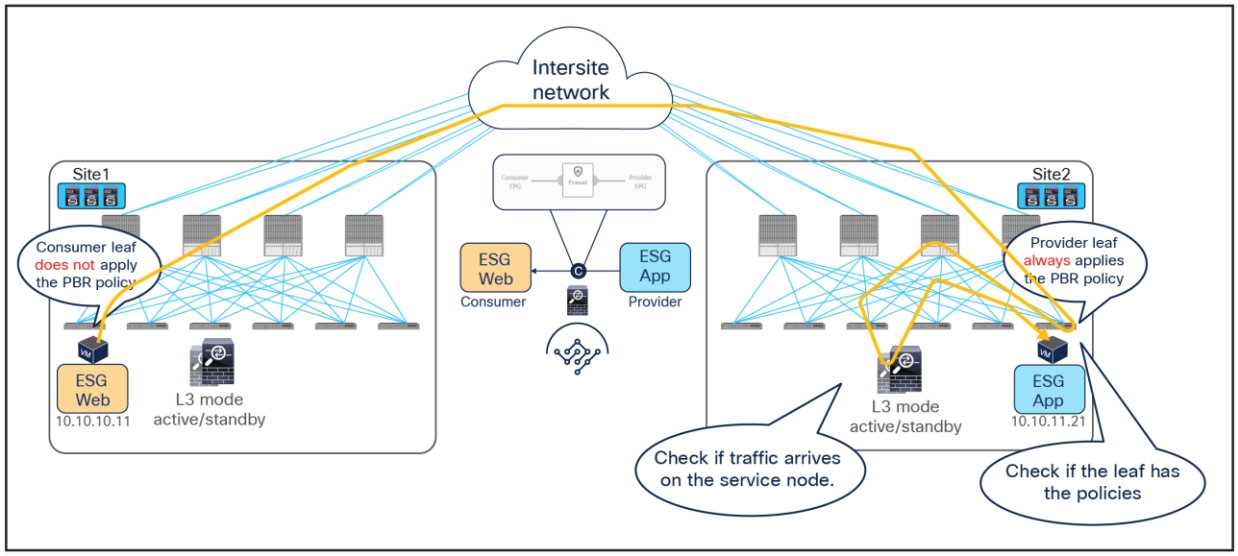


Figure 107.
Traffic flow example

© 2026 Cisco and/or its affiliates. All rights reserved.

If you see that consumer-to-provider traffic is received on the service node but not on the provider endpoint, please check the following, which are common mistakes:

- Service node routing table reaches the provider subnet (The service node must be able to reach the provider and consumer subnets).
- Service node security policy such as ACL permits the traffic.

Check policies on leaf nodes

If you do not see the traffic being received by the service node, you may need to take a look at the leaf nodes to see if policies are programmed on the switch nodes to permit or redirect the traffic.

Note: The policies are programmed based on ESG deployment status on the leaf. The show-command output in this section uses the leaf that has consumer ESG, provider ESG, and ESGs for the service node.

This sub-section also covers an important consideration for “Fabric-aware Policy Enforcement Mode”. Although the use of vzAny and combining multiple EPGs to an ESG will help to consume less TCAM resources, please be aware that enabling “Fabric-aware Policy Enforcement Mode” for PBR use cases will increase TCAM resource consumption in the VRF. Although this is to take different actions for the provider-to-consumer direction depending on the source site, this behavior is not specific to PBR contract. If “Fabric-aware Policy Enforcement Mode” is enabled on a VRF, a set of zoning-rules for each site is programmed for rules in the VRF even though it’s a permit contract. Please see the examples below for detail.

Example1: ESG-to-ESG contract with PBR

The figures below show the zoning-rule status before and after a service graph deployment in site1. In this example, the VRF scope ID is 3112962, the consumer ESG (Web ESG) class ID is 24, and the provider ESG (App ESG) class ID is 23.

Before enabling Fabric-aware Policy Enforcement Mode” and deploying the service graph, a leaf node has two permit zoning-rules that are for traffic between the Web ESG and the App ESG.

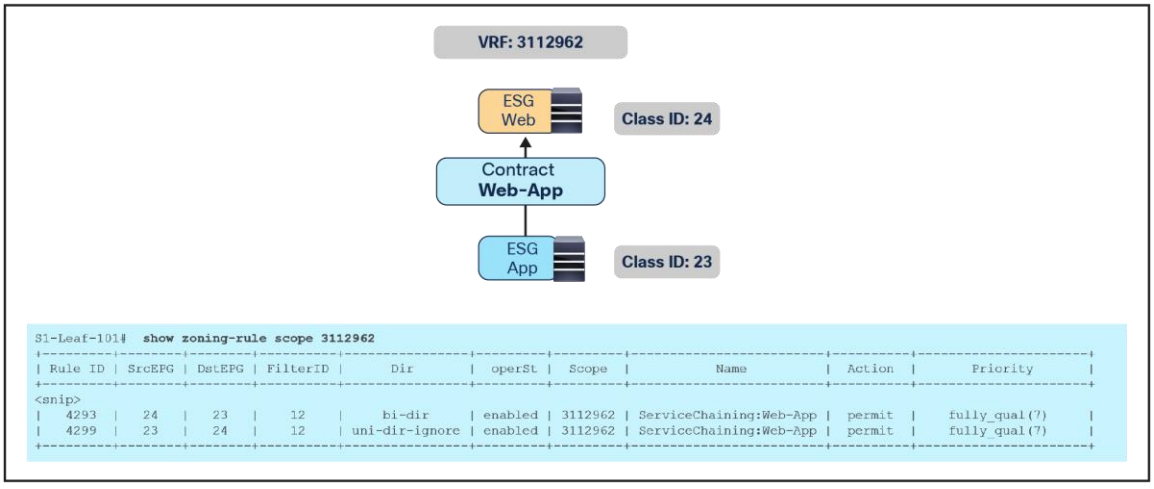


Figure 108.
ESG class IDs and zoning-rules (before service-graph deployment) without “Fabric-aware Policy Enforcement Mode”

© 2026 Cisco and/or its affiliates. All rights reserved.

Table 2. Permit rules before service-graph deployment without “Fabric-aware Policy Enforcement Mode”

Source class ID	Destination class ID	Action
24 (Web ESG)	23 (App ESG)	Permit
23 (App ESG)	24 (Web ESG)	Permit

When “Fabric-aware Policy Enforcement Mode” is enabled on the VRF, a set of zoning-rules is programmed for each site. Thus, before deploying the service graph, a leaf node has two permit zoning-rules for each site, which will be four rules if the contract is deployed for traffic between two sites.

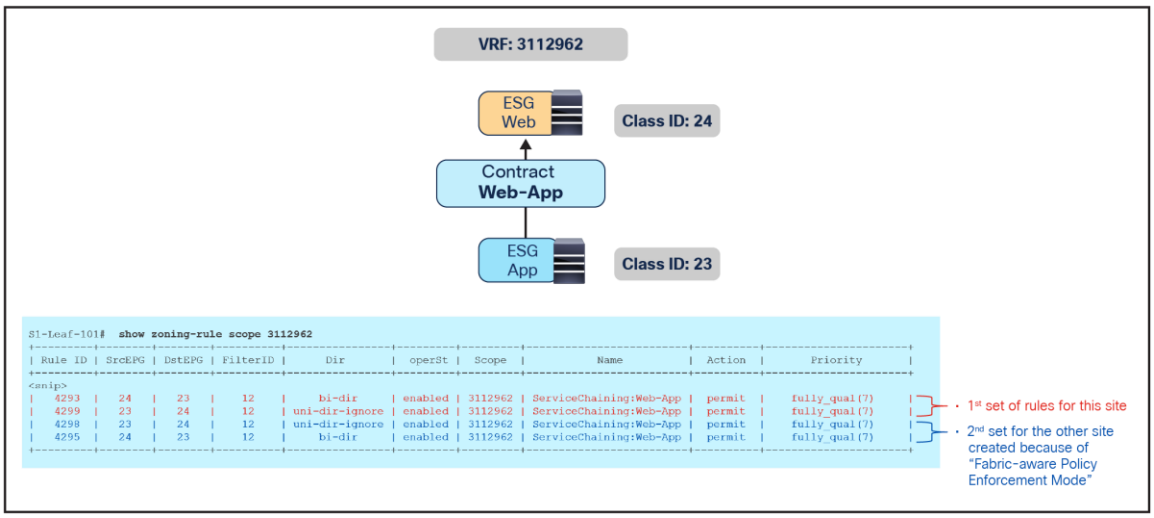


Figure 109.
ESG class IDs and zoning-rules (before service-graph deployment) with “Fabric-aware Policy Enforcement Mode”

Table 3. Permit rules before service-graph deployment without “Fabric-aware Policy Enforcement Mode”

Source class ID	Destination class ID	Action	Source site
24 (Web ESG)	23 (App ESG)	Permit	Site1 (local)
23 (App ESG)	24 (Web ESG)	Permit	Site1 (local)
24 (Web ESG)	23 (App ESG)	Permit	Site2
23 (App ESG)	24 (Web ESG)	Permit	Site2

© 2026 Cisco and/or its affiliates. All rights reserved.

Once the service graph is deployed, the zoning-rules get updated and the service node's class IDs are inserted based on the service graph configuration. Same with the example above, the set of rules is created for each source site. Please see the zoning-rules output in the figure below to understand the reason for this. ESG-to-ESG PBR uses a special behavior explained in Figure 18 to redirect traffic back to the PBR destination in the source site if the ingress provider leaf in the source site has not learned the destination endpoint. Thus, leaf nodes need to obtain different redirect zoning-rules for each site. In this example, 23-to-24 zoning-rules are for the provider-to-consumer traffic where the special behavior is required: one zoning-rule is to redirect traffic to the PBR destination in a local site (destgrp-8), and the other zoning-rule is to redirect traffic to the PBR destination in another site (destgrp-9).

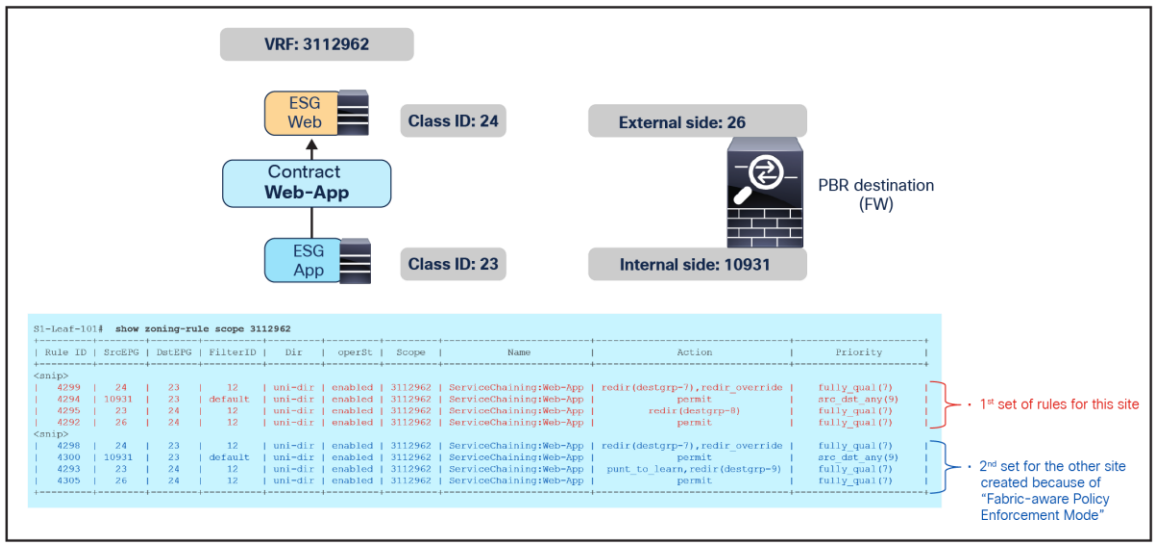


Figure 110.
ESG class IDs and zoning-rules (after service-graph deployment) with "Fabric-aware Policy Enforcement Mode"

Table 4. Permit and redirect rules after service-graph deployment with "Fabric-aware Policy Enforcement Mode"

Source class ID	Destination class ID	Action	Source site
24 (Web ESG)	23 (App ESG)	Redirect to destgrp-7 (FW-consumer in local site) with redir_override	Site1 (local)
10931 (FW-internal)	23 (App ESG)	Permit	Site1 (local)
23 (App ESG)	24 (Web ESG)	Redirect to destgrp-8 (FW-provider in local site)	Site1 (local)
26 (FW-external)	24 (Web ESG)	Permit	Site1 (local)
24 (Web ESG)	23 (App ESG)	Redirect to destgrp-7 (FW-consumer in local site) with redir_override	Site2
10931 (FW-internal)	23 (App ESG)	Permit	Site2
23 (App ESG)	24 (Web ESG)	Redirect to destgrp-9 (FW-provider in Site2) with punt_to_learn	Site2
26 (FW-external)	24 (Web ESG)	Permit	Site2

© 2026 Cisco and/or its affiliates. All rights reserved.

Note: The zoning-rule for ESG-to-ESG contract with PBR is created with action “redir_override”: this is required in the specific PBR deployment with Cisco ACI Multi-Site. With this action, the hardware creates two entries to take different actions depending on whether the destination (provider) is in the local leaf or not. If the destination is in the local leaf, the PBR policy is applied. If the destination is not in the local leaf, the traffic is just permitted so that the redirection can instead happen on the leaf in the site where the provider endpoint resides. That’s how to get a provider leaf to always apply PBR policy. The same mechanism is used for vzAny-to-ESG contract with PBR.

Important Note: Because of this behavior, in the case of ESG-to-ESG contract with PBR, it is critical to ensure that it is always possible to clearly identify a consumer and a provider side in zoning-rules for each given contract relationship between ESGs.

This means that the same ESG should never consume and provide the same contract and the definition of different contracts may be needed depending on the specific deployment scenario.

Also, if two different contracts were applied between the same pair of ESGs (so to be able to differentiate the provider and consumer ESG for each of them), it is critical to ensure that the zoning-rules created by those two contracts don’t have overlapping rules with same contract and filter priorities. Defining zoning-rules with the same priority that identify the same type of traffic could lead to a not deterministic forwarding behavior (creating asymmetric traffic through different firewalls). As a typical example, it would not work to create two contracts both using a “permit IP” rule to redirect all the traffic. If one contract is “permit IP” and the other contract is “permit ICMP only”, the zoning-rules created by the contract with “permit ICMP only” have higher priority because it’s more specific filter though the zoning-rule priorities are same. The table and figure below illustrate this example. In this case, ICMP traffic between Web and App ESGs is always redirected on the leaf where an endpoint in the Web ESG (the provider of the Contract2) resides whereas other traffic between Web and App ESGs is always redirected on the leaf where an endpoint in the App ESG (the provider of the Contract1) resides.

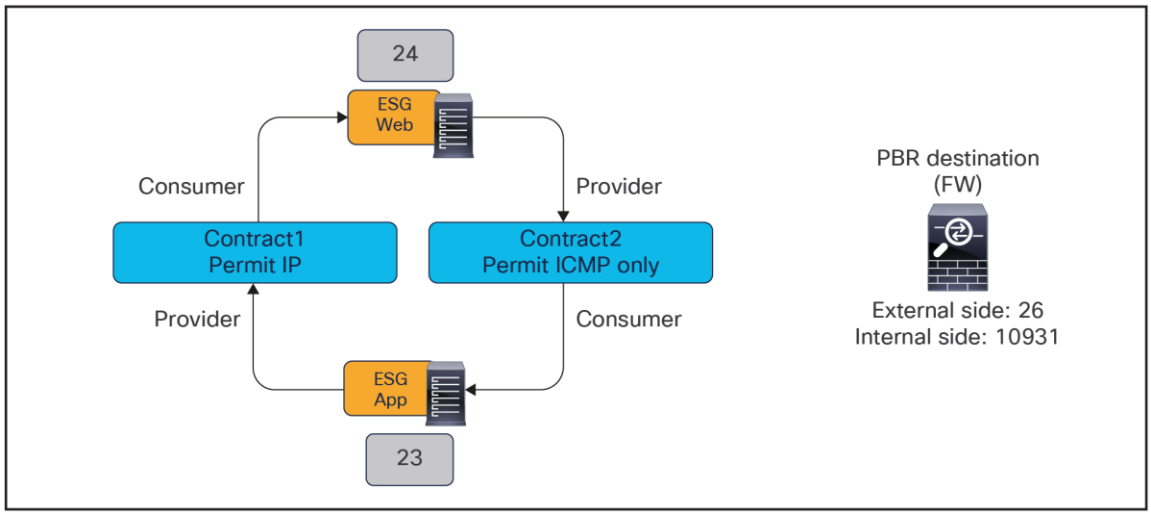


Figure 111.
ESG class IDs and contracts (after service-graph deployments)

© 2026 Cisco and/or its affiliates. All rights reserved.

Table 5. Permit and redirects rules with service graphs

	Source class ID	Destination class ID	Filter	Action	Zoning-rule priority	Source site
Contract1	24 (Web ESG)	23 (App ESG)	Permit IP	Redirect to destgrp-7 (FW-consumer in local site) with redir_override	7	Site1 (local)
	10931 (FW-internal)	23 (App ESG)	Permit IP	Permit	7	Site1 (local)
	23(App ESG)	24 (Web ESG)	Permit IP	Redirect to destgrp-8 (FW-provider in local site)	7	Site1 (local)
	26 (FW-external)	24 (Web ESG)	default (permit any)	Permit	9	Site1 (local)
Contract1	24 (Web ESG)	23 (App ESG)	Permit IP	Redirect to destgrp-7 (FW-consumer in local site) with redir_override	7	Site2
	10931 (FW-internal)	23 (App ESG)	Permit IP	Permit	7	Site2
	23(App ESG)	24 (Web ESG)	Permit IP	Redirect to destgrp-9 (FW-provider in Site2) with punt_to_learn	7	Site2
	26 (FW-external)	24 (Web ESG)	default (permit any)	Permit	9	Site2
Contract2	24 (Web ESG)	23 (App ESG)	Permit ICMP	Redirect to destgrp-7 (FW-consumer in local site)	7	Site1 (local)
	10931 (FW-internal)	23 (App ESG)	Permit ICMP	Permit	7	Site1 (local)
	23 (App ESG)	24 (Web ESG)	Permit ICMP	Redirect to destgrp-8 (FW-provider in local site) with redir_override	7	Site1 (local)
	26 (FW-external)	24 (Web ESG)	default (permit any*)	Permit	9	Site1 (local)

© 2026 Cisco and/or its affiliates. All rights reserved.

	Source class ID	Destination class ID	Filter	Action	Zoning-rule priority	Source site
Contract2	24 (Web ESG)	23 (App ESG)	Permit ICMP	Redirect to destgrp-7 (FW-consumer in local site)	7	Site2
	10931 (FW-internal)	23 (App ESG)	Permit ICMP	Permit	7	Site2
	23 (App ESG)	24 (Web ESG)	Permit ICMP	Redirect to destgrp-9 (FW-provider in Site2) with punt_to_learn	7	Site2
	26 (FW-external)	24 (Web ESG)	default (permit any*)	Permit	9	Site2

*By default, the zoning-rule for the traffic from provider side of the service node to the provider EPG uses default filter (permit any) even though the filter used in the contract is not default filter. This behavior can be changed by using "filter-from-contract" option in the Service Graph.

For more information about zoning-rules and priorities, please refer to the Contract priorities section in the ACI Contract Guide: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-743951.html#Contractpriorities.

The figure below shows how to check the destinations for a redirect destination group (destgrp). In addition to the PBR destinations in a local site such as destgrp-7 and destgrp-8 in this example, the PBR destination in a remote site is also listed.

```
S1-Leaf-101# show service redir info
======================================================================================================================================================
LEGEND
TL: Threshold(Low)  |  TH: Threshold(High)  |  HP: HashProfile  |  HG: HealthGrp  | BAC: Backup-Dest |  TRA: Tracking | RES: Resiliency | W: Weight
======================================================================================================================================================
List of Dest Groups
GrpID Name            destination                          HG-name                                          BAC W   operSt   operStQual      TL  TH  HP  TRA RES
===== ====            ===========                          =============                                    === ===  =======  ===========     === === === === ===
7     destgrp-7       dest-[192.168.12.1]-[vxlan-3112962]  ServiceChaining::L3-Firewall--ndo--implct        N   1    enabled  no-oper-grp     0   0   sym yes no
9     destgrp-9       dest-[192.168.11.2]-[vxlan-3112962]  ServiceChaining::L3-Firewall--ndo--implct::2      N   1    enabled  no-oper-grp     0   0   sym yes no
8     destgrp-8       dest-[192.168.11.1]-[vxlan-3112962]  ServiceChaining::L3-Firewall--ndo--implct        N   1    enabled  no-oper-grp     0   0   sym yes no
```

Figure 112.
Check redirect group

If you check the same information on the APIC and on the leaf nodes in site2, you will see similar outputs with different class IDs because each site uses different class IDs. With ACI Multi-Site, the spines have translation tables to change the class IDs for inter-site traffic so that policy can be maintained consistently across sites (namespace normalization).

© 2026 Cisco and/or its affiliates. All rights reserved.

The figure below shows the translation tables and class IDs in site1 and site2.

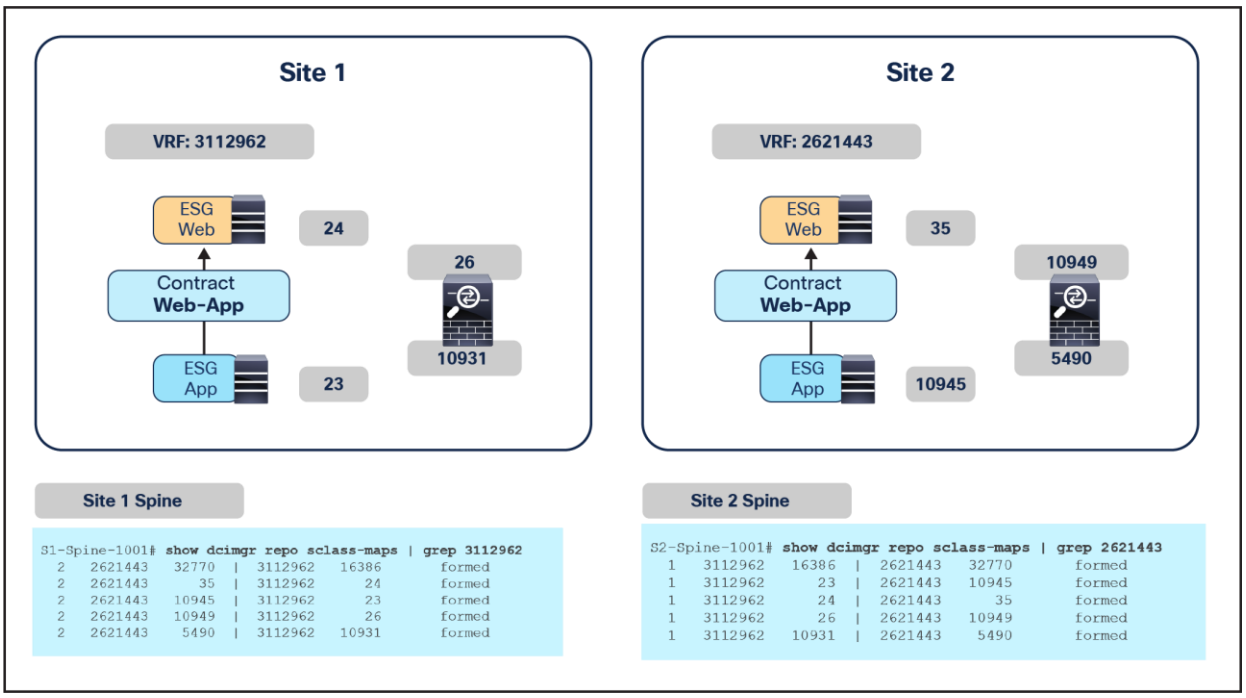


Figure 113.
Translation tables on spine nodes

Table 6. ESG class ID translation

ESG and Service EPG	Site1 class ID	Site1 VRF	Site2 class ID	Site2 VRF
Web ESG	24	3112962	35	2621443
App ESG	23	3112962	10945	2621443
FW-external	26	3112962	10949	2621443
FW-internal	10931	3112962	5490	2621443

Note: 16386 in site1 and 32770 in site2 are the VRF class IDs.

© 2026 Cisco and/or its affiliates. All rights reserved.

Example2: vzAny-to-vzAny contract with PBR

In the case of vzAny-to-vzAny contract with permit, by default, 0-to-0 permit zoning-rules are programmed on the leaf nodes; 0 represents vzAny. If there is no port-number specified in the filter used in the contract, it will be one 0-to-0 zoning-rule because the consumer-to-provider and provider-to-consume rules are identical.

Since “Fabric-aware Policy Enforcement Mode” is enabled on the VRF for vzAny PBR, a set of zoning-rules is programmed for each site. Please see the zoning-rules output in the figure below to understand the reason for this. vzAny PBR uses a special behavior explained in Figure 18 to redirect traffic back to the PBR destination in the source site if the ingress leaf in the source site has not learned the destination endpoint. Thus, leaf nodes need to obtain different redirect zoning-rules for each site. In this example, one zoning-rule is to redirect traffic to the PBR destination in a local site (destgrp-7), and the other zoning-rule is to redirect traffic to the PBR destination in another site (destgrp-8).

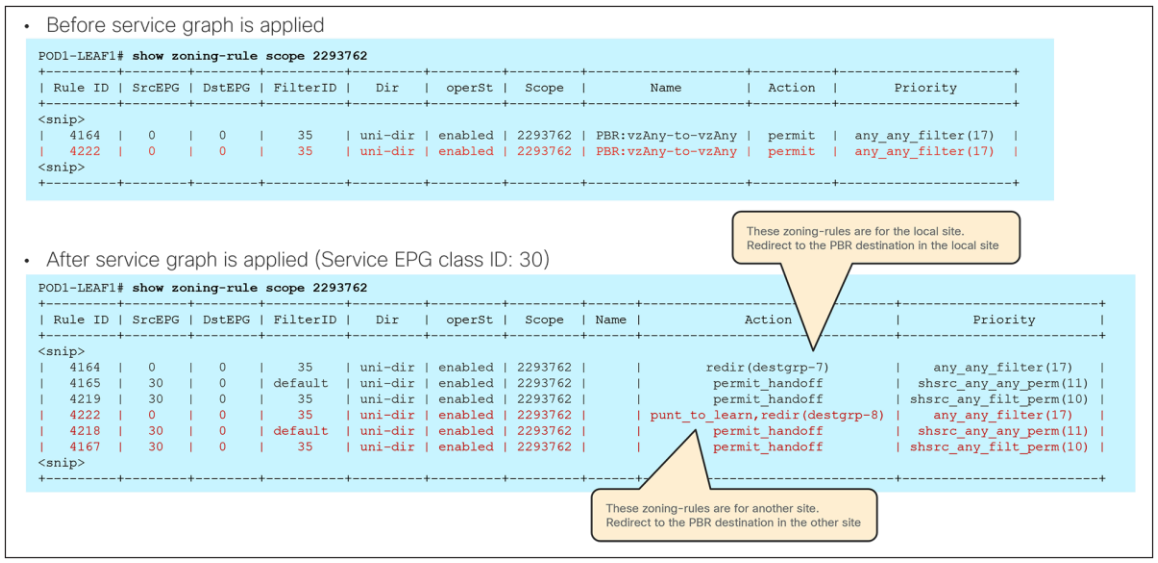


Figure 114.
Zoning-rules for vzAny-to-vzAny contract with PBR

© 2026 Cisco and/or its affiliates. All rights reserved.

Check where the policy is applied

If you don't see the traffic in the service node even though you see zoning-rules and translation tables accordingly programmed on the switch nodes, the traffic might be dropped somewhere else or the policy might not be enforced on the leaf node. To check specific forwarding information on each ACI switch node, ELAM (Embedded Logical Analyzer Module) Assistant is available under Operations tab on APIC.

Where policy is applied could be different depending on the endpoint-learning status. To identify if endpoint is learned through conversational learning or not, you can use the "show system internal epm" command.

```
POD1-LEAF1# show system internal epm endpoint vrf PBR:VRF1
VRF : PBR:VRF1 ::: Context id : 20 ::: Vnid : 2293762
<snip>
MAC : 0000.0000.0000 ::: Num IPs : 1
IP# 0 : 10.10.2.100 ::: IP# 0 flags :  ::: l3-sw-hit: No
Vlan id : 0 ::: Vlan vnid : 0 ::: VRF name : PBR:VRF1
BD vnid : 0 ::: VRF vnid : 2293762
Phy If : 0 ::: Tunnel If : 0x18010017
Interface : Tunnel23
Flags : 0x8000080004400 ::: sclass : 49171 ::: Ref count : 3
EP Create Timestamp : 10/16/2023 22:15:49.259045
EP Update Timestamp : 10/16/2023 23:42:58.535504
EP Flags : IP|sclass|timer|control-ep|

# vsh_lc -c "show system internal epm endpoint ip 10.10.2.100"

MAC : 0000.0000.0000 ::: Num IPs : 1
IP# 0 :  10.10.2.100

VRF name : PBR:VRF1 ::: VRF vnid : 2293762
phy if : 0 ::: tunnel if : 0x1801001f ::: Interface : Tunnel31
Ref count : 3 ::: sclass : 49154
 ::: Learns Src: EPM
EP Flags : IP|sclass|timer|control-ep|
Aging: Timer-type : control-ep ::: Timeout-left : 86067 ::: Hit-bit : Yes ::: Timer-reset count : 0
```

Figure 115.
Check to see if endpoint has been learned through conversational learning

FAQ

This section covers frequently asked questions regarding this solution.

Q. Can we have multiple ESGs consuming or providing the same contract with an attached service graph with PBR?

A. Yes, the same considerations apply as for a single-site deployment. Figure 116 illustrates an example. You can even mix a site-local ESG and a stretched ESG across sites, though the service EPG must always be stretched. In the example given in Figure 116, ESG2 and ESG4 are providers, and ESG1 and ESG3 are consumers.

© 2026 Cisco and/or its affiliates. All rights reserved.

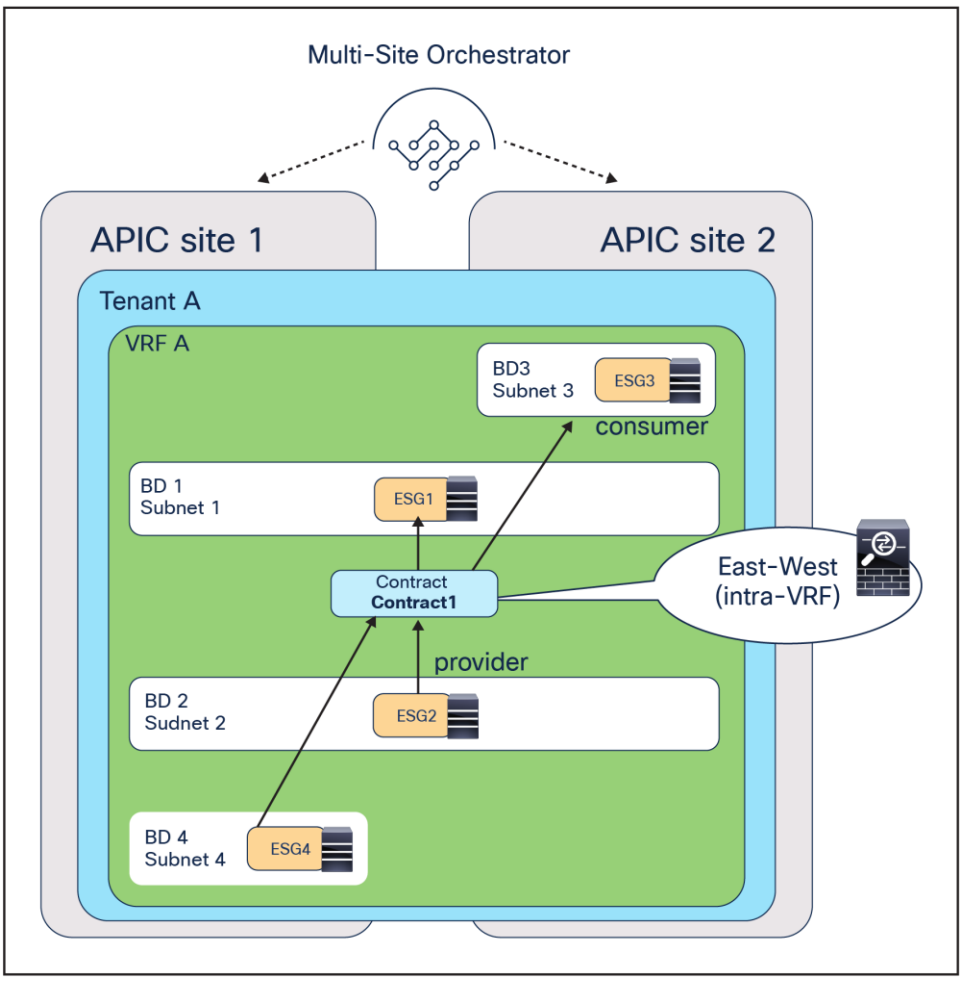


Figure 116.
Multiple ESGs consuming and providing the same contract with PBR

Q. Would an inter-tenant inter-VRF design work?

A. Yes except vzAny-to-vzAny that works for intra-VRF only.

Q. Can we use a managed-mode service graph?

A. No, Cisco ACI Multi-Site supports unmanaged-mode service graphs only. Starting from Cisco ACI Release 5.2(1), APIC supports unmanaged-mode service graphs only.

Q. Can we redirect traffic to a service node in a different site?

A. No, the use of a PBR destination in a remote site is not supported. Each site must deploy local service node(s).

Q. Do we have scale considerations?

A. Please take a look at the Cisco ACI verified scalability guide: https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html.

© 2026 Cisco and/or its affiliates. All rights reserved.

Q. Can the VIP of the load balancer be defined outside of the Self IP BD subnet?

A. Yes. In that case, L3Out is required for announcing VIP reachability to the fabric. Starting from NDO Release 4.2(3e), the use of an L3Out as a service-node connector is supported without PBR. If PBR is required on the connector, a BD must be used instead of an L3Out.

Q. What ESG selectors are available in Multi-Site?

A. The following ESG selectors are supported in Multi-Site: EPG Selector, IP Subnet Selector, Tag Selector, External EPG selector, external IP subnet selector. Please see Nexus Dashboard document for detail

Q. Why do “Fabric-aware Policy Enforcement Mode” and “L3 Multicast” need to be enabled on the VRF for ESG and vzAny PBR?

A. The reason “Fabric-aware Policy Enforcement Mode” needs to be enabled is because the VRF needs additional zoning-rules, which is explained in Figure 117. The reason “L3 Multicast” needs to be enabled is because conversational learning uses a unicast or multicast control packet, depending on the use cases. The figures below illustrate examples. For endpoint-to-endpoint traffic, a unicast control packet is used for conversational learning because the source class ID is translated by spine for endpoint-to-endpoint communication.

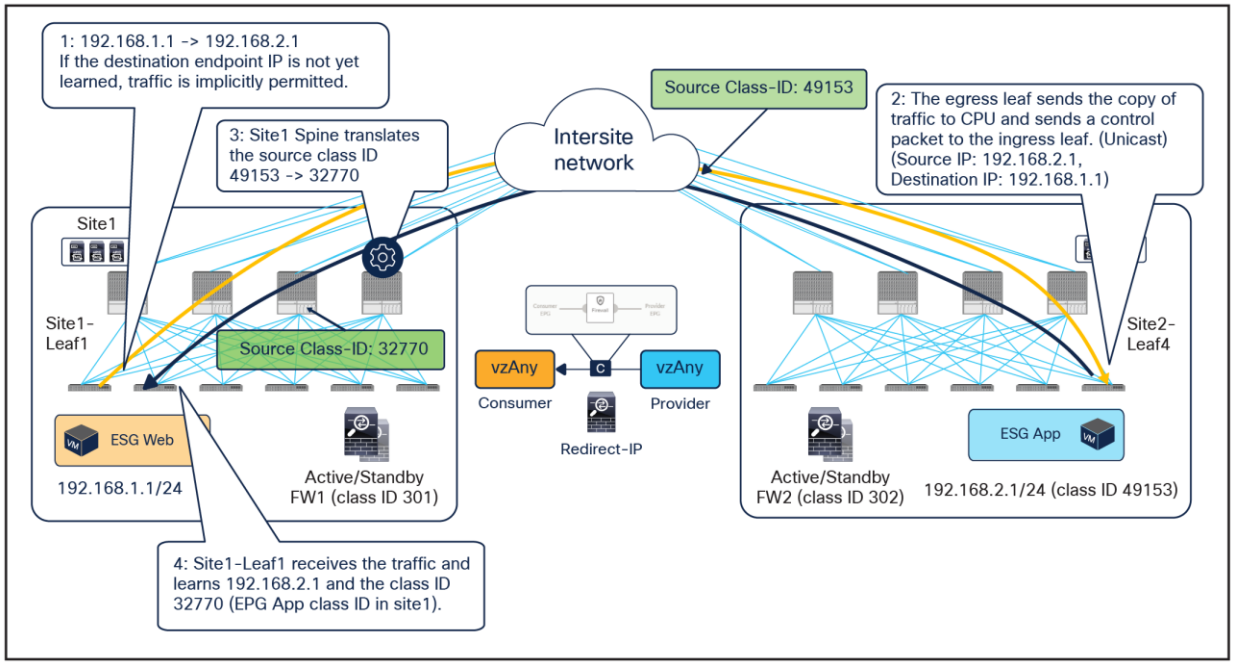


Figure 117.
Conversational learning for endpoint-to-endpoint communication (unicast)

© 2026 Cisco and/or its affiliates. All rights reserved.

Differently from endpoint-to-endpoint traffic, the class ID translation is not performed for the unicast traffic toward an inter-site L3Out (please see the inter-site L3Out section in Cisco ACI Multi-Site Architecture White Paper for details). Because of this, for external-to-endpoint traffic, the egress compute leaf cannot use a unicast control packet for conversational learning (since this conversational learning packet will be in the endpoint-to-external direction that does not carry the translated source class ID). Therefore, to circumvent this, the egress compute leaf uses a multicast control packet instead of unicast in the endpoint-to-external direction. The class ID translation is performed even if one of the multicast destinations is a border leaf with an inter-site L3Out.

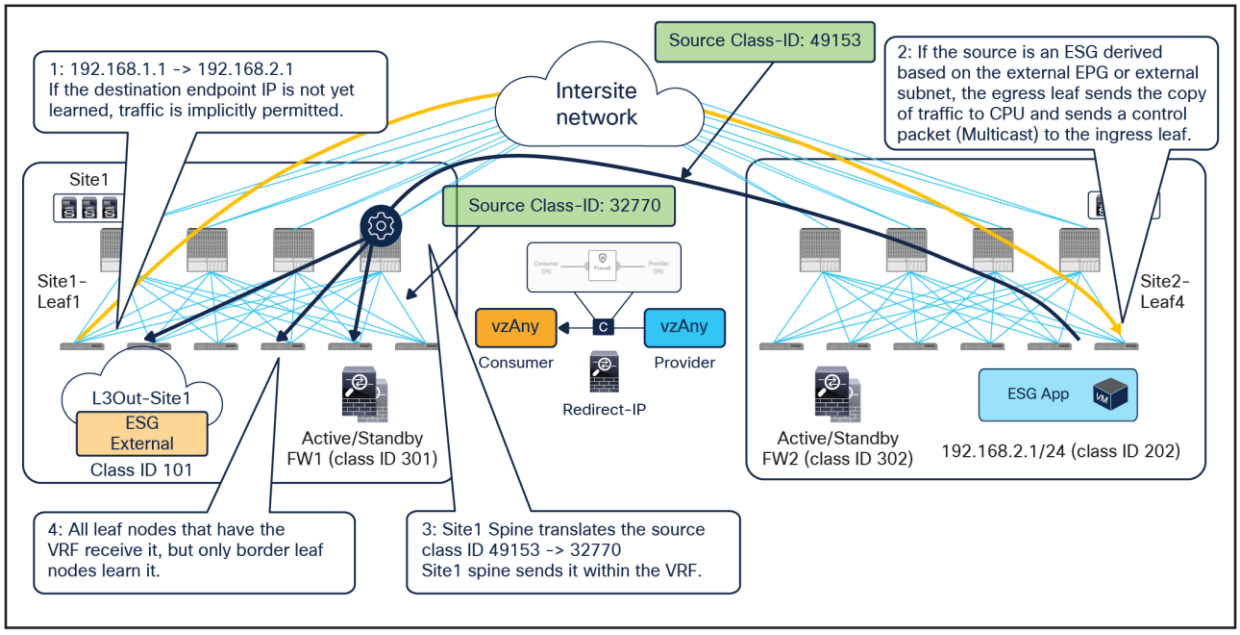


Figure 118.
Conversational learning for external-to-endpoint communication (multicast)

Conclusion

There are three different deployment models to integrate service nodes with Cisco ACI Multi-Site fabrics:

- Independent active/standby service-node pair in each site (recommended).
- Active/standby service-node pair connected to separate sites (not recommended).
- Active/active service-node cluster stretched across separate sites (not recommended).

© 2026 Cisco and/or its affiliates. All rights reserved.

The use of PBR is the recommended approach to integrate independent firewall pairs connected to separate sites. PBR destinations can be L1/L2/L3 service nodes. Examples and considerations are summarized in Table 7.

Table 7. Service integration modes for Cisco ACI Multi-Site fabric

Service node	Independent service node in each site (recommended)
Transparent (L1/L2) mode firewall	Yes • ACI is gateway; use PBR
Routed mode (L3) firewall	Yes • ACI is gateway; use PBR or • Connect firewall as an L3Out external device with host-route advertisement (north-south)
Routed mode load balancer	Yes • NAT on load balancer or • ACI is gateway; use PBR for return traffic.

Multi-Site PBR supports the following contract configurations:

- Intra-VRF and inter-VRF ESG-to-ESG contract (one/two-arm, multiple nodes service insertion with bidirectional and unidirectional PBR).
- Intra-VRF and inter-VRF vzAny-to-ESG contract (one/two-arm, multiple nodes service insertion with bidirectional and unidirectional PBR).
- Intra-VRF vzAny-to-vzAny contract (one-arm one-node service insertion with bidirectional PBR).
- Where PBR policy is applied differs based on contract configuration. Please see Table 1.

© 2026 Cisco and/or its affiliates. All rights reserved.

For more information

Cisco Application Centric Infrastructure White Papers:

https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-listing.html.

Cisco ACI Contract Guide White Paper:

https://www.cisco.com/c/en/us/products/collateral/networking/cloud-networking/application-centric-infrastructure/contract-guide.html

Cisco Application Centric Infrastructure Policy-Based Redirect Service Graph Design White Paper:

https://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739971.html.

ACI Multi-Pod White Paper:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-737855.html.

Cisco ACI Multi-Site Architecture White Paper:

https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/white-paper-c11-739609.html.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Printed in USA C11-6013751-01 11/26

© 2026 Cisco and/or its affiliates. All rights reserved.