



## **Cisco ACI Multi-Site Fundamentals Guide, Release 1.x**

**First Published:** 2017-08-10

**Last Modified:** 2018-11-20

### **Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
<http://www.cisco.com>  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883





## CONTENTS

---

### PREFACE

<b>Preface</b>	<b>v</b>
Audience	v
Documentation Conventions	v
Documentation Feedback	vi
Obtaining Documentation and Submitting a Service Request	vi

---

### CHAPTER 1

<b>New and Changed Information</b>	<b>1</b>
New and Changed Information	1

---

### CHAPTER 2

<b>About Cisco ACI Multi-Site</b>	<b>3</b>
About Cisco ACI Multi-Site	3
Terminology	4
User and Roles	5
Cisco ACI Multi-Site Schema and Templates	5

---

### CHAPTER 3

<b>Multi-Site Use Cases</b>	<b>9</b>
Cisco ACI Multi-Site Service Integration	9
Supported Use Cases	9
East-West Intra-VRF/Non-Shared Service	9
East-West Shared Service	11
North-South Intra-VRF/Non-Shared Service	13
Limitations	15
Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN	16
Limitations	17
Troubleshooting	17
Stretched Bridge Domain with Layer 2 Broadcast Extension	17

- Stretched Bridge Domain with No Layer 2 Broadcast Extension 19
- Stretched EPG Across Sites 21
- Stretched VRF with Inter-Site Contracts 23
- Shared Services with Stretched Provider EPG 25
- Migration of Cisco ACI Fabric to Cisco ACI Multi-Site 28
  - Setting up Cisco ACI Multi-Site with Multipod-Enabled Fabrics 29
  - Adding a Multipod Fabric as a Site on Cisco ACI Multi-Site 30
  - Converting a Single POD Site to Multipod 32



## Preface

---

This preface includes the following sections:

- [Audience, on page v](#)
- [Documentation Conventions, on page v](#)
- [Documentation Feedback, on page vi](#)
- [Obtaining Documentation and Submitting a Service Request, on page vi](#)

## Audience

This publication is for network administrators who install, configure, and maintain Cisco Nexus switches.

## Documentation Conventions

Command descriptions use the following conventions:

Convention	Description
<b>bold</b>	Bold text indicates the commands and keywords that you enter literally as shown.
<i>Italic</i>	Italic text indicates arguments for which the user supplies the values.
[x]	Square brackets enclose an optional element (keyword or argument).
[x   y]	Square brackets enclosing keywords or arguments separated by a vertical bar indicate an optional choice.
{x   y}	Braces enclosing keywords or arguments separated by a vertical bar indicate a required choice.
[x {y   z}]	Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.
variable	Indicates a variable for which you supply values, in context where italics cannot be used.

Convention	Description
string	A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.

Examples use the following conventions:

Convention	Description
<code>screen font</code>	Terminal sessions and information the switch displays are in screen font.
<b>boldface screen font</b>	Information you must enter is in boldface screen font.
<i>italic screen font</i>	Arguments for which you supply values are in italic screen font.
<>	Nonprinting characters, such as passwords, are in angle brackets.
[ ]	Default responses to system prompts are in square brackets.
!, #	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

## Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to . We appreciate your feedback.

## Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation* at: <http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html>

Subscribe to *What's New in Cisco Product Documentation*, which lists all new and revised Cisco technical documentation as an RSS feed and delivers content directly to your desktop using a reader application. The RSS feeds are a free service.



# CHAPTER 1

## New and Changed Information

This chapter contains the following sections:

- [New and Changed Information, on page 1](#)

## New and Changed Information

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

**Table 1: New Features and Changed Behavior in the Cisco ACI Multi-Site Fundamentals Guide**

Cisco ACI Multi-Site Version	Feature	Description	Where Documented
1.2(1)	Cisco ACI Multi-Site Service Integration	Added support for Layer 4 to Layer 7 service graphs between endpoint groups (EPGs) that are stretched across sites.	For more information, see <a href="#">Multi-Site Use Cases, on page 9</a> .
1.2(1)	Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN	Added support for direct connection between spines of 2 different sites without any IPN between the sites.	For more information, see <a href="#">Multi-Site Use Cases, on page 9</a> .
1.2(1)	Setting up Cisco ACI Multi-Site with Multipod Enabled Fabrics	Added support for setting up Cisco ACI Multi-Site with Multipod enabled fabrics.	For more information, see <a href="#">Multi-Site Use Cases, on page 9</a> .
1.0(2)	External EPGs	Added external EPGs, used to connect site local L3Outs, to enable communications between the EPGs in a stretched tenant and VRF.	For more information, see <a href="#">Multi-Site Use Cases, on page 9</a> .

<b>Cisco ACI Multi-Site Version</b>	<b>Feature</b>	<b>Description</b>	<b>Where Documented</b>
1.0(1)	--	This guide was released.	--



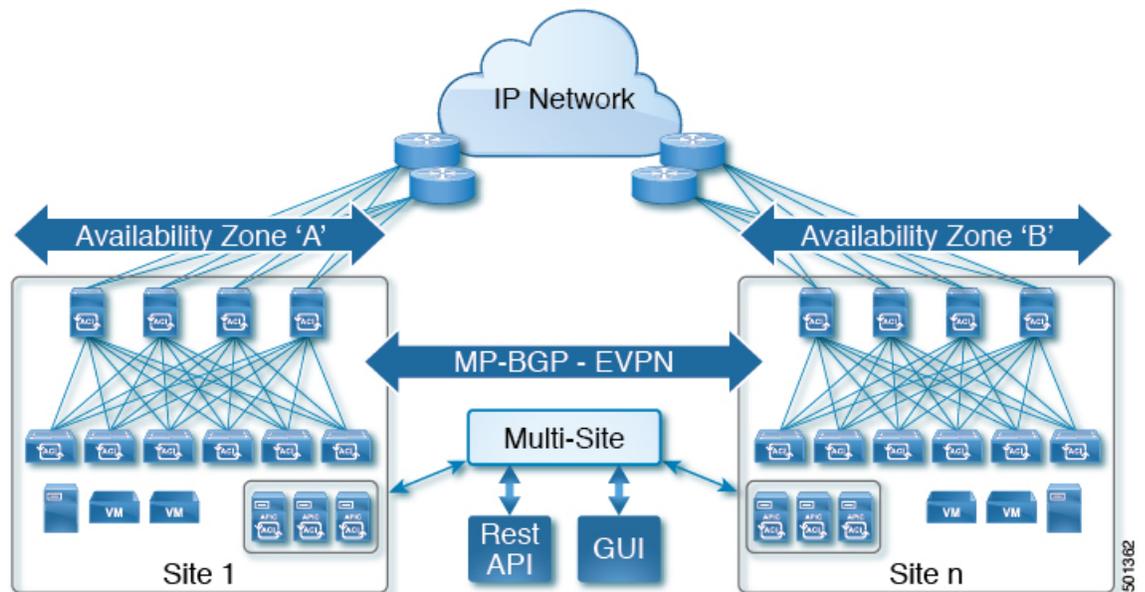
## CHAPTER 2

# About Cisco ACI Multi-Site

- [About Cisco ACI Multi-Site, on page 3](#)
- [Terminology, on page 4](#)
- [User and Roles, on page 5](#)
- [Cisco ACI Multi-Site Schema and Templates, on page 5](#)

## About Cisco ACI Multi-Site

Figure 1: Cisco ACI Multi-Site Architecture



As the newest advance on the Cisco ACI methods to interconnect networks, Cisco ACI Multi-Site is an architectural approach for interconnecting and managing multiple sites, each serving as a single fabric and availability zone. As shown in the diagram, the Multi-Site architecture has three main functional components:

- Two or more ACI fabrics built with Nexus 9000 switches deployed as leaf and spine nodes.
- One APIC cluster domain in each fabric.

- An inter-site policy manager, named Cisco ACI Multi-Site, which is used to manage the different fabrics and to define inter-site policies.

Multi-Site has the following benefits:

- Complementary with Cisco APIC, in Multi-Site each site is an availability zone (APIC cluster domain), which can be configured to be a shared or isolated change-control zone.
- MP-BGP EVPN is used as the control plane between sites, with data-plane VXLAN encapsulation across sites.
- The Multi-Site solution enables extending the policy domain end-to-end across fabrics. You can create policies in the Multi-Site GUI and push them to all sites or selected sites. Alternatively, you can import tenants and their policies from a single site and deploy them on other sites.
- Multi-Site enables a global view of site health.
- From the GUI of the Multi-Site Policy Manager, you can launch site APICs.
- Cross-site namespace normalization is performed by the connecting spine switches. This function requires Cisco Nexus 9000 Series switches with "EX" on the end of the name, or newer.
- Disaster recovery scenarios offering IP mobility across sites is one of the typical Multi-Site use cases.

For information about hardware requirements and compatibility, see *Cisco ACI Multi-Site Hardware Requirements Guide*.

For best practices for Multi-Site, see the *Deployment Best Practices* in [Cisco ACI Multi-Site Architecture White Paper](#).

For the Cisco ACI Multi-Site documentation set, see <http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html>.

## Terminology

As a complementary product with Cisco ACI, much of the Cisco ACI Multi-Site terminology is shared with ACI and APIC (for example, they both use the terms *fabric*, *tenant*, *contract*, *application profile*, *EPG*, *bridge domain*, and *L3Out*). For definitions of ACI terminology, see *Cisco Application Centric Infrastructure Fundamentals*.

### Micro-services architecture

In its first implementation, the Cisco ACI Multi-Site (inter-site policy manager) is represented by a cluster of three Virtual Machines (VMs) running on ESXi hosts. These ESXi hosts do not need to be connected to the ACI leaf nodes, because it is only required to establish IP connectivity between the VMs and the OOB IP addresses of the different APIC cluster nodes.

### Namespace

Each fabric maintains separate data in its name space, including such objects as the TEP pools, Class-IDs (EPG identifiers) and VNIDs (identifying the different Bridge Domains and the defined VRFs). The site-connecting spine switches (EX or later) perform the necessary namespace translation (normalization) between sites.

### Schema

Profile including the site-configuration objects that will be pushed to sites.

**Site**

APIC cluster domain or single fabric, treated as an ACI region and availability zone. It can be located in the same metro-area as other sites, or spaced world-wide.

**Stretched**

Objects (tenants, VRFs, EPGs, bridge-domains, subnets or contracts) are stretched when they are deployed to multiple sites.

**Template**

Child of a schema, a template contains configuration-objects that are shared between sites or site-specific.

**Template Conformity**

When templates are stretched across sites, their configuration details are shared and standardized across sites. To maintain template conformity, it is recommended to only make changes in the templates, using the Multi-Site GUI and not in a local site's APIC GUI.

## User and Roles

The Cisco ACI Multi-Site provides access according to a user's role through role-based access control (RBAC). Roles are used for both local and external authentication. The following user roles are available in Cisco ACI Multi-Site.

- Power User—A power user can perform all the operations as an *admin* user.
- Site and Tenant Manager—A site and tenant manager can manage sites, tenants, and associations.
- Schema Manager—A schema manager can manage all schemas regardless of tenant associations.
- Schema Manager - Restricted —A restricted schema manager can manage schemas that contain at least one tenant to which the user is explicitly associated.
- User and Role Manager—A user and role manager can manage all the users, their roles, and passwords.

**Admin User**

In the initial configuration script, the admin account is configured and the *admin* is the only user when the system starts. The initial password for the *admin* user is set by the system. You must change the *admin* password during the first log in.

- The *admin* user is assigned the role of a Power User.
- Use the *admin* user to creating other users and perform all other Day-0 configurations.
- The account status of the *admin* user cannot be set to **Inactive**.

## Cisco ACI Multi-Site Schema and Templates

**Cisco ACI Object Model**

At the top level, the Cisco ACI object model is built on a group of one or more tenants, allowing the network infrastructure administration and data flows to be segregated.

## Policy Types

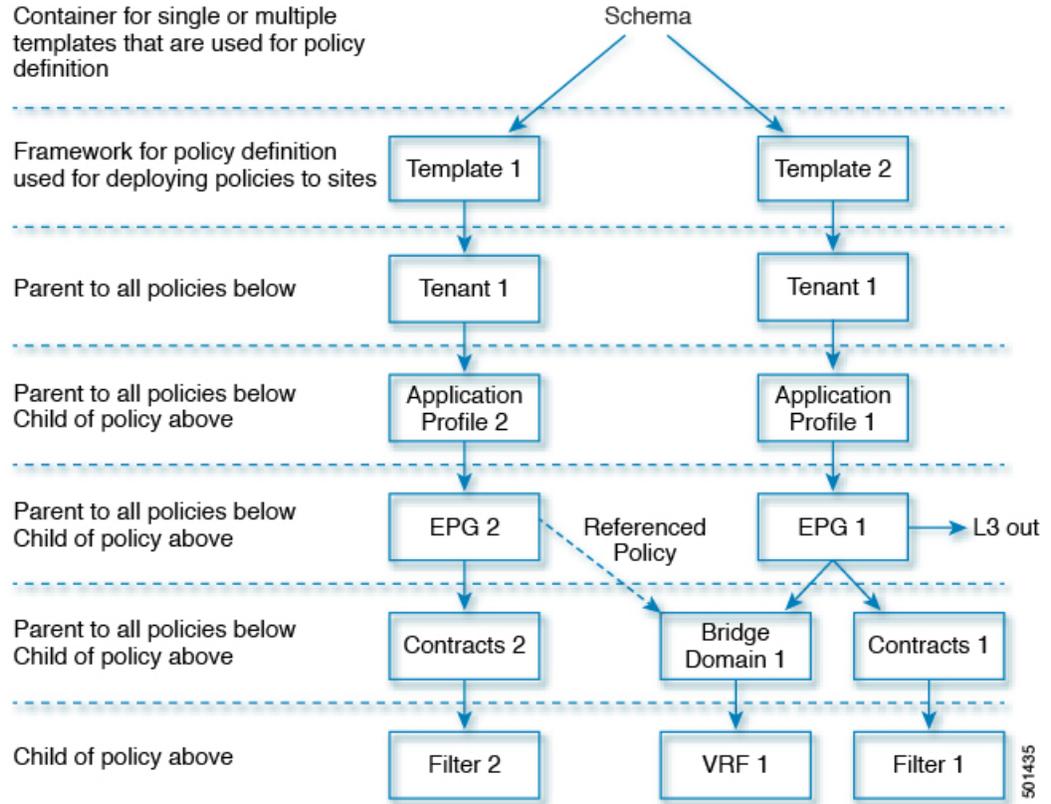
See the following section on the terminology and conceptual information on different policy types:

- **Schemas:** Schemas are the containers for single or multiple templates that are used for defining the policies. Templates are the framework for defining and deploying the policies to the sites.
- **Tenants:** A tenant is a logical container for application policies that enable an administrator to exercise domain-based access control. A tenant represents a unit of isolation from a policy perspective, but it does not represent a private network. Tenants can represent a customer in a service provider setting, an organization or domain in an enterprise setting, or just a convenient grouping of policies.  
  
Tenant is the parent policy to all the policies, for example, Application Profiles, EPG, Contract, Bridge Domains, VRFs, and Filters.
- **Application Profile:** The application profile is a set of requirements that an application instance has on the virtualizable fabric. The policy regulates connectivity and visibility among endpoints within the scope of the policy.
- **EPG:** An EPG is a managed object that is a named logical entity that contains a collection of endpoints. Endpoints are devices that are connected to the network directly or indirectly. They have an address (identity), a location, attributes (such as version or patch level), and can be physical or virtual. Knowing the address of an endpoint also enables access to all its other identity details. EPGs are fully decoupled from the physical and logical topology. Endpoint examples include servers, virtual machines, network-attached storage, or clients on the Internet. Endpoint membership in an EPG can be dynamic or static.
- **Contracts:** Contracts define inbound and outbound permit, deny, and QoS rules and policies such as redirect. Contracts allow both simple and complex definition of the way that an EPG communicates with other EPGs, depending on the requirements of the environment. Although contracts are enforced between EPGs, they are connected to EPGs using provider-consumer relationships. Essentially, one EPG provides a contract, and other EPGs consume that contract.
- **Bridge Domains:** A bridge domain (fvBD) represents a Layer 2 forwarding construct within the fabric. The following figure shows the location of bridge domains in the management information tree (MIT) and their relation to other objects in the tenant.
- **Virtual Routing and Forwarding (VRF):** A Virtual Routing and Forwarding (VRF) object (fvCtx) or context is a tenant network (called a private network in the APIC GUI). A tenant can have multiple VRFs. A VRF is a unique Layer 3 forwarding and application policy domain. The following figure shows the location of VRFs in the management information tree (MIT) and their relation to other objects in the tenant.
- **Filters:** Filters are specific rules for the policy between two EPGs. Filters consist of inbound and outbound rules: permit, deny, redirect, log, copy, and mark.

## Model of Schemas and Templates

See the following illustration for simplifying the object model of Schemas and Templates:

Figure 2: Framework for Cisco ACI Multi-Site Schema and Templates



See the relation between different policy types:

- Application Profiles is the parent policy for EPGs.
- EPG is the parent policy for Contracts and Bridge Domains.
- Contracts is the parent policy for Filters.
- Bridge Domains is the parent policy for VRFs.





## CHAPTER 3

# Multi-Site Use Cases

---

- [Cisco ACI Multi-Site Service Integration, on page 9](#)
- [Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN, on page 16](#)
- [Stretched Bridge Domain with Layer 2 Broadcast Extension, on page 17](#)
- [Stretched Bridge Domain with No Layer 2 Broadcast Extension, on page 19](#)
- [Stretched EPG Across Sites, on page 21](#)
- [Stretched VRF with Inter-Site Contracts, on page 23](#)
- [Shared Services with Stretched Provider EPG, on page 25](#)
- [Migration of Cisco ACI Fabric to Cisco ACI Multi-Site, on page 28](#)

## Cisco ACI Multi-Site Service Integration

Starting in Release 1.2(1), there are many variants of service graph deployment in single ACI fabric. For Cisco ACI Multi-Site deployment, not all use cases are supported in current release. Supported uses are mentioned below in the document.

To support the use cases mentioned below, following topology is required for service nodes for all the use cases.

- Each site has individual Active/Standby Service Node pair.
- Layer 4 to Layer 7 devices are in unmanaged mode.
- Policy based redirect (PBR) policies are required to redirect traffic to service nodes.

## Supported Use Cases

### East-West Intra-VRF/Non-Shared Service

This is the use case for east-west communication between endpoints in the same VRF instance across sites. EPG WEB and EPG APP are deployed across sites and the firewall (FW) needs to be inserted between them. This is a common design for traffic within the application.

Requirements:

- Policy-based redirect (PBR) policies are required on the consumer and provider connectors of the service node.

- Provider EPGs should have the subnet defined under it and should be unique (similar to intra-VRF route leaking). This is to ensure that the first packet does not reach to the provider's leaf switch without the policy being applied, as you want the traffic to select the service node in the consumer site always. The policies are always applied on the consumer's leaf switch. For this purpose, provider's subnet would be leaked to consumer's leaf switch.
- Redirect policies are always applied on the consumer leaf switch, which means the FW device cluster chosen will always be local to the site of the consumer for traffic flows.

The figures below shows the example for traffic from EPG WEB to EPG APP and vice versa.

Figure 3: Across Site Incoming Traffic

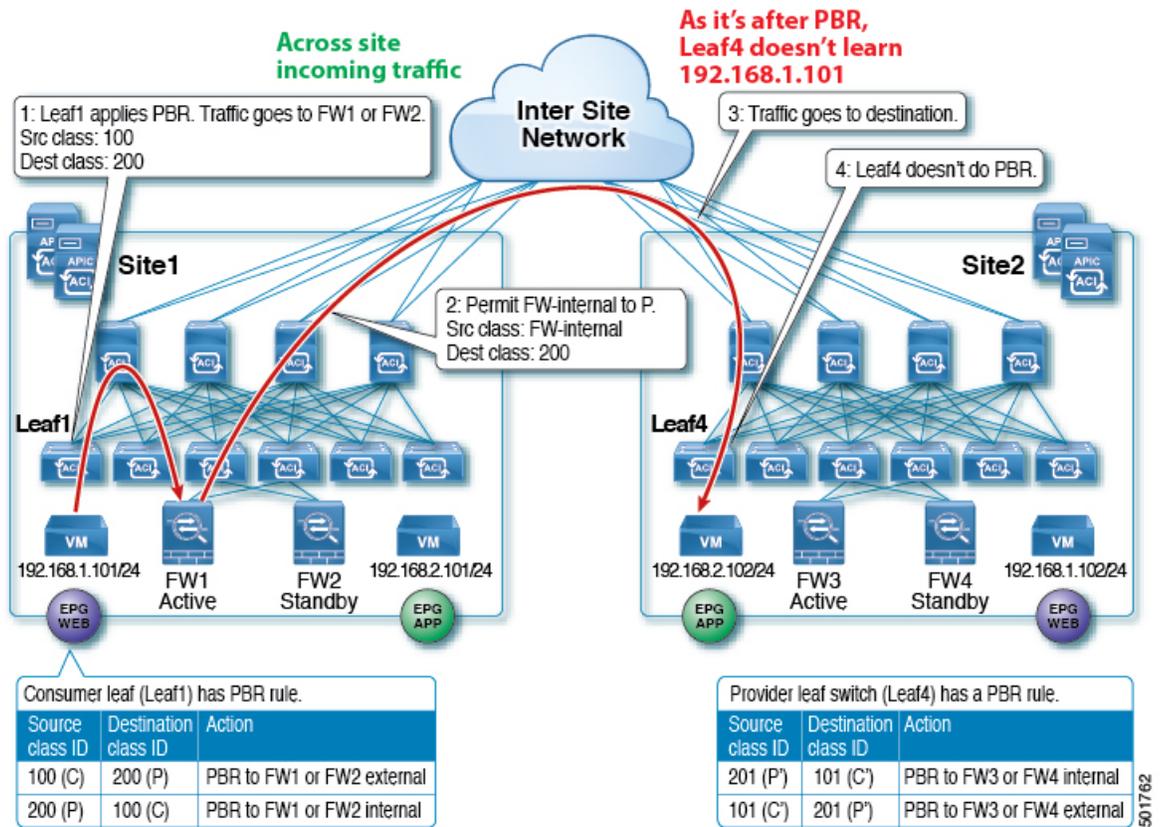
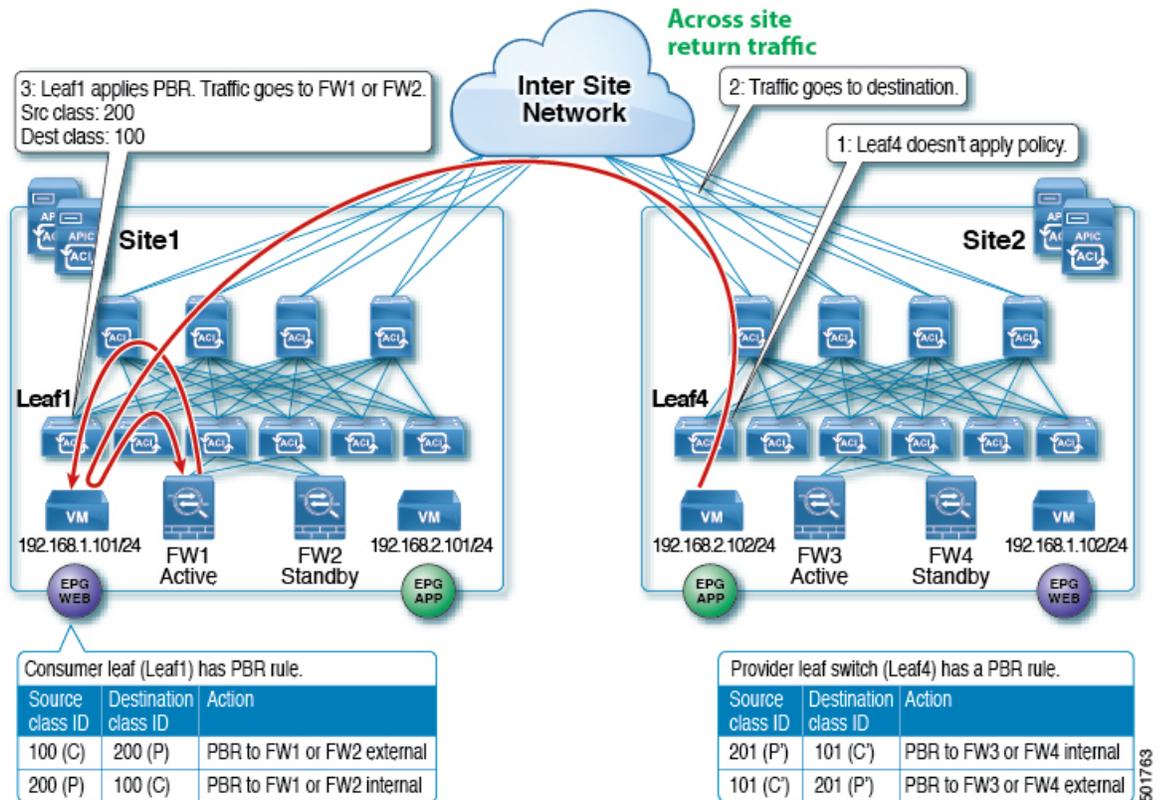


Figure 4: Across Site Return Traffic



501763

## East-West Shared Service

This is the use case for east-west communication between endpoints in the different VRFs across sites. The figure shows the example, EPG WEB in VRF1 and EPG APP in VRF2 are deployed across sites and firewall needs to be inserted between them. This is a common design for the shared service use case. For example, the NFS service is shared to multiple servers in different VRFs.

Requirements:

- PBR policies are required on the consumer and provider connectors of the service node.
- Provider EPGs should have the subnet defined under it and should be unique (similar to intra-VRF route leaking). This is to ensure that first packet does not reach to the provider's leaf switch without the policy being applied, as you want the traffic to select the service node in the consumer site always. The policies are always applied on the consumer's leaf switch. For this purpose, provider's subnet would be leaked to consumer's leaf.
- Redirect policies are always applied on the consumer leaf, which means the firewall (FW) device cluster chosen will always be local to the site of the consumer for traffic flows.

Figure 5: Across Site Incoming Traffic

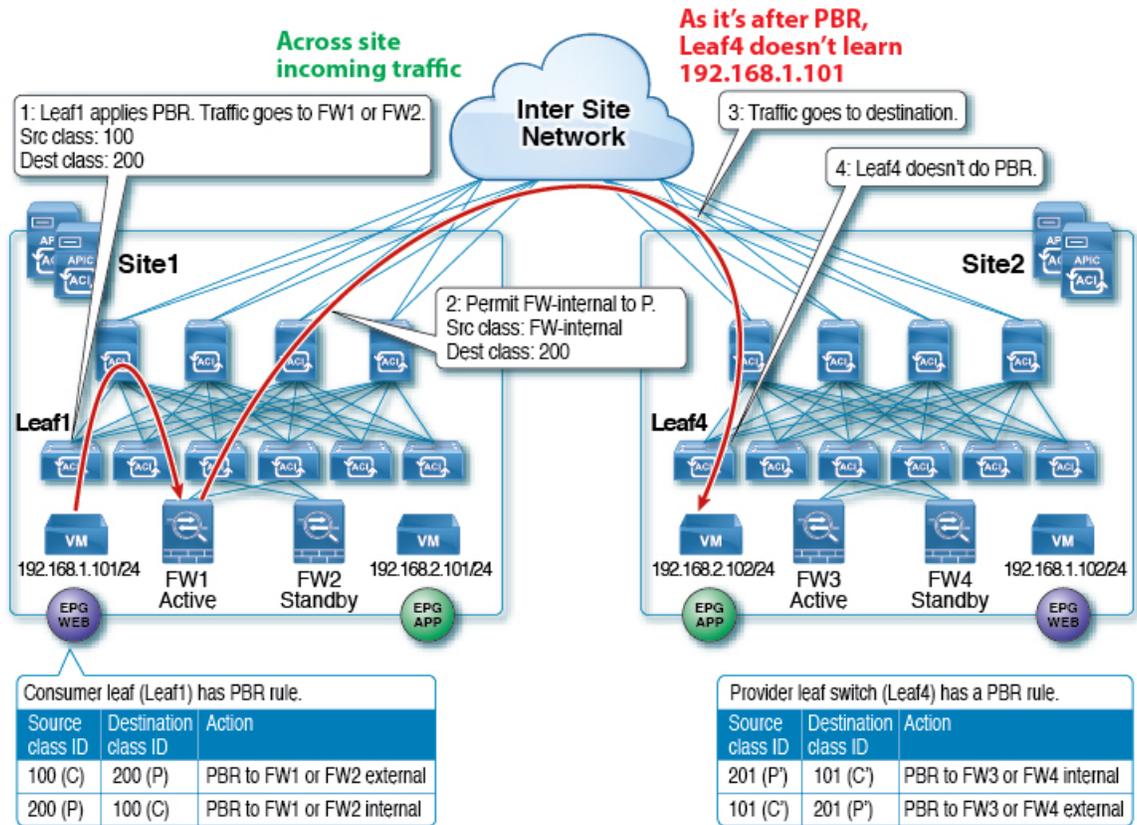
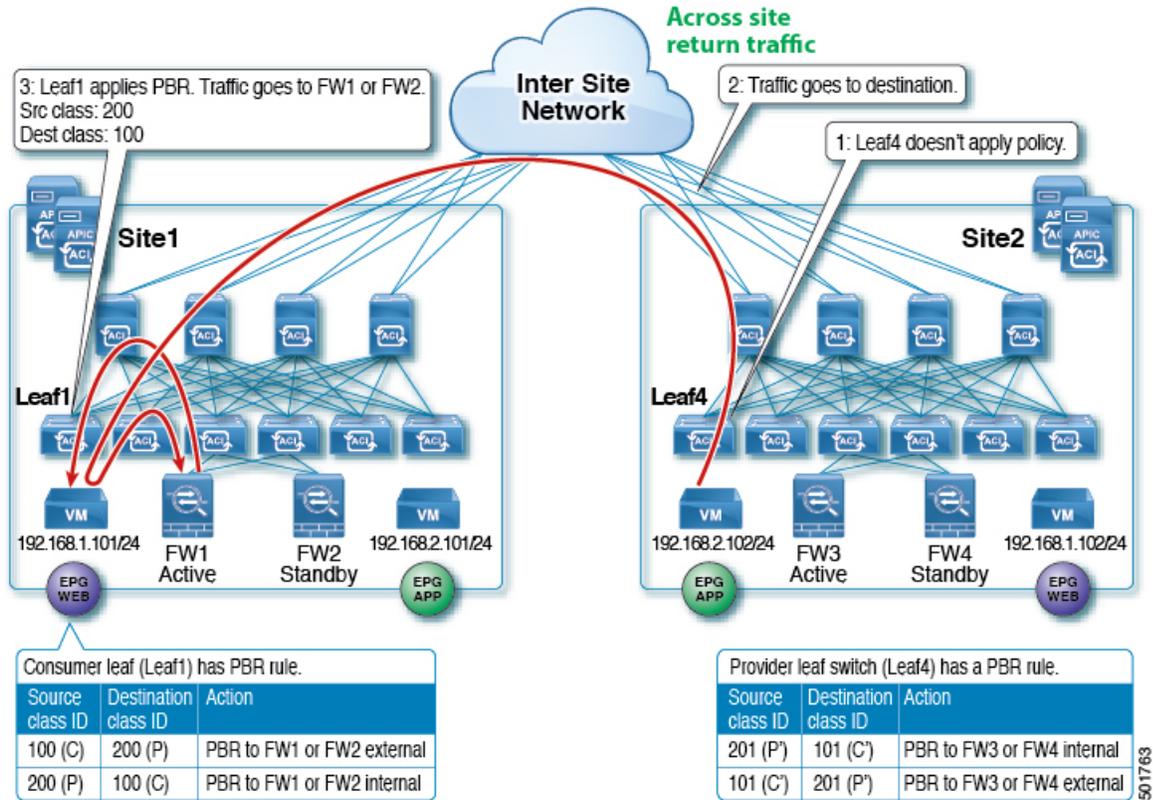


Figure 6: Across Site Return Traffic



501763

### North-South Intra-VRF/Non-Shared Service

This is the use case for north-south communication between endpoints in the datacenter and outside. Each site has L3out accessing to EPG WEB that is deployed across sites and the firewall needs to be inserted between them. This is a common design for the web front-end that is accessible from client outside of the data center.

Requirements:

- Both L3Out and EPG-WEB are same VRF.
- VRF needs to be in ingress mode (which is by default). In egress mode does not work, because of asymmetric traffic flow.
- PBR policies are needed on both connectors of the service graph.

Figure 7: Intra-VRF, ingress enforcement L3out is the consumer - IncomingTraffic

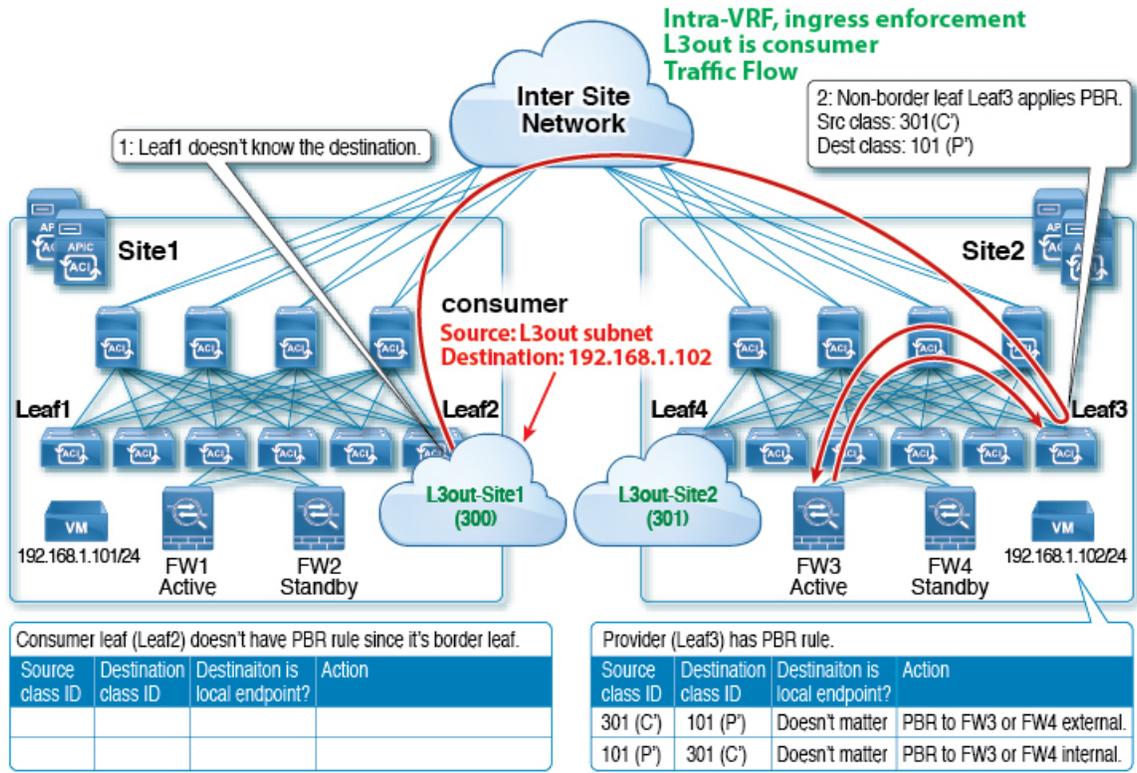
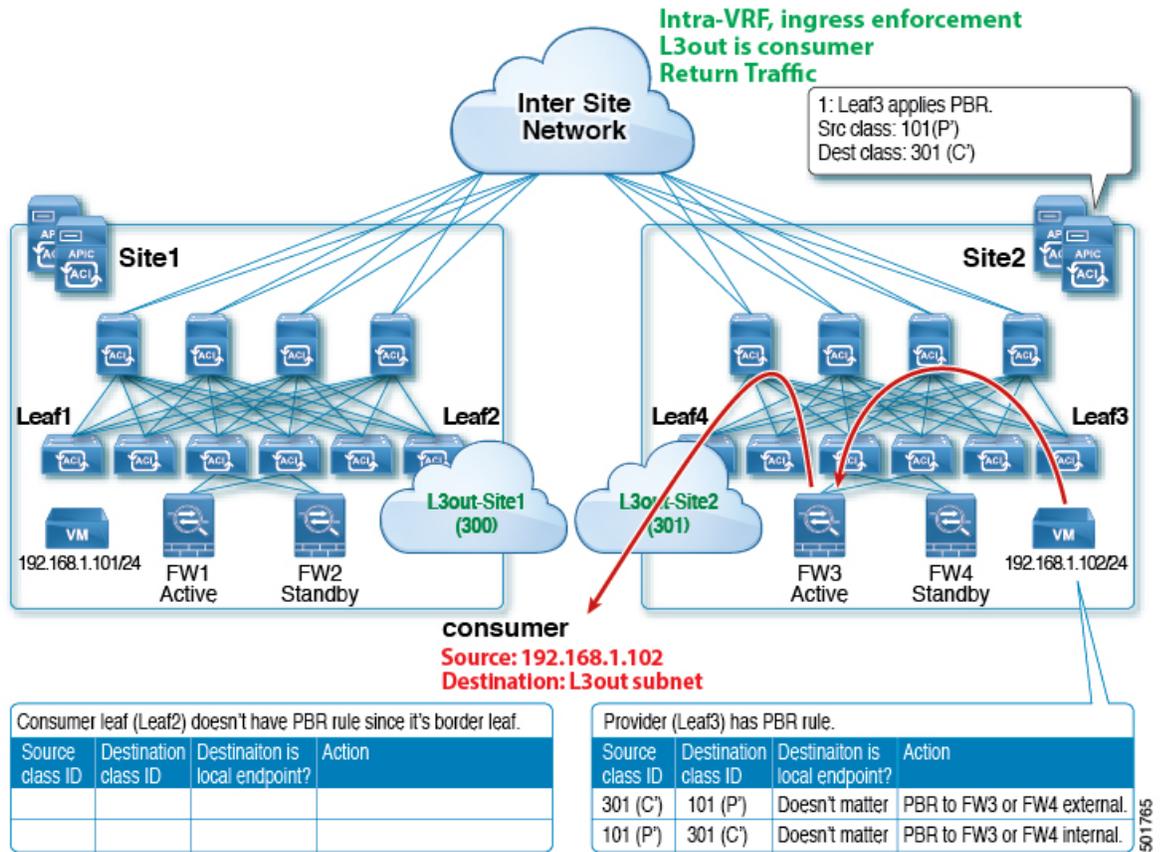


Figure 8: Intra-VRF, ingress enforcement L3out is the consumer - Return Traffic



## Limitations

In Multi-Site deployments, when the contract is deployed with a service graph, Scope of the contract should be defined in such a way so that all providers are in the same scope. This creates single graph instance for all the providers.

In Release 1.2(1), Multi-Site does not support handling multiple graph instances created due to provider being in a scope that is bigger than contract scope. While deciding on the contract scope note the following:

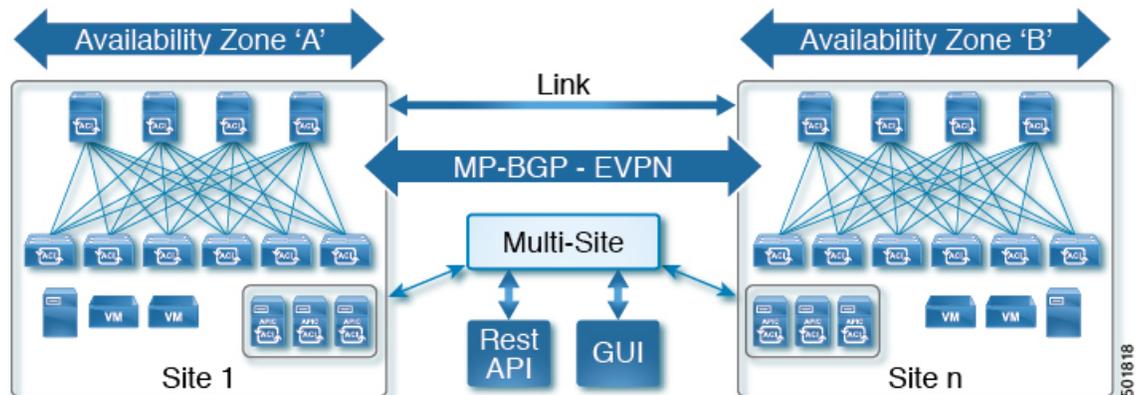
- If all the providers are in a single VRF, contract scope can be set to any available option.
- If providers are in multiple VRFs, but in a single Tenant, scope should be set to “Tenant”.
- If providers are in multiple tenants, this scenario is not supported with a single contract.

To overcome this limitation multiple contracts can be created with the same service graph attached.

# Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN

This Cisco ACI Multi-Site use case provides support for direct connection between spines of 2 different sites without any IPN between the sites.

**Figure 9: Multi-Site Back to Back Spine – Basic Setup**



This use case enables:

- Support for direct connection between spines of 2 different sites without any IPN between the sites
- Support for only a single POD per site deployments
- Requires unique fabric names across sites

Design:

- LLDP will detect spine to spine connection and will create a wiring issue on that port
- DHCP relay will not be configured on the link
- When the LLDP detects unique fabric names and when the spines on both sides are discovered, the port will be put back in-service except for the following:
  - ISIS will not be enabled on the link
  - Infra VLAN will not be learned from the link
  - LLDP TLV will be between the sites and will be ignored
- Spine-to-spine link will be treated as external subinterface
- The configuration and data path will be same as a regular Multi-Site set up

## Limitations

- With back-to-back connectivity, we recommend that you deploy multiple spines in each site to provide inter-site connectivity. From each of these spines, provide multiple links to each of the spines in each of the other sites.
- In the hybrid case where IPN is also used, we recommend that all sites have to be connected to the IPN in a fault-tolerant fashion to avoid transit situation.
- Only two sites are supported with back-to-back spine.
- No new configuration required in APIC for this use case.

## Troubleshooting

This section describes troubleshooting techniques.

- In APIC, check if l3extOut is configured for this interface in both sites.
- If there is no reachability between the two site spines, perform the following:

- Make sure there are no wiring issues, the port is up and switchingSt is enabled:

```
dev-infral-spine1# cat /mit/sys/lldp/inst/if-[eth1--1]/summary | grep wiringIssues
wiringIssues :
dev-infral-spine1#
```

- Make sure the IP address is assigned from the l3extOut configuration and OSPF session is up:

```
IP Interface Status for VRF "overlay-1"
eth1/53.7, Interface status: protocol-up/link-up/admin-up, iod: 63, mode: external
```

- Check the svc\_ifc\_policyelem.log\* file in the SPINE that is connected to the other site:

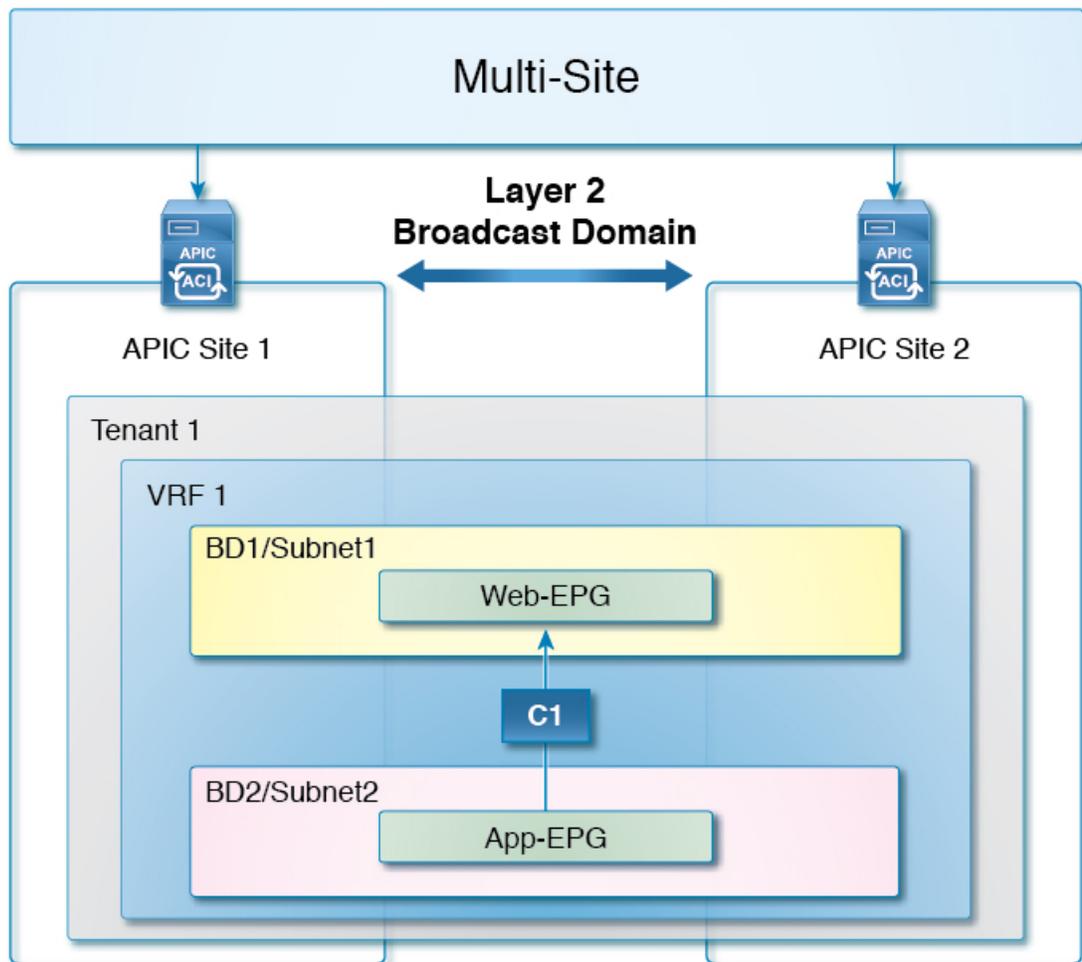
```
$ zgrep "back to back spine ignore wiring check." \
/var/sysmgr/tmp_logs/dme_logs/svc_ifc_policyelem.log*
```

## Stretched Bridge Domain with Layer 2 Broadcast Extension

This is the most basic Cisco ACI Multi-Site use case, in which a tenant and VRF are stretched between sites. The EPGs in the VRF (with their bridge domains (BDs) and subnets), as well as their provider and consumer contracts are also stretched between sites.

In this use case, Layer 2 broadcast flooding is enabled across fabrics. Unknown unicast traffic is forwarded across sites leveraging the Head-End Replication (HER) capabilities of the spine nodes that replicate and send the frames to each remote fabric where the Layer 2 BD has been stretched.

Figure 10: Stretched Bridge Domain with Layer 2 Broadcast Extension



This use case enables:

- Same application hierarchy deployed on all sites with common policies. This allows seamlessly deploying workloads belonging to the various EPGs across different fabrics and governing their communication with common and consistent policies.
- Layer 2 clustering
- Live VM migration
- Active/Active high availability between the sites
- Using Service Graphs to push shared applications between sites is not supported.

#### Prerequisites for this Use Case

- Sites have been added, APIC controllers are active, and communications are established.
- The tenant to be stretched has been created.
- The Multi-Site Site and Tenant Manager account is available

Single profile including the objects in the following table, pushed to multiple sites:

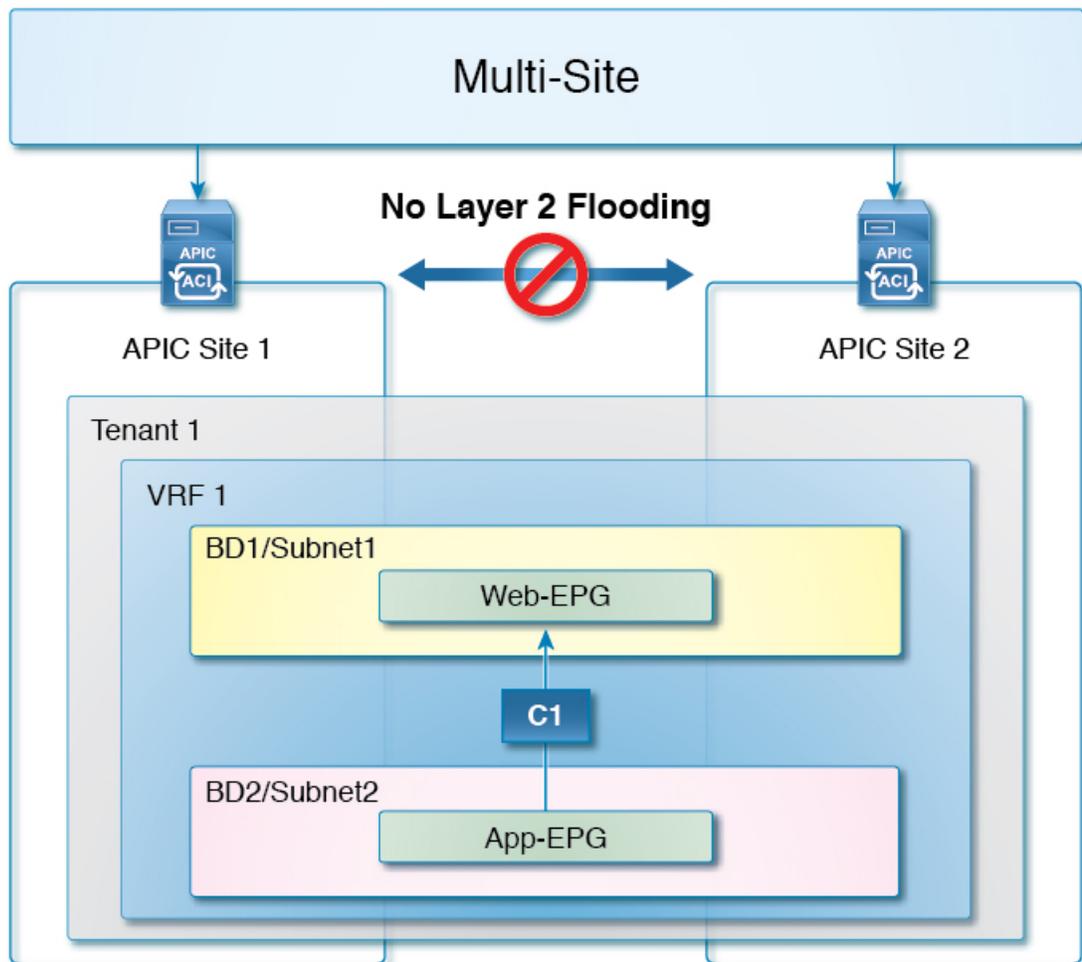
**Table 2: Features to be Configured for this Use Case**

Configuration	Description	Stretched or Local
Tenant	Imported from APIC or created in Multi-Site	Stretched
Site L3Outs	Configured in the APIC GUI and linked in the stretched tenant and VRF, site-specific templates	Local
VRF	VRF for the tenant	Stretched
Bridge Domain	Layer 2 stretching enabled Layer 2 flooding enabled Subnets to be shared added	Stretched
EPGs	EPGs in the BD	Stretched
Contracts	Include the filters needed to govern EPG communication	Stretched
External EPGs	Network Mappings of Site L3Outs (Cisco ACI Multi-Site, Release 1.0(1)) Site Connections of Site L3Outs through External EPGs (Multi-Site, Release 1.0(2))	Local, but linked to other sites

## Stretched Bridge Domain with No Layer 2 Broadcast Extension

This Cisco ACI Multi-Site use case is similar to the first use case where a tenant, VRF, and their EPGs (with their bridge domains and subnets) are stretched between sites.

Figure 11: Stretched Bridge Domain with No Layer 2 Broadcast Extension



However, in this use case, Layer 2 broadcast flooding is localized at each site. Layer 2 broadcast, multicast and unknown unicast traffic is not forwarded across sites over replicated VXLAN tunnels.

This use case enables:

- Control plane overhead is reduced by keeping Layer 2 flooding local
- Inter-site IP mobility for disaster recovery
- "Cold" VM Migration
- Using Service Graphs to push shared applications between sites is not supported.

#### Prerequisites for this Use Case

- Sites have been added, APIC controllers are active, and communications are established.
- The tenant to be stretched has been created.
- The Multi-Site Site and Tenant Manager account is available

Profile with the objects in the following table, pushed to multiple sites:

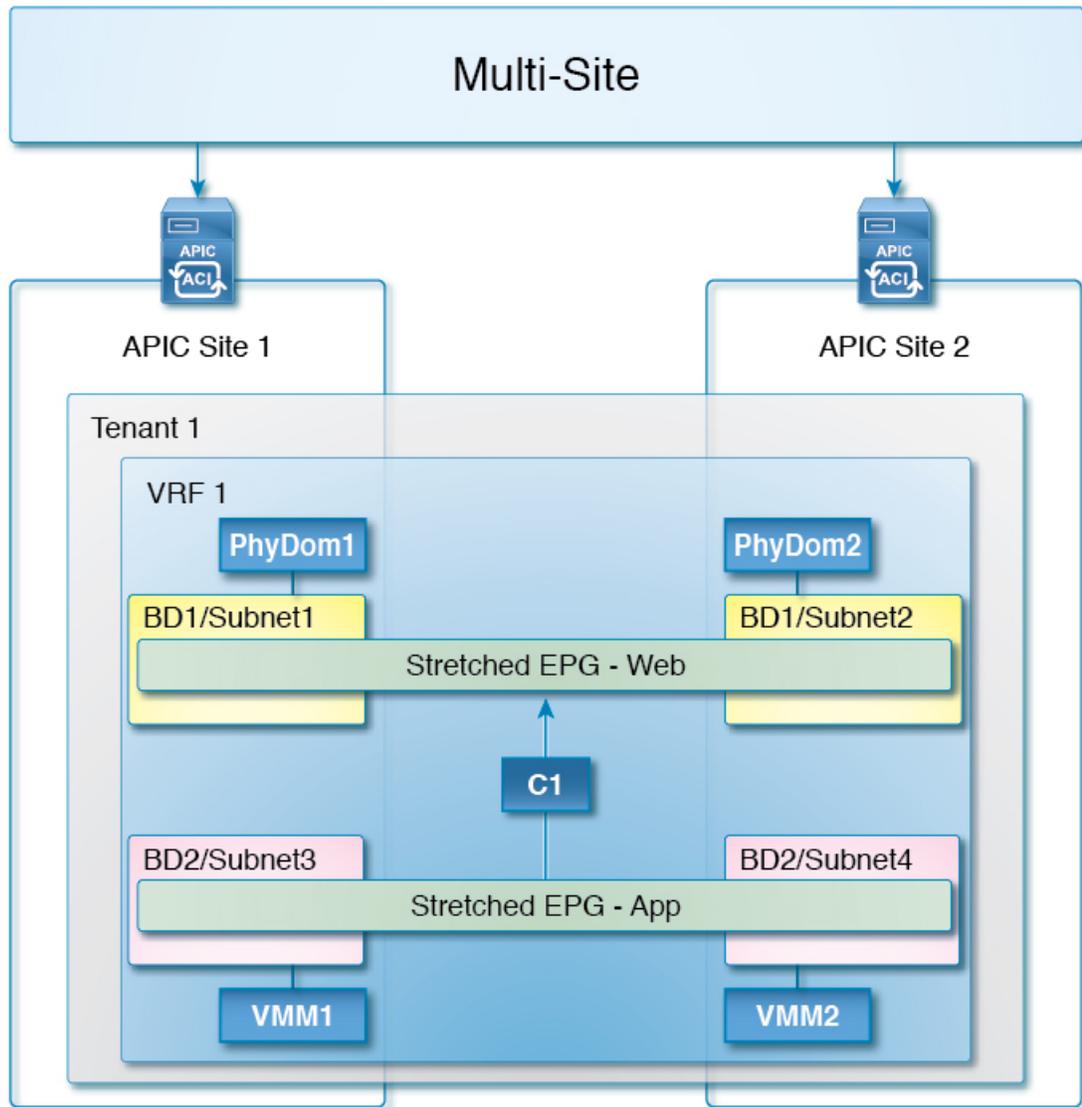
**Table 3: Features to Be Configured for this Use Case**

Configuration	Description	Stretched or Local
Tenant and VRF	Imported from APIC or created in Multi-Site	Stretched
Site L3Outs	Configured in the APIC GUI and linked in the stretched tenant and VRF, site-specific templates	Local
Bridge Domain	Layer 2 stretching enabled Layer 2 flooding disabled Subnets to be shared added	Stretched
EPGs	All EPGs in the BD	Stretched
Contracts	Include whatever filters and contracts are needed to govern EPG communication	Stretched
External EPGs	Network Mappings of Site L3Outs (Cisco ACI Multi-Site, Release 1.0(1)) Site Connections of Site L3Outs through External EPGs (Multi-Site, Release 1.0(2))	Local, but linked to other sites

## Stretched EPG Across Sites

This Cisco ACI Multi-Site use case provides endpoint groups (EPGs) stretched across multiple sites. Stretched EPG is defined as an endpoint group that expands across multiple sites where the underlying networking, site local, and bridge domain can be distinct.

Figure 12: Stretched EPG Across Sites



This use case enables Layer 3 forwarding to be used among all sites.

#### Prerequisites for this Use Case

- Sites have been added, APIC controllers are active, and communications are established.
- The relevant tenants have been created.
- The Multi-Site Site and Tenant Manager account is available
- A physical domain and VMM domain must exist on APIC.

Profiles pushed to single or multiple sites, including the objects in this table:

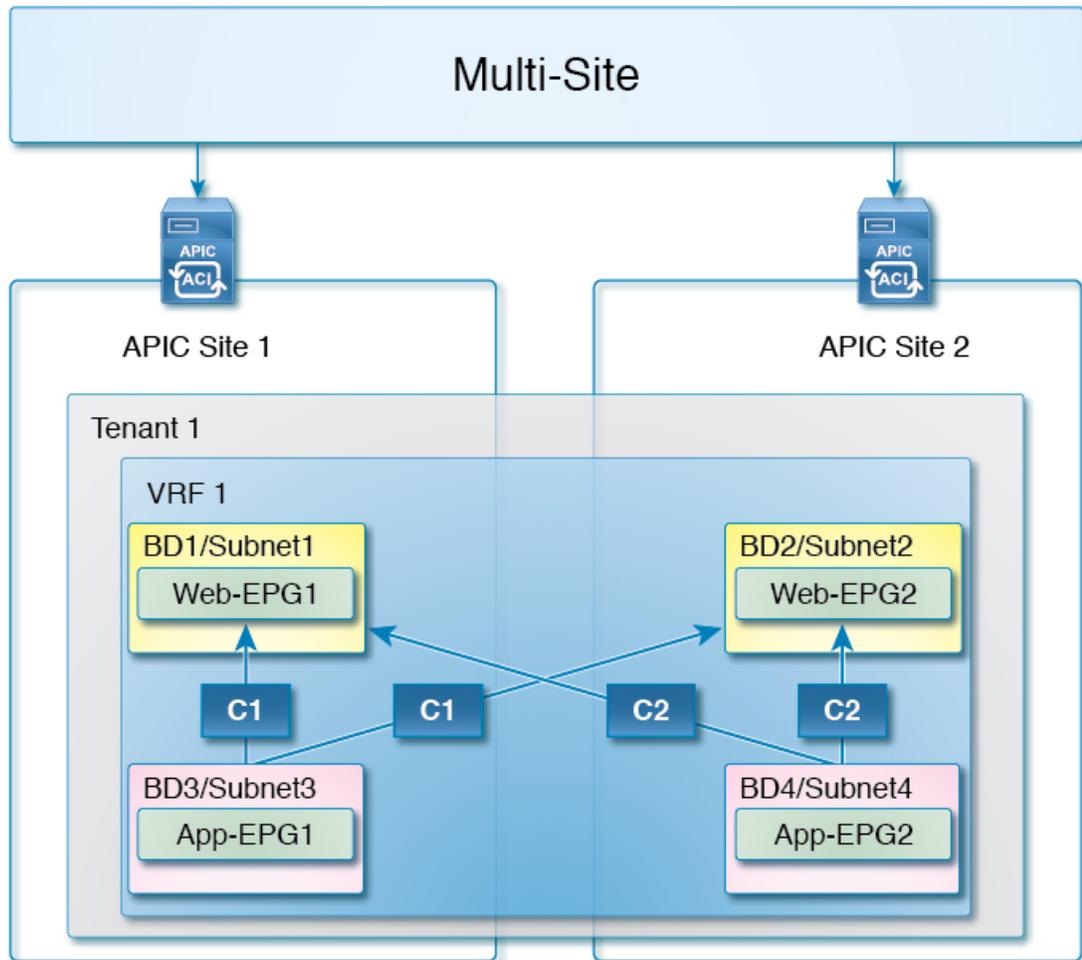
**Table 4: Features to be Configured for this Use Case**

Configuration	Description	Stretched or Local
Tenant, VRF and EPGs	Imported from APIC or created in Multi-Site.	Stretched
Site L3Outs	Configured in the APIC GUI and linked in the stretched tenant and VRF, site-specific templates	Local
Bridge Domains (DBs)	Layer 2 stretching disabled.	Stretched
Subnets	Unique for each BD on the local site.	Local
Contract	Contracts configured on site where they are provided	Local
External EPGs	Network Mappings of Site L3Outs (Cisco ACI Multi-Site, Release 1.0(1)) Site Connections of Site L3Outs through External EPGs (Multi-Site, Release 1.0(2))	Local, but linked to other sites

## Stretched VRF with Inter-Site Contracts

This Multi-Site use case provides inter-site communication between endpoints connected to different Bridge Domains (BDs) that are part of the same stretched VRF. VRF Stretching is a convenient way to manage EPGs across sites (and the contracts between them).

Figure 13: VRF Stretching with Inter-site Contracts



In the diagram above, the App-EPGs provide the C1 and C2 contracts across the sites, and the Web-EPGs consume them across the sites.

This use case has the following benefits:

- The tenant and VRF are stretched across sites, but EPGs and their policies (including subnets) are locally defined.
- Because the VRF is stretched between sites, contracts govern cross-site communication between the EPGs. Contracts can be consistently provided and consumed within a site or across sites.
- Traffic is routed within and between sites (with local subnets) and static routing between sites is supported.
- Separate profiles are used to define and push local and stretched objects.
- No Layer 2 stretching and local Layer 2 Broadcast domains.
- “Cold” VM migration, without the capability of preserving the IP address of the migrated endpoints.
- Using Service Graphs to push shared applications between sites are not supported.

**Prerequisites for this Use Case**

- Sites have been added, APIC controllers are active, and communications are established.
- The tenants to be stretched have been created.
- The Multi-Site Site and Tenant Manager account is available.

Profiles pushed to single or multiple sites, including the objects in this table:

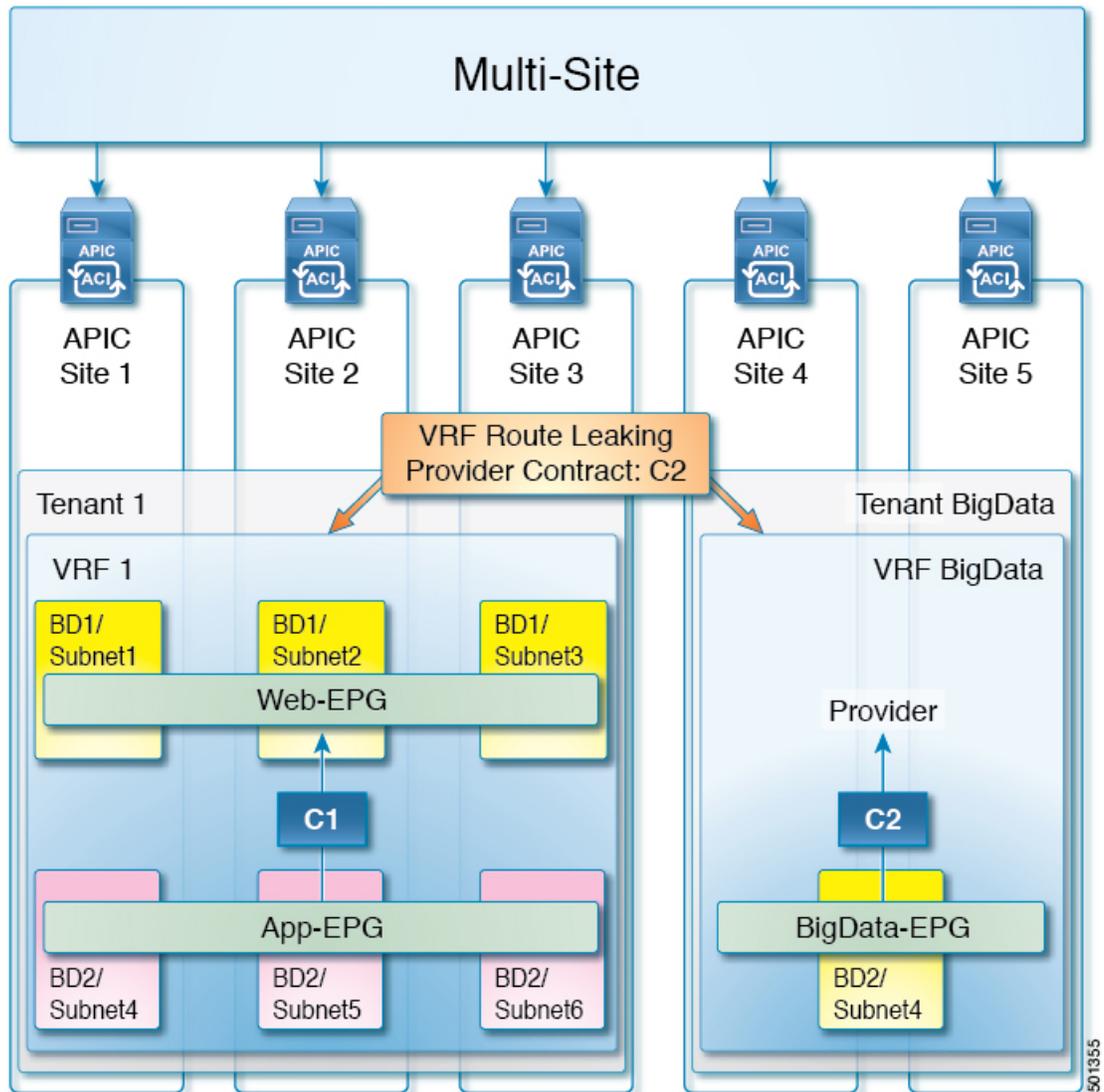
**Table 5: Features to be Configured for this Use Case**

Configuration	Description	Stretched or Local
Tenant and VRF	Imported from APIC or created in Multi-Site	Stretched
Site L3Outs	Configured in the APIC GUI and linked in the stretched tenant and VRF, site-specific templates	Local
EPGs providing contracts	EPGs for each site that provides services.	Local
EPGs consuming contracts	EPGs that consume the provided contracts, may be in the same site or multiple sites	Local
Bridge Domains for each EPG	Layer 2 stretching disabled Layer 2 flooding disabled	Local
Contracts	Contracts configured on site where they are provided	Local, but shared
External EPGs	Network Mappings of Site L3Outs (Cisco ACI Multi-Site, Release 1.0(1)) Site Connections of Site L3Outs through External EPGs (Multi-Site, Release 1.0(2))	Local, but linked to other sites

## Shared Services with Stretched Provider EPG

In this use case, the Provider EPGs in one group of sites offer shared services and the EPGs in another group of sites consume the services. All sites have local EPGs and bridge domains.

Figure 14: Shared Services with Stretched Provider EPG



In the diagram above, Site 4 and Site 5 (with BigData-EPG, in Tenant BigData/VRF BigData), provides shared data services, and the EPGs in Site 1 to Site 3, in Tenant 1/VRF 1, consume the services.

In the Shared Services usecase of Multi-Site, at the VRF boundary routes are leaked between VRFs for routing connectivity and by importing contracts across sites.

This use case has the following benefits:

- Shared services enable communications across VRFs and tenants while preserving the isolation and security policies of the tenants.
- A shared service is supported only with non-overlapping and non-duplicate subnets.
- Each group of sites has a different tenant, VRF, and one or more EPGs stretched across it.
- Site groups can be configured to use Layer 2 Broadcast extensions or to localize Layer 2 flooding.

- Stretched EPGs share the same bridge domain, but the EPGs have subnets that are configured under the EPG, not under the bridge domain.
- The provider contract must be set to global scope.
- VRF route leaking enables communication across the VRFs.
- Using Service Graphs to push shared applications between sites is not supported.

#### Prerequisites for this Use Case

- Sites have been added, APIC controllers are active, and communications are established.
- The relevant tenants have been created.
- The Multi-Site Site and Tenant Manager account is available

Schemas, with templates, pushed to groups of sites, including the objects in this table:

**Table 6: Features to be Configured for this Use Case**

Configuration	Description	Stretched or Local
Shared service provider schema, with multiple templates	Shared template, includes the following objects: <ul style="list-style-type: none"> <li>• Tenant</li> <li>• VRF</li> <li>• Provider Contract with global scope.</li> <li>• EPG with subnet set to <b>Advertised Externally and Shared Between VRFs.</b></li> </ul> Site-Specific templates, including bridge domains (optionally set for Layer 2 extension) and external EPGs	Stretched (pushed to all sites in the provider group)

Configuration	Description	Stretched or Local
Shared service consumer schema with multiple templates	<p>Shared template, includes the following objects:</p> <ul style="list-style-type: none"> <li>• Tenant</li> <li>• VRF</li> <li>• EPG with subnet set to <b>Advertised Externally and Shared Between VRFs</b>.</li> </ul> <p><b>Note</b> For the consumer EPGs, the subnets can alternatively be added in the BDs.</p> <ul style="list-style-type: none"> <li>• Consumer Contract (same name as the provided contract).</li> </ul> <p>Site-Specific templates, including bridge domains (optionally set for Layer 2 extension) and external EPGs</p>	Stretched or local
VRF route leaking	Contracts must be configured to enable VRF route leaking.	Configured cross-site

## Migration of Cisco ACI Fabric to Cisco ACI Multi-Site

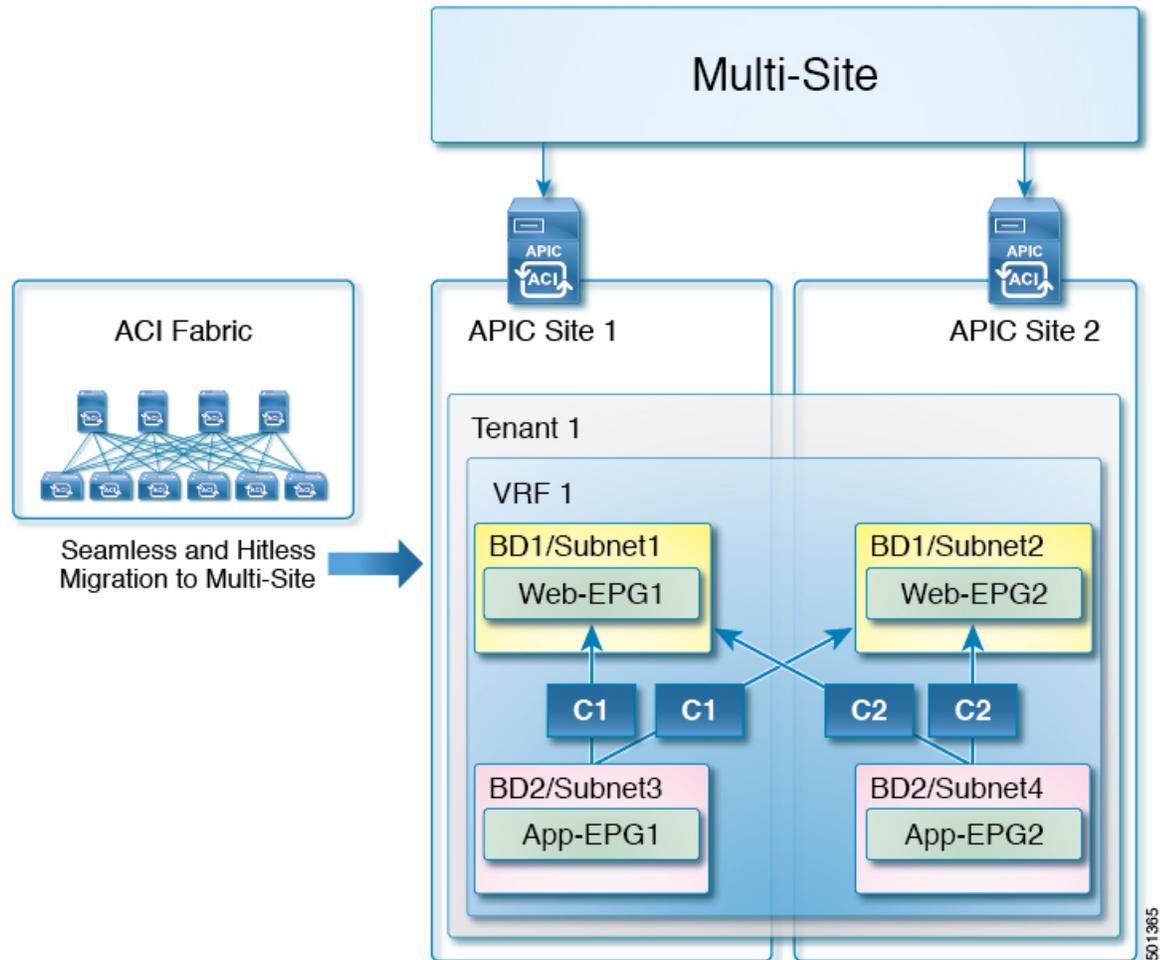
This is a common Cisco ACI Multi-Site use case, in which a tenant is migrated or imported from Cisco ACI fabric to Cisco ACI Multi-Site.

This use case is targeted for Brownfield to Greenfield and Greenfield to Greenfield types of deployments. The Brownfield to Brownfield use case is only supported in this release if both Cisco APIC sites are deployed with the same configuration. Other Brownfield to Brownfield use cases will be deployed in a future Cisco ACI Multi-Site release.

For Brownfield configurations, two scenarios are considered for deployments:

- A single or multiple pod ACI fabric is in place already. You can add another site in a Multi-Site configuration.
- Two ACI fabrics are in place already, the objects (tenants, VRFs, and EPGs) across sites are initially defined with identical names and policies, and they are connected leveraging a traditional L2/L3 DCI solution. You can convert this configuration to Multi-Site as explained in the following configuration diagram:

Figure 15: Migration of Cisco ACI Fabric to Cisco ACI Multi-Site



501365

## Setting up Cisco ACI Multi-Site with Multipod-Enabled Fabrics

Starting in release 1.2(1), two use cases add support for setting up Cisco ACI Multi-Site with multipod-enabled fabrics.

Guidelines and limitations for these two use cases:

- Only the following switches will be connected to the IPN/ISN:
  - Cisco Nexus 93180LC-EX, 93180YC-EX, and 93108TC-EX switches.
  - Cisco Nexus 9504, 9408, and 9516 switches with the following line cards:
    - X9736C-EX
    - X97160YC-EX
    - X9732C-EX
    - X9732C-EXM

- Remove IPN links from old generation spine switches.
- The same IPN/ISN will be used for multipod and Multi-Site.
- In a Cisco ACI Multi-Site deployment, you cannot use an overlapping tunnel endpoints (TEP) pool range and GIPO pool range on the 2 sites using a single IPN/ISN.

When a tenant is imported from the Cisco APIC GUI, all the objects associated with the tenant are imported in Cisco ACI Multi-Site:

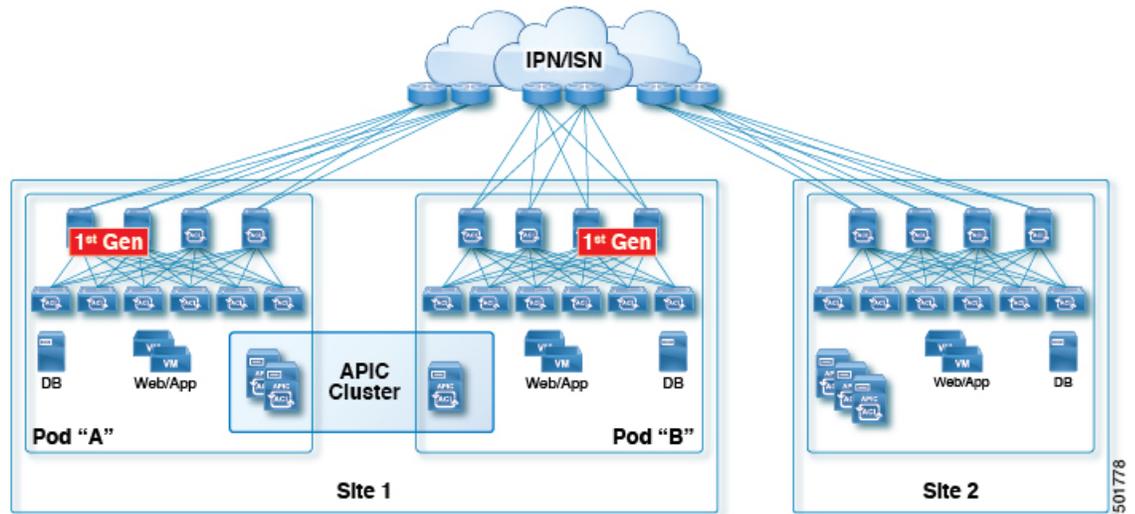
**Table 7: Features to be configured for these use cases**

Configuration	Description	Stretched or Local
Tenant	Create a tenant in Cisco ACI Multi-Site and import the tenant policies from the Cisco APIC	Stretched
VRF	VRF instance for the tenant	Stretched
Bridge Domain	Layer 2 stretching enabled Layer 2 flooding enabled Subnets to be shared added	Stretched
EPGs	EPGs in the BD	Stretched
Contracts	Include the filters needed to govern EPG communication	Stretched
Site L3Outs	Configured in the Cisco APIC and linked with external EPGs	Local

## Adding a Multipod Fabric as a Site on Cisco ACI Multi-Site

This section describes an overview of how to add a multipod fabric as a site on Cisco ACI Multi-Site.

Figure 16: Cisco ACI fabric with multiple PODs as a site in Cisco ACI Multi-Site



Fabric with multiple pods add as a site in Cisco ACI Multi-Site.

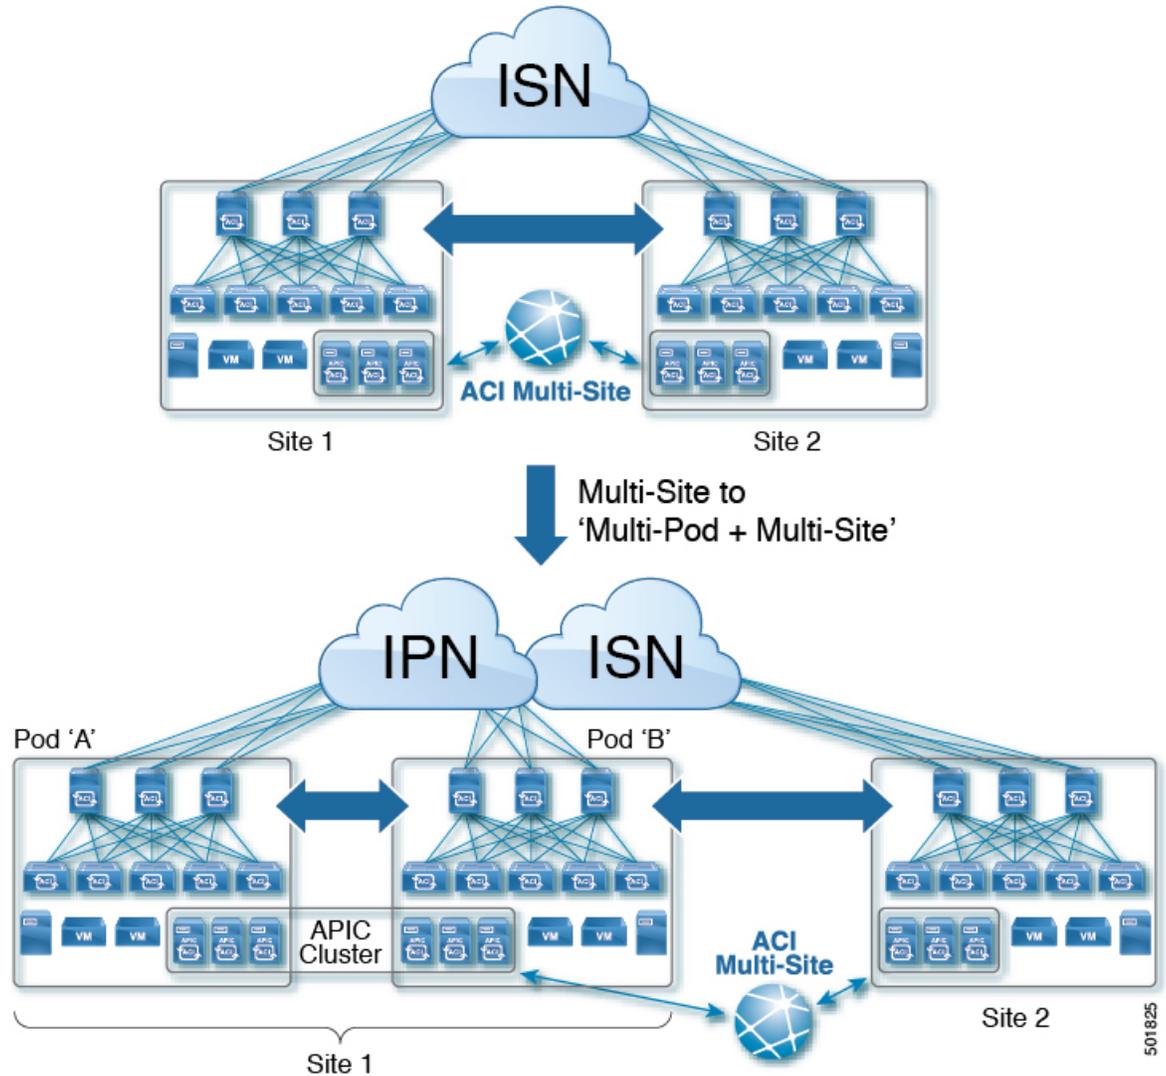
This is an overview of the procedure:

- Add a multipod-enabled fabric as a site in Cisco ACI Multi-Site.
  - Cisco ACI Multi-Site discovers common configurations for Cisco ACI Multi-Site and multipod, such as spine to IPN links configuration, OSPF information, BGP information and auto-populates in the Cisco ACI Multi-Site infra configuration.
- Provide Cisco ACI Multi-Site-specific configurations such as MCAST TEP, MSITE DP-TEP, or MSITE CP-TEP and enable Multi-Site for the site in Cisco ACI Multi-Site infra page.
  - You can also configure for Cisco ACI Multi-Site the same DP-TEP/CP-TEP that you configured for multipod.
- Apply the infra configuration in Cisco ACI Multi-Site.
  - Cisco ACI Multi-Site configures Cisco APIC with Cisco ACI Multi-Site-specific configurations and common configurations for Cisco ACI Multi-Site and multipod, such as spine to IPN links config, OSPF information, and BGP information, and will not configure multipod-specific configuration.
  - Cisco ACI Multi-Site uses the same infra L3Out used for multipod to configure Cisco ACI Multi-Site. Cisco ACI Multi-Site determines it based on fabricExtCtrlPeering=yes and fabricExtIntersiteCtrlPeering=yes under l3extInfraNodeP in the infra L3Out.
  - You can configure GOLF for an L3Out that you use with Cisco ACI Multi-Site by using one of the following configurations:
    - One L3Out for Cisco ACI Multi-Site, multipod, and GOLF, and different (zero or more) L3Outs for GOLF.
    - One L3Out for Cisco ACI Multi-Site, multipod and different (zero or more) L3Outs for GOLF.

## Converting a Single POD Site to Multipod

This section describes an overview of how to convert a single POD site to multipod.

Figure 17: Covering a single POD site in Multi-Site to a multipod site



Converting a single POD site to Multipod.

This is an overview of the procedure:

- Use the same spine nodes and uplinks for both communications.
- Use Cisco APIC to configure multipod. Use the same infra L3Out used for Cisco ACI Multi-Site for multipod also.
- You can use the same control-plane (CP) tunnel endpoints (TEP) and data-plane (DP) TEP for both multipod and Cisco ACI Multi-Site, or you can define separate TEPs for multipod and Cisco ACI Multi-Site.

- After configuring Cisco ACI Multi-Site, click on the "refresh" icon in the Cisco ACI Multi-Site infra page to discover the new pods.
- In Cisco ACI Multi-Site, provide Cisco ACI Multi-Site-specific configurations, such as Cisco ACI Multi-Site DP-TEP per pod and Cisco ACI Multi-Site CP-TEP.
- Apply infra.

