



Cisco ACI Multi-Site Fundamentals Guide, Release 3.1(x)

First Published: 2020-05-11

Last Modified: 2024-07-10

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

CHAPTER 1	New and Changed Information	1
	New and Changed Information	1

CHAPTER 2	About Cisco ACI Multi-Site	3
	About Cisco ACI Multi-Site	3
	Terminology	4
	Users, Roles, and Permissions	5
	Cisco ACI Multi-Site Schema and Templates	7

CHAPTER 3	Cisco ACI Multi-Site Use Cases	9
	Cisco ACI Multi-Site Service Integration	9
	Single-Node Service Graphs	10
	East-West FW Service Graph	10
	North-South FW Service Graph	12
	East-West LB Service Graph	13
	North-South LB Service Graph	15
	Two-Node Service Graphs	16
	East-West FW and IPS Service Graph	16
	East-West FW and LB Service Graph	18
	External EPG with Shared L3Out	20
	Configuring External EPG for Shared L3Out	23
	Shared Security Import Subnet Examples	24
	Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN	25
	Stretched Bridge Domain with Layer 2 Broadcast Extension	27
	Stretched Bridge Domain with No Layer 2 Broadcast Extension	29
	Stretched EPG Across Sites	31

Stretched VRF with Inter-Site Contracts 33

Shared Services with Stretched Provider EPG 35

Migration of Cisco ACI Fabric to Cisco ACI Multi-Site 38

 Setting up Cisco ACI Multi-Site with Multipod-Enabled Fabrics 39

 Adding a Multipod Fabric as a Site on Cisco ACI Multi-Site 40

 Converting a Single POD Site in Multi-Site to a Multipod Site 41



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 the organization and features in this guide from the release the guide was first published to the current release. The table does not provide an exhaustive list of all changes made to the guide.

Table 1: Latest Updates

Release	New Feature or Update	Where Documented
3.0(1)	--	--




CHAPTER 2

About Cisco ACI Multi-Site

- [About Cisco ACI Multi-Site, on page 3](#)
- [Terminology, on page 4](#)
- [Users, Roles, and Permissions, on page 5](#)
- [Cisco ACI Multi-Site Schema and Templates, on page 7](#)

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 a region (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. 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.

Users, Roles, and Permissions

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

- Power User—A role that allows the user to perform all the operations.
- Site Manager—A role that allows the user to manage sites, tenants, and associations between them.
- Schema Manager—A role that allows the user to manage all schemas regardless of their tenant associations.
- Schema Editor—A role that allows the user to manage schemas that contain at least one tenant to which the user is explicitly associated.
- User Manager—A role that allows the user to manage all the users, their roles, and passwords.

Each role above is associated with a set of permissions, which in turn are used to show relevant and hide irrelevant elements from the user's view of the Orchestrator GUI. For example, the User Manager role has only the user-related permissions associated with it and as such the user with that role will only see **Users** and **Admin** tabs in the GUI.

User Roles and Permissions

The following table lists the Cisco ACI Multi-Site permissions allowed with each available user role. The *Attribute-Value (AV)* column specifies the user configuration string required when configuring an external authentication server for use with the Multi-Site Orchestrator. External authentication is covered in more detail in the *Administrative Operations* chapter.

Table 2: User Roles

User Role	Permissions	Attribute-Value (AV) Pair
Power User	<ul style="list-style-type: none"> • Dashboard • Sites • Schemas • Tenants • Users • Troubleshooting Reports 	shell:misc-roles=powerUser
Site Manager	<ul style="list-style-type: none"> • Dashboard—Sites • Sites • Tenants 	shell:misc-roles=siteManager
Schema Manager	<ul style="list-style-type: none"> • Dashboard—Sites and Schema Health • Schemas 	shell:misc-roles=schemaManager
Schema Editor	<ul style="list-style-type: none"> • Dashboard—Sites and Schema Health • Schemas 	shell:misc-roles=schemaEditor
User Manager	<ul style="list-style-type: none"> • Users 	shell:misc-roles=userManager

Admin User

In the initial configuration script, a default `admin` user account is configured and is the only user account available when the system starts. The initial password for the `admin` user is set by the system and you are prompted to change it after the first log in.

- The `admin` user's default password is `welcome2misc!`
- The `admin` user is assigned the Power User role.
- 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**.

Read-Only Access

Each of the user roles above can be assigned in read-only mode. When read-only permissions are granted, the user can view any fabric objects available to that role just like before, but they cannot make any changes to those objects.

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

Cisco ACI Multi-Site Use Cases

- [Cisco ACI Multi-Site Service Integration, on page 9](#)
- [External EPG with Shared L3Out, on page 20](#)
- [Cisco ACI Multi-Site Back-to-Back Spine Connectivity Across Sites Without IPN, on page 25](#)
- [Stretched Bridge Domain with Layer 2 Broadcast Extension, on page 27](#)
- [Stretched Bridge Domain with No Layer 2 Broadcast Extension, on page 29](#)
- [Stretched EPG Across Sites, on page 31](#)
- [Stretched VRF with Inter-Site Contracts, on page 33](#)
- [Shared Services with Stretched Provider EPG, on page 35](#)
- [Migration of Cisco ACI Fabric to Cisco ACI Multi-Site, on page 38](#)

Cisco ACI Multi-Site Service Integration

Starting with Release 2.0(1), Cisco ACI Multi-Site supports service graphs with a load balancer and two-node service graphs with a load balancer and a firewall, in addition to the previously supported single-node graphs with a firewall.

Previous releases provided single-node service graphs support by applying PBR policies on the consumer's site for East-West traffic. In order to support two-node graphs in East-West scenario, PBR policies are now applied on the provider's site. While it prevents traffic from bouncing between sites in return data path, it requires a subnet to be configured under the consumer EPGs. In North-South scenario, PBR policies are still applied on the non-border leaf as they were in previous release.

To support the use cases described in this chapter, the following topology is required for service nodes:

- Each site has individual active/standby service node pair
- Layer 4 to Layer 7 devices are in un-managed mode
- VRFs are stretched across sites
- Consumer and provider EPGs have cross-site contract


In addition to the topology requirements above, keep in mind the following considerations:

- External and internal connector of a node can be same logical interface
- In case of East-West traffic, a subnet must be configured under consumer EPGs
- In case of East-West inter-VRF traffic, a subnet must be configured on both consumer and provider EPGs

- In case of North-South traffic, policies are applied on the non-border leaf
- Shared service scenario is supported for East-West traffic, but not for North-South traffic

A sample topology used throughout the use-cases in this chapter is shown below:

Figure 3: Cisco ACI Multi-Site Service Integration Topology



Single-Node Service Graphs

East-West FW Service Graph

This is the use case for East-West communication with a Firewall (FW) between endpoints in the same VRF or different VRFs across sites.


The following two local PBR policies are required for the Service Graph:

- PBR policy on FW's external connector to redirect consumer-to-provider traffic to FW's external interface
- PBR policy on FW's internal connector to redirect provider-to-consumer traffic to FW's internal interface

The following figures show incoming traffic packet flow from consumer on Site1 to provider on Site2:

- Consumer leaf does not apply any rules, forwards traffic to the provider
- Provider leaf applies policy and send traffic to FW2
- Finally, traffic is sent to the provider EPG


Figure 4: East-West FW Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Provider leaf applies policy to redirect traffic to FW2
- Traffic is then sent to consumer on Site1
- Consumer leaf does not apply any rules, forwards traffic to consumer EPG

Figure 5: East-West FW Reverse Traffic




North-South FW Service Graph

This is the use case for North-South communication with a Firewall (FW) between endpoints in the same VRF across sites.

The following figures show incoming traffic packet flow from consumer on Site1 to provider on Site2:

- Consumer border leaf does not apply any rules, forwards traffic to the provider
- Non-border leaf on provider's site applies policy and sends traffic to FW2's external interface
- Finally, traffic is sent to the EPG


Figure 6: North-South FW Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Non-border leaf on provider's site applies policy to redirect traffic to FW2's internal connector
- Traffic is then sent out the Site2's L3Out

Figure 7: North-South FW Reverse Traffic



East-West LB Service Graph

This is the use case for East-West communication with the Load-Balancer (LB) between endpoints in the same VRF or different VRFs across sites. Service Graphs with LB are different from the ones with a Firewall (FW), because in this case the traffic is destined for the VIP of the LB. This use-case describes a scenario where the LB is in one site with local provider EPG and consumer is in another site.


The following figures shows incoming traffic packet flow from consumer on Site1 to provider (EPG App) on Site2:

- Consumer leaf does not apply any rules, forwards traffic to the provider
- Traffic is forwarded to LB2's VIP
- Finally, traffic is sent to the provider EPG



Note The example in this section uses no SNAT on the load-balancer. PBR is for return traffic to LB, as such if LB does SNAT, PBR is not necessary. Also, keep in mind that in case of no SNAT and PBR, the LB's VIP and its real servers must be in same site.


Figure 8: East-West LB Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Provider leaf applies policy to redirect traffic to LB2
- Traffic is then sent to consumer on Site1
- Consumer leaf does not apply any rules, forwards traffic to consumer EPG

Figure 9: East-West LB Reverse Traffic



North-South LB Service Graph

This is the use case for North-South communication with a Load-Balancer (LB) between endpoints in the datacenter and outside. The following diagram shows the packet flow for a scenario where L3Out traffic enters from the Site that is not hosting the LB for which the traffic is directed (VIP is in different site). In this we have L3Out as Consumer and regular EPG as provider. In this case policy is always applied on the provider site's non-border leaf.


The following figures show incoming traffic packet flow from consumer on Site1 to provider on Site2:

- Consumer border leaf does not apply any rules, forwards traffic to VIP on the Site2
- Non-border leaf on provider's site applies policy and traffic is forwarded to the LB
- Finally, traffic is sent to the provider EPG from LB



Note The example in this section uses no SNAT on the load-balancer. PBR is for return traffic to LB, as such if LB does SNAT, PBR is not necessary. Also, keep in mind that in case of no SNAT and PBR, the LB's VIP and its real servers must be in same site.


Figure 10: North-South LB Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Non-border leaf on provider's site applies policy to redirect traffic to LB
- Traffic is then sent out the Site2's L3Out

Figure 11: North-South LB Reverse Traffic



Two-Node Service Graphs


East-West FW and IPS Service Graph

This is the use case for East-West communication with a Firewall (FW) and an Intrusion Prevention System (IPS) between endpoints in the same VRF or different VRFs across sites.

The following figures shows incoming traffic packet flow from consumer on Site1 to provider on Site2:

- Consumer leaf does not apply any rules, forwards traffic to the provider
- Provider leaf applies policy and send traffic to FW2's external interface
- Traffic is then redirected back to the provider leaf and then to IPS2's external interface
- Finally, traffic is sent to the provider EPG


Figure 12: East-West FW/IPS Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Provider leaf applies policy to redirect traffic to IPS2's internal connector
- Traffic is then redirected to FW2's internal connector
- Traffic is then sent to consumer on Site1
- Consumer leaf does not apply any rules, forwards traffic to consumer EPG

Figure 13: East-West FW/IPS Reverse Traffic




East-West FW and LB Service Graph

This is the use case for East-West communication with the Firewall (FW) and Load-Balancer (LB) between endpoints in the same VRF or different VRFs across sites. This is a common design for traffic within the application that requires the server load-balancing for high availability and scale.

The following figures shows incoming traffic packet flow from consumer on Site1 to provider on Site2:

- Consumer leaf does not apply any rules, forwards traffic to the provider
- Provider leaf where the LB2's VIP is connected applies policy and send traffic to FW2's external interface
- Traffic is then redirected to LB2
- Finally, traffic is sent to the provider EPG


Figure 14: East-West FW/LB Incoming Traffic



The following figures shows return traffic packet flow from provider on Site2 to consumer on Site1:

- Provider leaf applies policy to redirect traffic to LB2
- Traffic is then redirected to FW2's internal connector
- Traffic is then sent to consumer on Site1
- Consumer leaf does not apply any rules, forwards traffic to consumer EPG

Figure 15: East-West FW/LB Reverse Traffic



External EPG with Shared L3Out

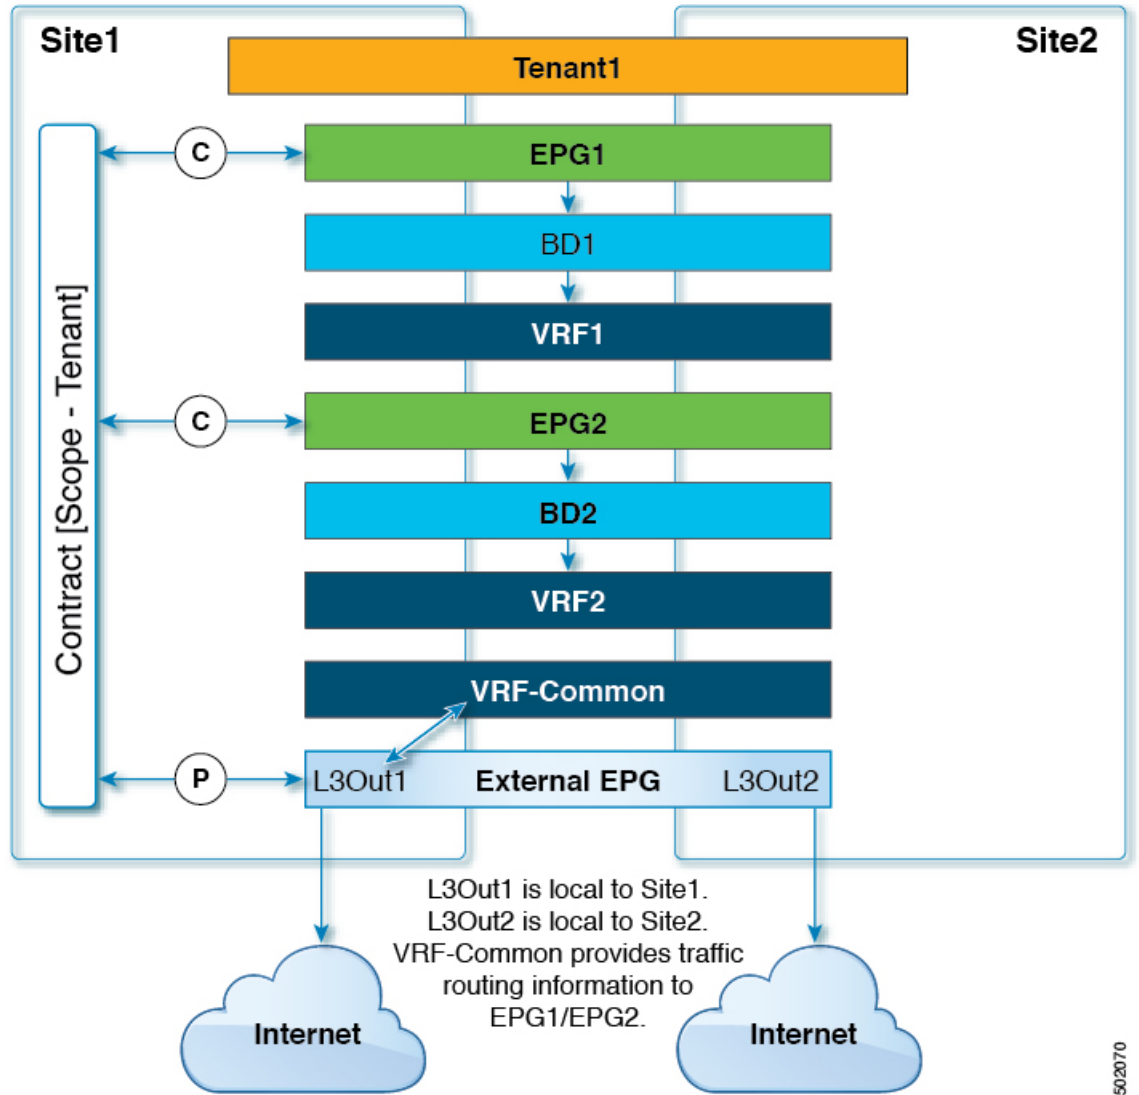
Starting with Release 2.0(1), Cisco ACI Multi-Site supports shared services, with consumer and provider EPGs in different VRFs and with L3Out External EPG as a provider or consumer. Previous versions of Multi-Site supported this use-case only when the L3Out and consumer EPGs were in the same VRF.

The most common use-case for this is an External EPG deployed in the `common` tenant that provides Internet service, while the other tenants, or consumer EPGs, use it for Internet access. But in addition, this feature also enables the following use-cases:



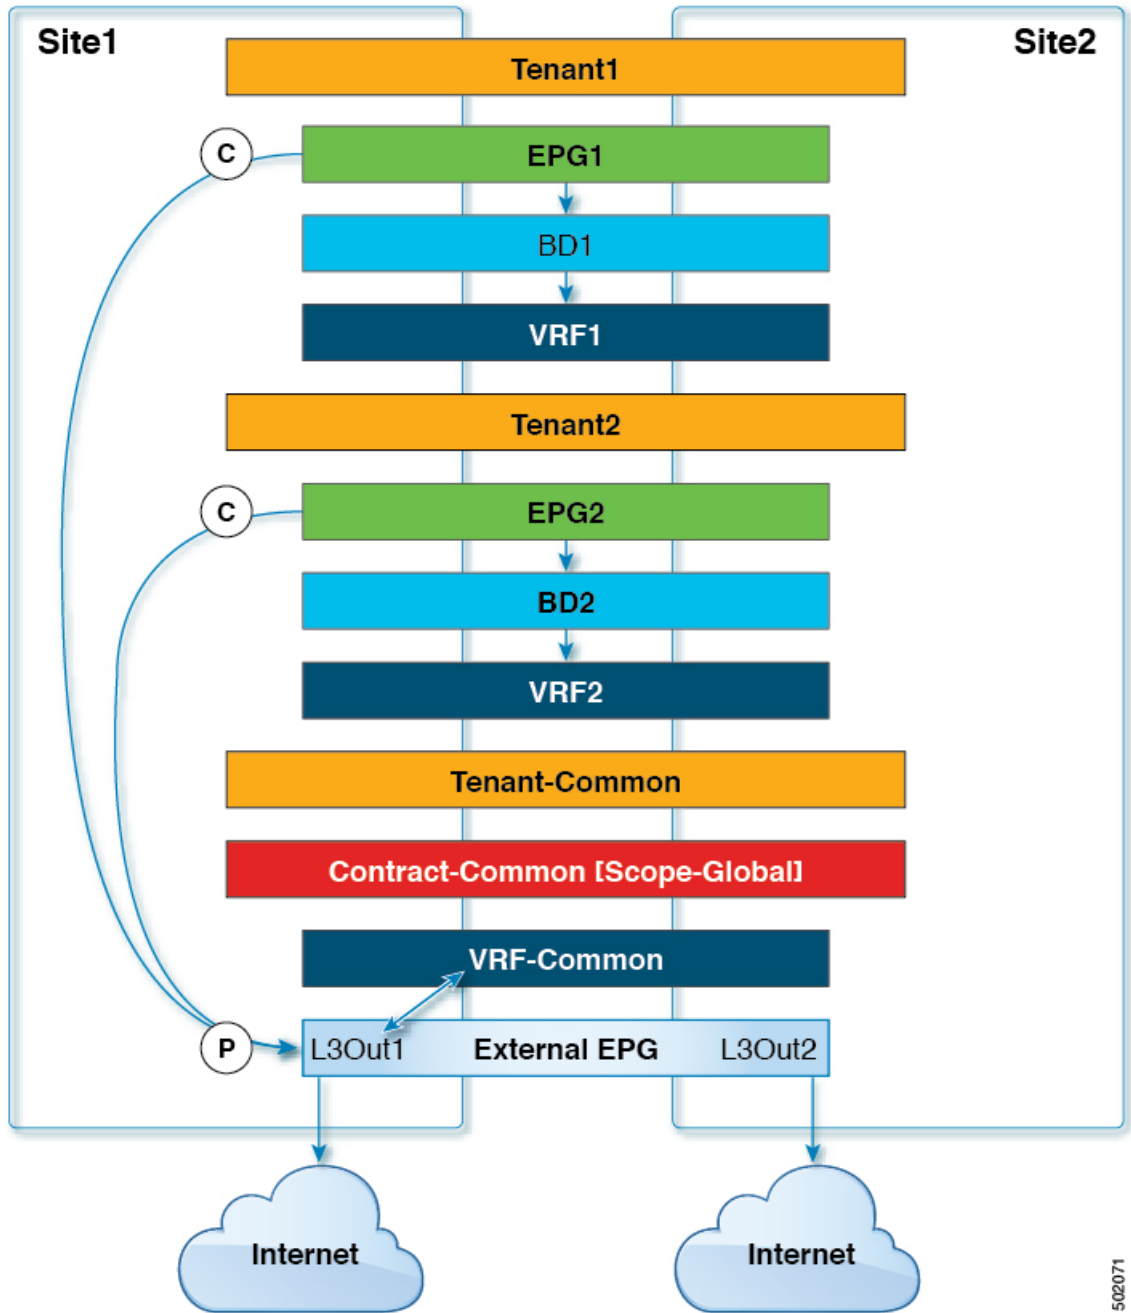
Note The external EPG in the following examples is shown as a provider, however the same applies for cases where external EPG is a consumer instead.

Figure 16: EPGs, VRFs, Bridge Domains, and External EPG under User Tenant



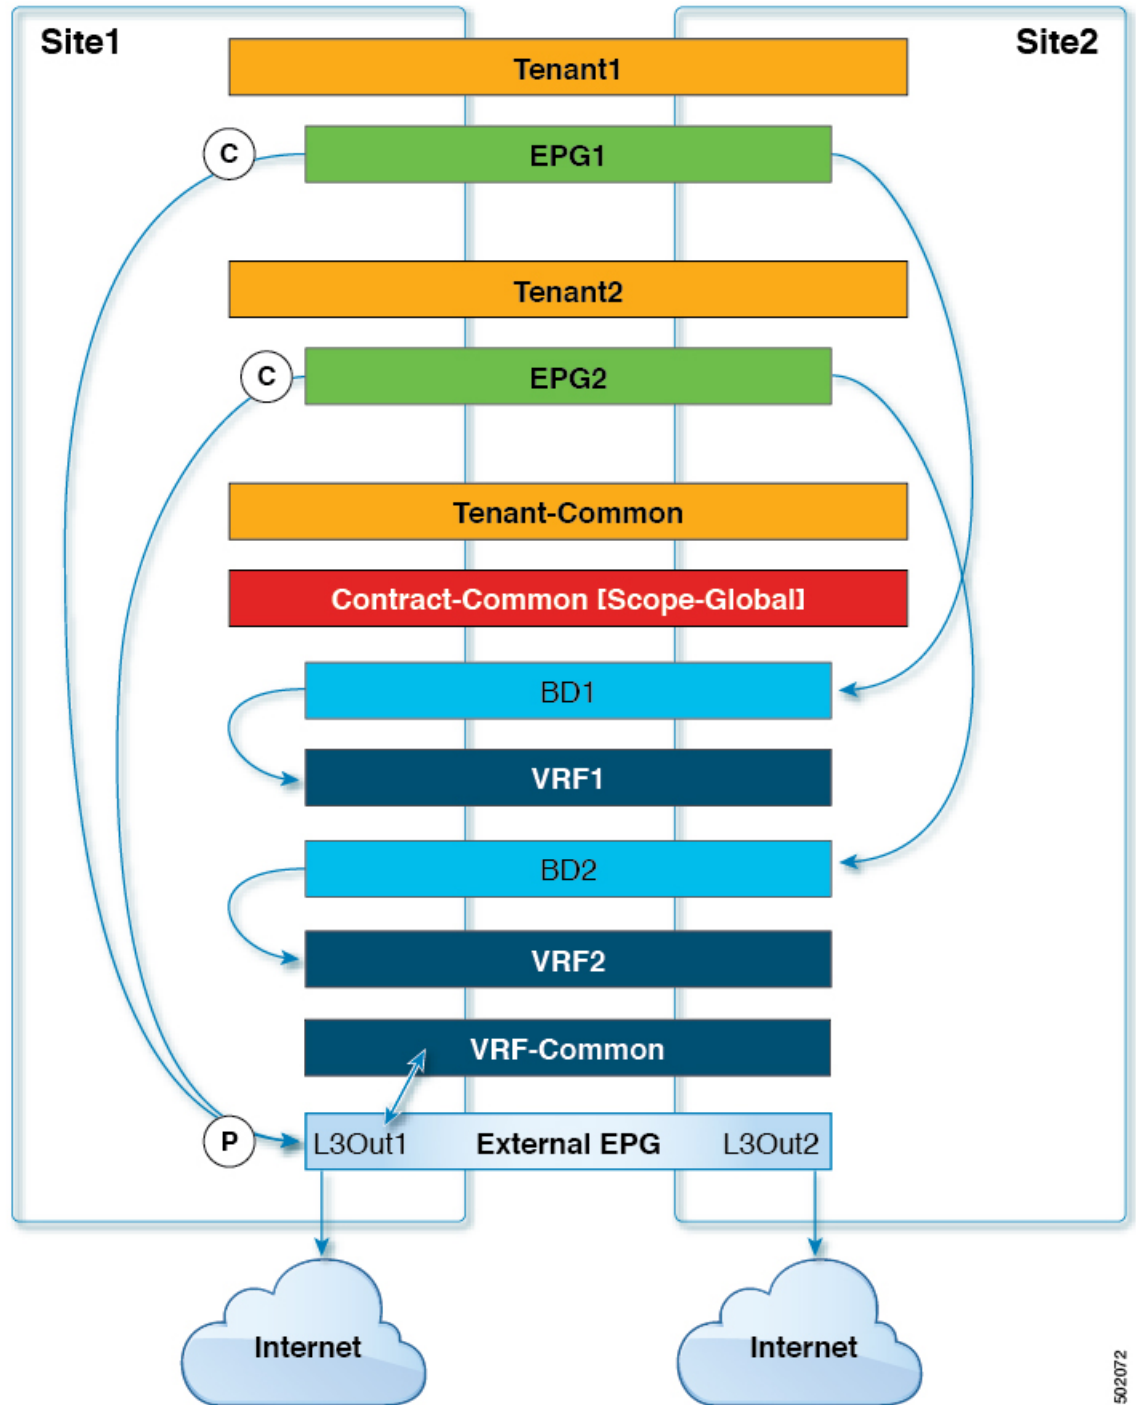
502070

Figure 17: EPGs, VRFs, and Bridge Domains under a User Tenant. External EPG under commonTenant



502071

Figure 18: EPGs under a User Tenant, but VRF, BDs, and External EPG under common Tenant



502072

Configuring External EPG for Shared L3Out

The following steps describe how to configure route control properties and route control profile name on the External EPG at the template level.

-
- Step 1** Log in to the Multi-Site GUI.
- Step 2** From the left-hand sidebar, select the **Schemas** view.
- Step 3** Select a Schema.
- Step 4** Select the External EPG you want to configure.
- Step 5** In the right-hand pane, click **+SUBNET** to add a subnet.

Fill in the following fields:

- **Classification Subnet** – Routes from the L3Out matching this subnet and external traffic within the defined prefix is classified as part of this external EPG.
- **Shared Route Control Subnet** – Determines whether the defined route is leaked to the VRF with which it is shared.
- **Shared Security Import Subnet** – Provides for a more granular control of the routing and policy planes. For additional information on specific examples for this setting, see [Shared Security Import Subnet Examples, on page 24](#).
- **Aggregate Shared Routes** – Determines whether all prefixes that fall within the defined route are leaked to the VRF with which it is shared. Aggregate Shared Routes can be enabled only when **Shared Route Control Subnet** is enabled

The route can be leaked into the other private network, but no ACLs will be installed for the route in the other network. Such a scenario is possible when the shared route is in a bigger subnet, while security is applied on a separate smaller subnet.

Shared Security Import Subnet Examples

This section provides examples of using **Shared Security Import Subnet** setting to configure more granular control of the routing and policy planes.

The three use-case examples below assume L3Out received the following three routes: 10.0.1.0/24, 10.0.2.0/24, and 10.0.3.0/24.

Use-Case 1

Two external EPGs are configured with the following settings:

- External EPG1 with 10.0.1.0/24 subnet, **Shared Route Control Subnet** and **Shared Security Import Subnet** settings enabled, and a contract to allow only TCP traffic
- External EPG2 with 10.0.2.0/24 subnet, **Shared Route Control Subnet** and **Shared Security Import Subnet** settings enabled, and a contract to allow only UDP traffic

In this case, only 10.0.1.0/24 and 10.0.2.0/24 prefixes are leaked to the other VRF and different contracts can be specified to allow only TCP based traffic for the 10.0.1.0/24 subnet and only UDP traffic for the 10.0.2.0/24 subnet.

Use-Case 2

Three external EPGs are configured with the following settings:

- External EPG1 with 10.0.0.0/16 subnet, **Shared Route Control Subnet** and **Aggregate Shared Routes** settings enabled
- External EPG2 with 10.0.1.0/24 subnet, only **Shared Security Import Subnet** setting enabled, and a contract to allow only TCP traffic
- External EPG3 with 10.0.2.0/24 subnet, only **Shared Security Import Subnet** setting enabled, and a contract to allow only UDP traffic

In this case, aggregated 10.0.0.0/16 defined in EPG1 will ensure that subsets 10.0.1.0/24, 10.0.2.0/24, and 10.0.3.0/24 are leaked to the other VRF, but ACLs are installed only for 10.0.1.0/24 and 10.0.2.0/24.

For 10.0.1.0/24 and 10.0.2.0/24, different contracts can be specified to allow only TCP traffic for the 10.0.1.0/24 subnet and only UDP traffic for the 10.0.2.0/24 subnet. No ACLs will be present for 10.0.3.0/24 resulting in policy drop even though route is leaked to the other VRF.

Use-Case 3

Three external EPGs are configured with the following settings:

- External EPG1 with 10.0.0.0/16 subnet, **Shared Security Import Subnet** setting enabled, and a contract to allow all traffic
- External EPG2 with 10.0.1.0/24 subnet, **Shared Route Control Subnet** setting enabled
- External EPG3 with 10.0.2.0/24 subnet, **Shared Route Control Subnet** setting enabled

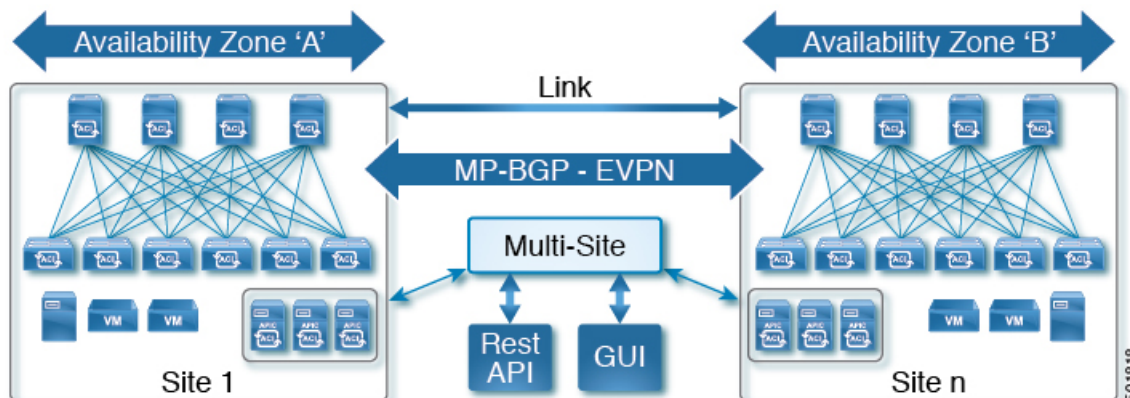
In this case, only 10.0.1.0/24 and 10.0.2.0/24 are leaked into the other VRF as they are marked as shared route, while 10.0.3.0/24 is not leaked. There will be a single ACL to allow all traffic within subnet 10.0.0.0/16 to the L3Out.

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. 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

Figure 19: Multi-Site Back to Back Spine – Basic Setup



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.
- Only two sites are supported with back-to-back spine.
- No new configuration required in APIC for this use case.

Troubleshooting

- 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-infra1-spine1# cat /mit/sys/lldp/inst/if-[eth1--1]/summary | grep wiringIssues
wiringIssues :
dev-infra1-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:

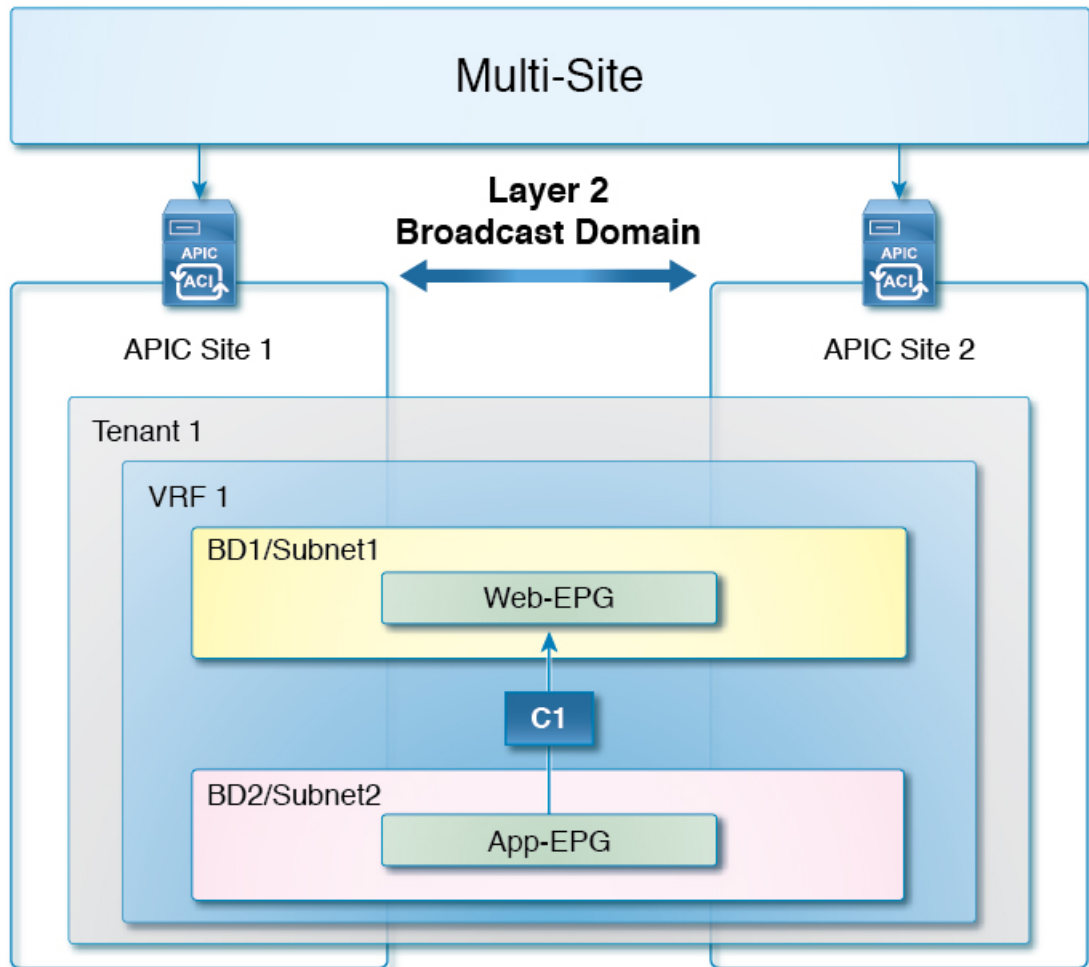
```
$ zgrep "back to back spine ignore wiring check." \  
/var/sysmgr/tmp_logs/dme_logs/svc_ifc_policyelem.log*
```

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 20: 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 3: 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 21: 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:

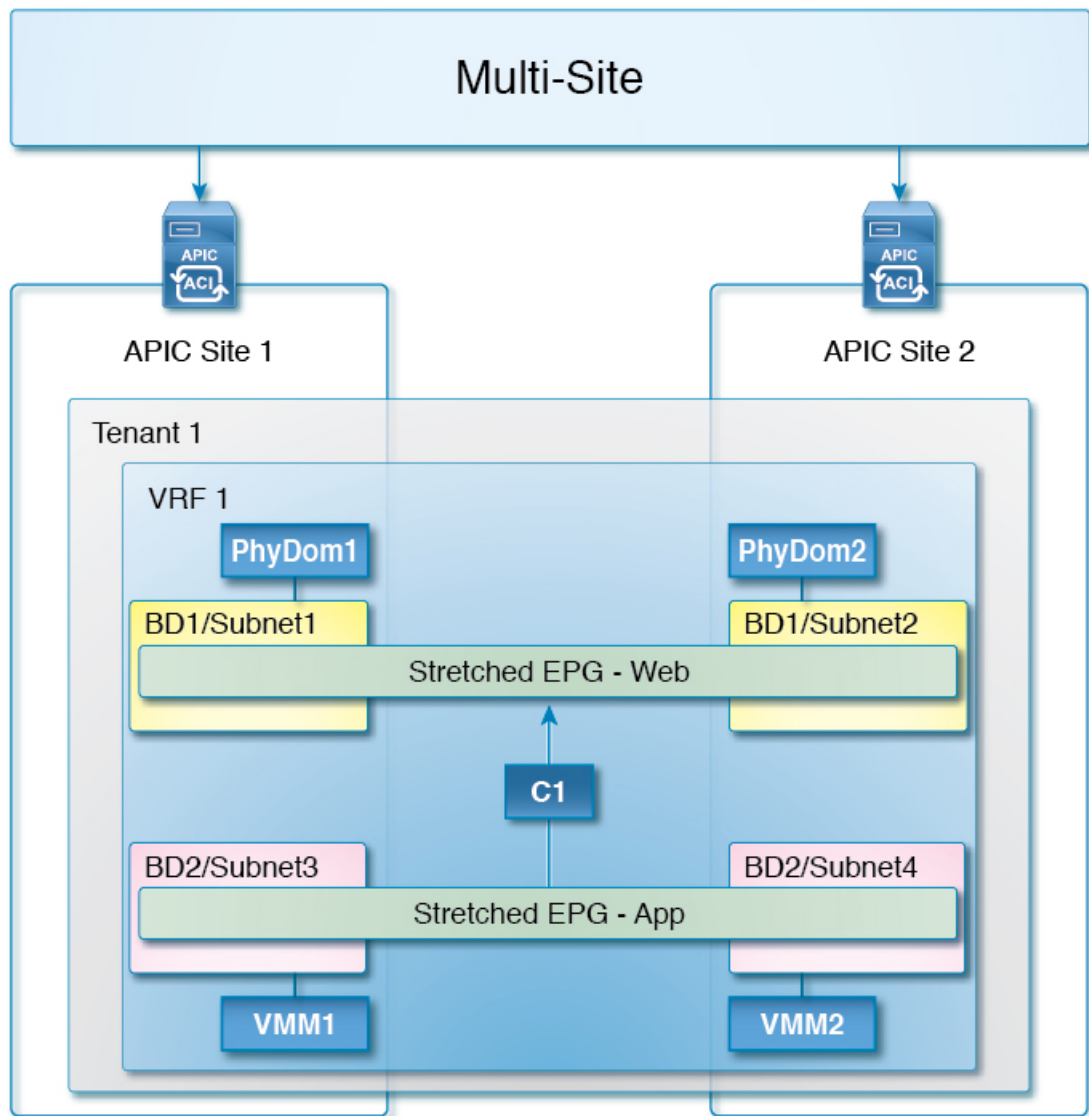
Table 4: Features to Be Configured for this Use Case

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 22: 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:

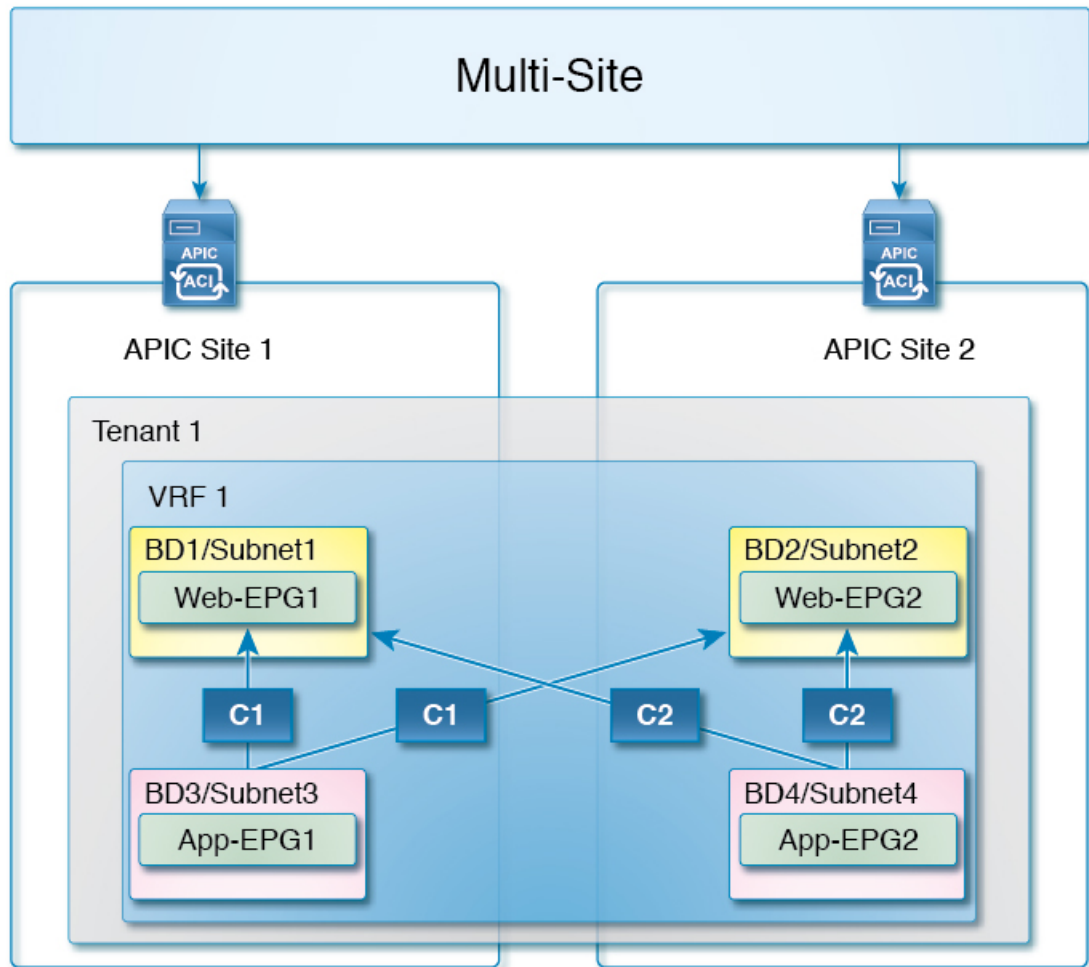
Table 5: Features to be Configured for this Use Case

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 23: 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 6: 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 24: 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 7: Features to be Configures 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"> • Tenant • VRF • Provider Contract with global scope. • EPG with subnet set to Advertised Externally and Shared Between VRFs. 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"> • Tenant • VRF • EPG with subnet set to Advertised Externally and Shared Between VRFs. <p>Note For the consumer EPGs, the subnets can alternatively be added in the BDs.</p> <ul style="list-style-type: none"> • Consumer Contract (same name as the provided contract). <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 25: Migration of Cisco ACI Fabric to Cisco ACI Multi-Site



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 8: 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 26: Cisco ACI fabric with multiple PODs 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 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 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 Multi-Site the same DP-TEP/CP-TEP that you configured for multipod.
- Deploy the infra configuration in Cisco ACI Multi-Site.
 - Cisco ACI Multi-Site configures Cisco APIC with Multi-Site-specific configurations and common configurations for Multi-Site and multipod, such as spine 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 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 in the same L3Out that you use for Multi-Site and multipod. The supported configurations are:
 - One L3Out for Multi-Site, multipod, and GOLF, and different (zero or more) L3Outs for GOLF.
 - One L3Out for Multi-Site, multipod and different (zero or more) L3Outs for GOLF.

Converting a Single POD Site in Multi-Site to a Multipod Site

This section describes an overview of how to convert a single POD site in Multi-Site to a multipod site.

Figure 27: Converting a single POD site in Multi-Site to a multipod site



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 Multi-Site for multipod also.
- You can use the same BGP-EVPN Router-ID and Overlay TEP for both multipod and Multi-Site, or you can define separate Router-ID and TEP for multipod and 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 Multi-Site-specific configurations, such as Overlay TEP and BGP-EVPN Router-ID.
- Deploy infra.