Cisco ACI Multi-Site Configuration Guide, Release 2.2(x)

First Published: 2019-06-24

Last Modified: 2020-04-23

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

CONTENTS

PREFACE

Preface ix

- Document Conventions ix
- Related Documentation x
- Documentation Feedback xi
- Obtaining Documentation and Submitting a Service Request xi

CHAPTER 1

New and Changed Information 1

- New and Changed Information 1

CHAPTER 2

Overview of the GUI 3

- Overview of the Cisco ACI Multi-Site Orchestrator GUI 3
- Dashboard 4
- Sites Page 7
- Schemas Page 8
- Tenants Page 10
- Users Page 12
- Policies Page 13
- Admin Page 13

CHAPTER 3

Infrastructure Management 21

- Cisco ACI Multi-Site and Cisco APIC Interoperability Support 21
- Multi-Site Orchestrator Communication Ports 22
- Configuring Fabric Access Policies for All APIC Sites 23
  - Configuring Fabric Access Global Policies 23
  - Configure Fabric Access Interface Policies on Each APIC 24
- Configuring Sites That Contain Remote Leaf Switches 26

Multi-Site and Remote Leaf Guidelines and Limitations **26**

Configuring Routable Subnets for Remote Leaf Switches **26**

Enabling Direct Communication for Remote Leaf Switches **27**

Adding Sites **27**

Configuring Infra Prerequisites and Guidelines **28**

Configuring Infra: General Settings **29**

Refreshing Site Connectivity Information **29**

Configuring Infra Site-Specific Settings **30**

Configuring Infra: Pod Settings **31**

Configuring Infra: Spine Switches **32**

Deploying Infra Configuration **33**

Deleting Sites Using Multi-Site Orchestrator GUI **33**

Cisco ACI CloudSec Encryption **33**

Requirements and Guidelines **34**

CloudSec Encryption Terminology **35**

CloudSec Encryption and Decryption Handling **36**

CloudSec Encryption Key Allocation and Distribution **38**

Configuring Cisco APIC for CloudSec Encryption **40**

Configuring Cisco APIC for CloudSec Encryption Using GUI **41**

Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI **41**

Configuring Cisco APIC for CloudSec Encryption Using REST API **42**

Enabling CloudSec Encryption Using Cisco ACI Multi-Site Orchestrator GUI **43**

Rekey Process During Spine Switch Maintenance **43**

Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI **44**

Disabling and Re-Enabling Re-Key Process Using REST API **44**

Multi-Site Cross Launch to Cisco APIC **45**

Cross-Launch to Cisco APIC from Sites **45**

Cross-Launch to Cisco APIC from Schemas **46**

Cross-Launch to Cisco APIC from the Property Pane **46**

Viewing Cisco ACI Multi-Site-Managed Objects Using the Cisco APIC GUI **46**

CHAPTER 4 **Tenant Management** **49**

Managing Tenants Guidelines **49**

Adding Tenants **50**

Configuring Global Contracts Across Tenants or VRFs 50
Configuring Intra-EPG Isolation Using the Multi-Site GUI 52
Configuring Microsegmented EPGs Using the Multi-Site GUI 53
Associating EPGs with Domains Using the Multi-Site GUI 54
Displaying All the Tenants in an Aggregated View 55

CHAPTER 5 **User Management** 57
Users, Roles, and Permissions 57
Guidelines for User Management 59
Adding a User 59
Managing Users 60

CHAPTER 6 **Schema Management** 61
Schema Design Considerations 61
Schema Design for Cisco Cloud APIC Use-Cases 67
Creating a Schema Template 67
Importing Schema Elements From APIC Sites 67
Configuring an Application Profile 68
Configuring a VRF for the Tenant 69
Configuring a Bridge Domain 69
Configuring a Filter for Contracts 70
Configuring a Contract 70
Configuring an External EPG 71
Configuring an L3Out 71
Viewing Schemas 72
Migrating Objects Between Templates 72
Intersite L3Out 73
Intersite L3Out Guidelines and Limitations 73
Configuring Routable TEP Addresses 75
Creating or Importing Intersite L3Out and VRF 75
Configuring an External EPG to Use the Intersite L3Out 76
Creating a Contract for Intersite L3Out 77
Configuring Intersite L3Out for Application EPGs 78
Configuring Transit L3Out Across Sites 80

Shared Services with Intersite L3Out 81
EPG Preferred Groups 82
Configuring EPGs for Preferred Group 83
Layer 3 Multicast 84
Enabling Layer 3 Multicast 84
Shadow EPGs and BDs 85

CHAPTER 7 Administrative Operations 87
Viewing Site Status 87
Viewing Schema Health 87
Viewing Faults for Individual Sites 88
DHCP Relay Policy 89
Guidelines and Limitations 89
Creating DHCP Relay Policies 90
Creating DHCP Option Policies 91
Assigning DHCP Policies 92
Creating DHCP Relay Contract 93
Verifying DHCP Relay Policies in APIC 94
Editing or Deleting Existing DHCP Policies 95
System Logs 95
Generating Troubleshooting Report and System Logs 95
Enabling Log Streaming to an External Log Analyzer 96
Configuration Backup and Restore 97
Backup and Restore Guidelines 97
Remote Backups 98
Configuring a Remote Location for Backups 99
Moving Existing Backups to a Remote Location 100
Adding an NFS Share to Store Backups 100
Creating Backups 101
Restoring Backups 102
Downloading Backups 103
Importing Backups 103
Custom SSL Certificates 103
Adding Custom Certificate Authority 104

- Adding Custom Keyring 104
- Activating Custom Keyring 105
- Custom Certificates Troubleshooting 105

External Authentication 106

- Guidelines for Configuring External Authentication Servers 106
- Adding RADIUS or TACACS+ as Authentication Provider 107
- Adding LDAP as Authentication Provider 108
- Creating Login Domains 109
- Editing, Deleting, or Deactivating Login Domains 110
- Remote User Logon 110

System Configuration Settings 110

- System Alias and Banner 111
- Login Attempts and Lockout Time 111
- Proxy Server 112

Preface

This preface includes the following sections:

- Document Conventions, on page ix
- Related Documentation, on page x
- Documentation Feedback, on page xi
- Obtaining Documentation and Submitting a Service Request, on page xi

Document Conventions

Command descriptions use the following conventions:

Convention	Description
bold	Bold text indicates the commands and keywords that you enter literally as shown.
Italic	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.
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
`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	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.

This document uses the following conventions:

Note Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Warning IMPORTANT SAFETY INSTRUCTIONS

This warning symbol means danger. You are in a situation that could cause bodily injury. Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Use the statement number provided at the end of each warning to locate its translation in the translated safety warnings that accompanied this device.

SAVE THESE INSTRUCTIONS

Related Documentation

The following documentation provides additional information on Cisco ACI Multi-Site:

- *Cisco ACI Multi-Site Fundamentals Guide*
- *Cisco ACI Multi-Site Orchestrator Installation and Upgrade Guide*
- *Cisco ACI Multi-Site Configuration Guide*
- *Cisco ACI Multi-Site REST API Configuration Guide*
- *Cisco ACI Multi-Site Troubleshooting Guide*

All these documents are available at the following URL: http://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html

Documentation Feedback

To provide technical feedback on this document, or to report an error or omission, please send your comments to apic-docfeedback@cisco.com. 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 and Changed Information

Cisco ACI Multi-Site Version	Feature	Description	Where Documented
2.2(3)	Proxy configuration	A proxy server can be configured and enabled for Multi-Site Orchestrator connections to cloud sites.	For more information, see Proxy Server, on page 112.
2.2(1)	Cisco ACI Multi-Site and Cisco APIC Interoperability Support	Mixed operation support has been added to allow different release versions of Multi-Site Orchestrator and APIC to function within the same environment	For more information, see DHCP Relay Policy, on page 89.
2.2(1)	Object Migration Between Schemas and Templates	Bridge Domains and EPGs can now be moved between templates within the same or different schemas	For more information, see Migrating Objects Between Templates, on page 72.
2.2(1)	Intersite L3Out	A site can use another site's L3Out	For more information, see Intersite L3Out, on page 73.

Cisco ACI Multi-Site Version	Feature	Description	Where Documented
2.2(1)	DHCP Relay	A DHCP relay can be configured for endpoints in one EPG to access a DHCP server in another EPG or external EPG	For more information, see DHCP Relay Policy, on page 89.
2.2(1)	Remote Backup	Remote back up locations can be configured in the Orchestrator GUI.	For more information, see Remote Backups, on page 98.

CHAPTER 2

Overview of the GUI

- Overview of the Cisco ACI Multi-Site Orchestrator GUI, on page 3
- Dashboard, on page 4
- Sites Page, on page 7
- Schemas Page, on page 8
- Tenants Page, on page 10
- Users Page, on page 12
- Policies Page, on page 13
- Admin Page, on page 13

Overview of the Cisco ACI Multi-Site Orchestrator GUI

The Cisco ACI Multi-Site (Multi-Site) Orchestrator GUI is a browser-based graphical interface for configuring and monitoring your Cisco ACI, APIC, and Cloud APIC deployments.

The GUI is arranged according to function. For example, the **Dashboard** page contains an overview of your fabrics and their health, the **Sites** page provides information on each site and allows you to add sites, the **Schemas** page allows you to create and configure schemas, and so on.

The functionality of each Multi-Site Orchestrator GUI page is described in the following sections:

- Dashboard, on page 4
- Sites Page, on page 7
- Schemas Page, on page 8
- Tenants Page, on page 10
- Users Page, on page 12
- Policies Page, on page 13
- Admin Page, on page 13

The top of each page shows the controller status indicating how many controllers are operational, the **Get Started** menu icon, the **Settings** icon, and the **User** icon.

The **Get Started** menu provides easy access to a number of common tasks you may want to perform, such as adding sites or schemas, configuring policies, or performing administrative tasks.

The **Settings** icon allows you to access overview information about your Multi-Site Orchestrator, such as the currently running version, what's new in the current release, system logs, and Swagger API documentation.

- Clicking the **What's New in This Release** link displays a short summary of the new features in your release, as well as links to the rest of the Multi-Site documentation.
- Clicking the **System Logs** link allows you to configure and download system event logs, which is described in more detail in the *Administrative Operations* chapter in this guide.
- Clicking the **View Swagger Docs** link gives you access to the set of Swagger API object and method references. Using the Swagger API is described in more detail in the *Cisco ACI Multi-Site REST API Configuration Guide*.

The **User** icon allows you to view information about the currently logged in user, such as password updates, preferences, bookmarks. It also allows you to log out of the Orchestrator GUI.

- The **Reset Password** link allows you to update the currently logged in user's password
- The **Preferences** link allows you to change a few GUI options.
- The **Bookmarks** link opens the list of all the bookmarked schemas you save while using the Orchestrator. You can bookmark a schema by clicking the bookmark icon in the top right corner of the screen while viewing or editing the schema.

When working with fabric objects, a **Display Name** field is used throughout the Orchestrator's GUI whenever the objects are shown. You can specify a display name when creating the objects, however due to object naming requirements on the Cisco APIC, any invalid characters are removed and the resulting **Internal Name** is used when pushing the objects to sites. The **Internal Name** that will be used when creating the tenant is typically displayed below the **Display Name** textbox.

Dashboard

The Multi-Site dashboard displays the list of all of your site implementations in addition to their current functionality and health.

The following screen shot shows the Multi-Site dashboard display:

Figure 1: Multi-Site Dashboard

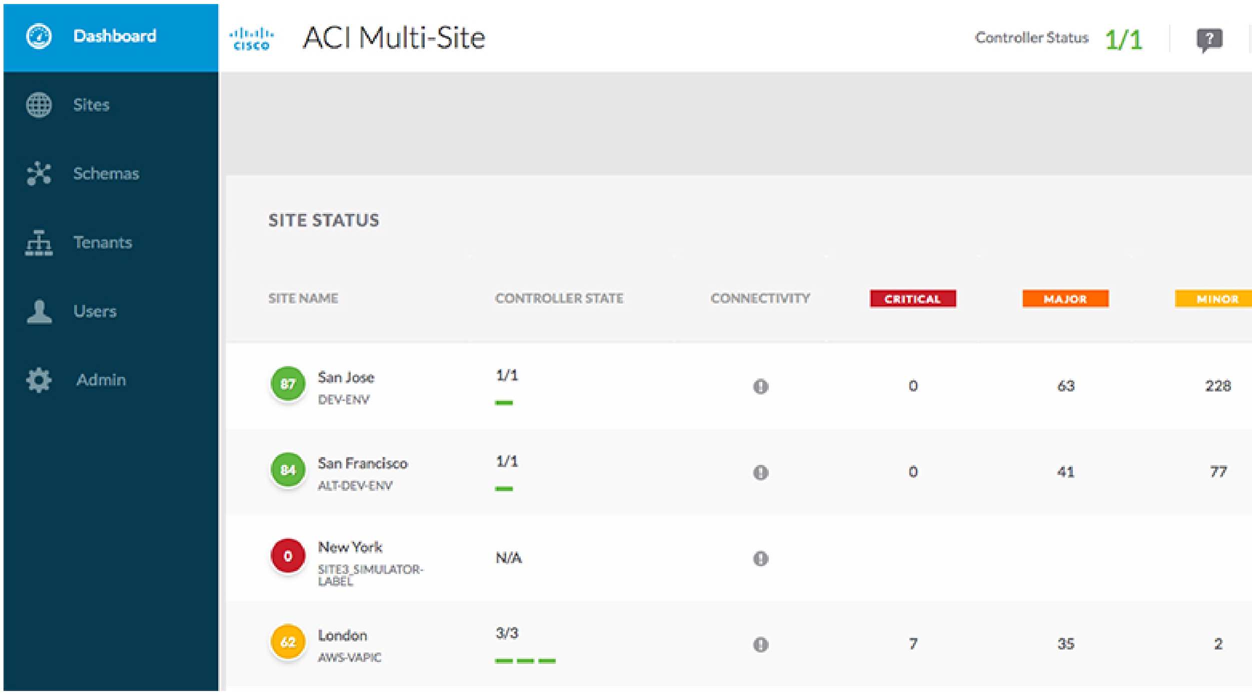

SITE STATUS

SITE NAME	CONTROLLER STATE	CONNECTIVITY	CRITICAL	MAJOR	MINOR
87 San Jose DEV-ENV	1/1	●	0	63	228
84 San Francisco ALT-DEV-ENV	1/1	●	0	41	77
0 New York SITE3_SIMULATOR-LABEL	N/A	●			
62 London AWS-VAPIC	3/3	●	7	35	2

The **Dashboard** has the following functional areas:

- **Site Status**: The site status table lists your sites according to name and location. The table also indicates the current health status for your implementation according to a descriptive color code.
  - The Controller State column indicates the number of controllers available and running. You can have a maximum number of 3 controllers in your Multi-Site implementation. For example, if one out of the 3 controller is down it is represented as 2/3.
  - The Connectivity column provides an operational status of the BGP sessions and the dataplane unicast and multicast tunnels that are connected to the peer sites for each site in the dashboard. This functionality is available starting with Cisco ACI Multi-Site, Release 1.0(2).

    When one or more BGP sessions or tunnels fail to establish, ACI Multi-Site provides the information about which exact local spines and remote spines failed to establish the BGP session or the tunnel. ACI Multi-Site should be enabled in the site in the infrastructure configuration, for the BGP sessions and the dataplane unicast and multicast tunnels to be established to the peer sites.

    BGP Sessions

    - When the BGP peering type is full-mesh in **Infra**-> **General Settings**, the spine node in a site with the BGP peering enabled will establish the BGP sessions to all the spine nodes with the BGP peering enabled in all the peer sites.
    - When the BGP peering type is route-reflector in **Infra**-> **General Settings**, the spine node in a site with both BGP peering enabled and route-reflector enabled, will establish the BGP sessions to all the spine nodes with the BGP peering enabled in all the peer sites. In the route-reflector mode, at least the local spine node or the remote spine node or both should have the route-reflector enabled. Otherwise, the BGP session is not established between them.

    - If the local and the remote ASNs are different, then it is eBGP. Therefore, the sessions between those sites are always full-mesh, irrespective of the BGP peering type and the route-reflector configuration.

  Unicast and Multicast Tunnels: A spine node in a site that is connected to ISN and has infrastructure configuration, will establish a tunnel to all the spine nodes that are connected to ISN in the peer sites.

The color codes indicate the following conditions:

- **Critical** (red)
- **Major** (orange)
- **Minor** (yellow)
- **Warning** (green)

The numbers in the color indicator columns indicate the number of faults per site.

- **+ Add Site:** enables you to add another site to our implementation. When you click **+ Add Site**, you must provide the following site details information on the **Connection Settings** page:
  - **Name**: the name of the site
  - **Labels**: the label identifier of the site. Multiple labels can be associated to a site.
  - **APIC Controller URL**: you can add more APIC controllers with a distinguishing URL of a cluster.
  - **Username** and **Password**: APIC login info with admin level privileges.
  - **Specify Domain For Site**: click the switch to on and provide the domain name if default authentication domain is configured in APIC.

    After you have entered your details for your new site, click the **Save** button.

- **Schema Health**: provides a listing of your schemas with locales and health.
  - Click the magnifying glass icon and enter a schema name to search for a subject schema.
  - Click **+ Add Schema** to start the procedure for adding a new schema to your site.
  - Click the site locale in the **Schema Health** table to view the schema details and status for a template.

    The **Schema Health** table provides a heat map type of display; that is, the health of the subject schema is displayed according to color. Schemas that span two columns (i.e, locales) indicate a stretched condition.

    - Click the color highlighted table cell to further discover what policies are incorporated into the subject schema. On the schema details page, you can click the arrow to go into the schema builder and update the policy details in the subject schema.
    - The color coded slider enables you to select a range for identifying schemas whose health require further review. For example, you can adjust the slider value to between 80 and 100. Then all of your schema implementations that fall within that specific range are displayed on the accompanying Schema Health table.

Sites Page

The Multi-Site **Sites** page displays all of the sites in your implementation. An example of the **Sites** page is shown in the following screen shot:

Figure 2: Multi-Site Sites Page

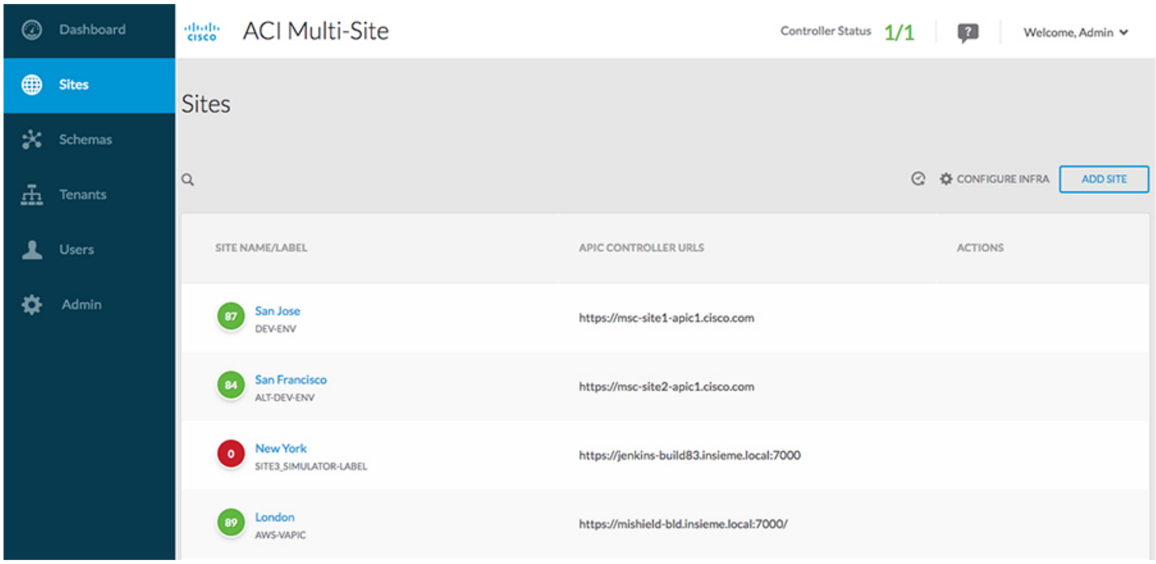

SITE NAME/LABEL	APIC CONTROLLER URLS	ACTIONS
87 San Jose DEV-ENV	https://msc-site1-apic1.cisco.com	
84 San Francisco ALT-DEV-ENV	https://msc-site2-apic1.cisco.com	
0 New York SITE3_SIMULATOR-LABEL	https://jenkins-build83.insieme.local:7000	
89 London AWS-VAPIC	https://mishield-bld.insieme.local:7000/	

The **Sites** page consists of the following two panes:

- **Site Name or Label**: the site status table lists your sites and then indicates the current health status for your implementation according to the following color coded identifiers:
  - **Critical** (red)
  - **Major** (orange)
  - **Minor** (yellow)
  - **Warning** (green)

  When you click a specific site, you can view or edit the site's details on the **Connection Settings** display:
  - **Name**
  - **Labels**
  - **APIC Controller URL**
  - **Username** and **Password**
  - **Specify Domain For Site**
  - **APIC Site ID**

    If you have made changes to the listed fields, click the **Save** button.

- **APIC Controller URLs**: the associated APIC URLs for your Multi-Site implementation

- **Configure Infra**: click this area to configure your Fabric infrastructure connectivity. For more information, refer to the Cisco Application Policy Infrastructure Controller (APIC) page.
- **Add Site**: click the **Add Site** button to add a site to your implementation. The following details are required for adding a site:
    - **Name:** the site name.
    - **Label:** select an existing or create a new label.
    - **APIC Controller URL**: the existing URL - click **+** to add a new APIC Controller URL.
    - **Username**: the site username.
    - **Password**: the unique site password for access.
    - **Specify Domain for Site**: click the selector to **On** to specify a domain for the site.
- **Actions**: drop down menu list option to edit, delete, or open a subject site in the APIC user interface.

Audit Logs

Click the **Audit Log** icon next to the **Configure Infra** tab to list the log details for the Sites page. The **Audit Logs: Sites List** page is displayed.

The table on the page displays the following details:

- **Date**
- **Action**
- **Details**
- **User**

Click the **Most Recent** tab to select the audit logs during a particular time period. For example, when you select the range from November 14, 2017 to November 17, 2017 and click **Apply**, the audit log details for this time period are displayed on the **Audit Logs** page.

Click the **Filter** icon next to the **Most Recent** tab to filter the log details using the following criteria:

- User: Select one user name or all users and click **Apply** to filter the log details using the user name.
- Action: Select the action, for example, created, updated, or deleted, and click **Apply** to filter the log details according to the action.

Schemas Page

The **Multi-Site Schemas** page lists all schemas that are associated with your implementation.

The following screen shot shows an example display:

Figure 3: Multi-Site Schemas Page

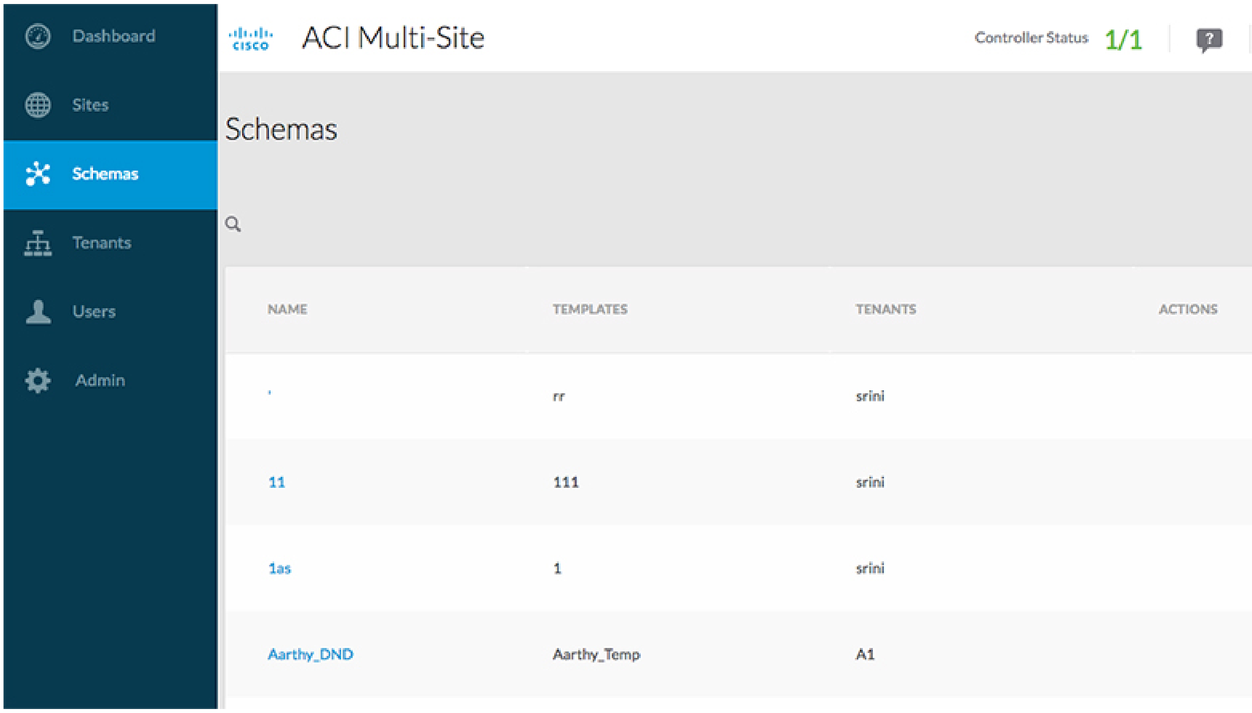

NAME	TEMPLATES	TENANTS	ACTIONS
'	rr	srini	
11	111	srini	
1as	1	srini	
Aarthy_DND	Aarthy_Temp	A1	

Use the magnifying glass and associated field to search for a specific schema. Use schemas to configure or import tenant policies, including the VRF, application profile with EPGs, filters and contracts, bridge domains, and external EPGs.

The Schemas table shows the following information in tabular form:

- **Name**: click the schema name to view or update the settings for the subject schema.
- **Templates**: displays the name of the template that is used for the schema. Templates are analogous to profiles in the ACI context, which group policies. You can create templates for stretched objects or site-specific objects.
- **Tenants**: displays the name of the tenant that is used for the subject schema.
- **Actions**: click the **Action** field with the associated schema to either edit or delete the subject schema.

Click the **Add Schema** button to add a new schema to your implementation. Further details on creating a schema are described in Schema Management, on page 61.

Audit Logs

Click the **Audit Log** icon next to the **Add Schema** tab to list the log details for the Schemas page. The **Audit Logs: Schemas List** page is displayed.

The table on the page displays the following details:

- **Date**
- **Action**
- **Details**

- **User**

Click the **Most Recent** tab to select the audit logs during a particular time period. For example, when you select the range from November 10, 2017 to November 14, 2017 and click **Apply**, the audit log details for this time period are displayed on the **Audit Logs** page.

Click the **Filter** icon next to the **Most Recent** tab to filter the log details using the following criteria:

- **User**: Select one username or all users and click **Apply** to filter the log details using the username.
- **Action**: Select the action, for example, created, updated, or deleted, and click **Apply** to filter the log details according to the action.

For more information about creating schemas, refer to Schema Management, on page 61.

Tenants Page

The Multi-Site **Tenants** page lists all of the tenants that comprise your implementation.

The following screen shot provides an example:

Figure 4: Multi-Site Tenants Page

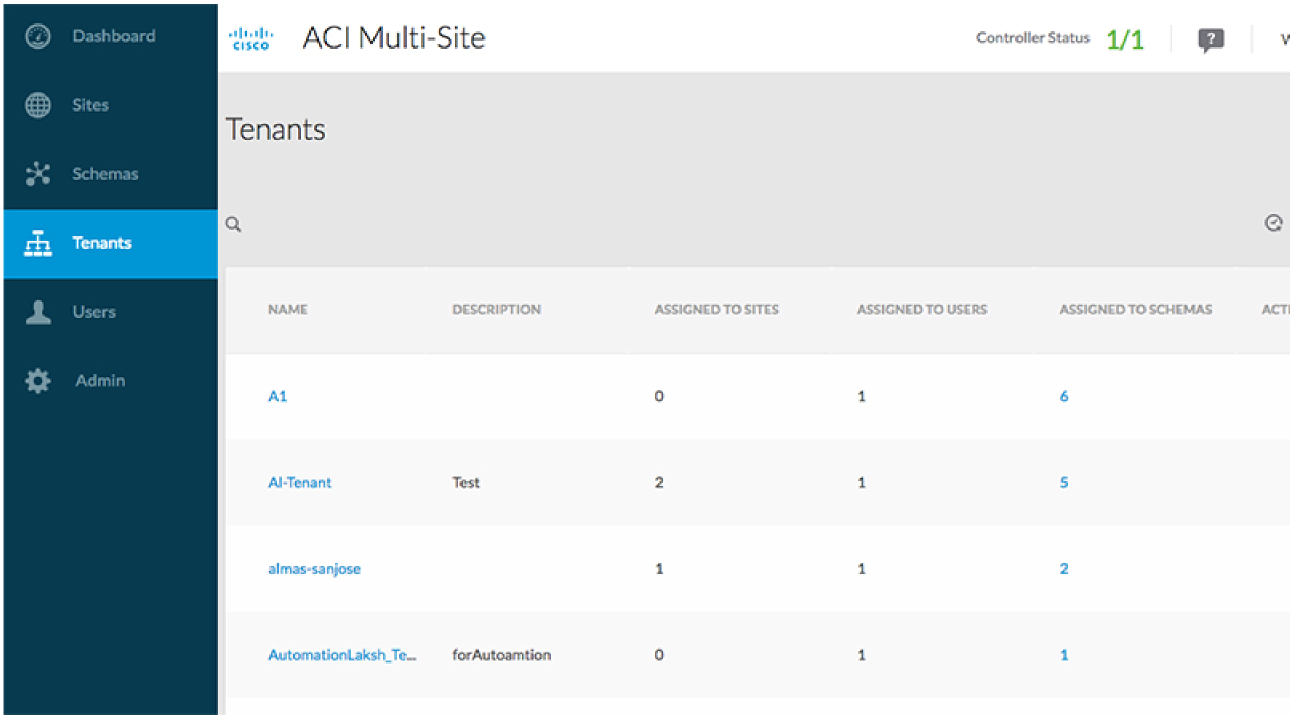

Tenants

NAME	DESCRIPTION	ASSIGNED TO SITES	ASSIGNED TO USERS	ASSIGNED TO SCHEMAS	ACTI
A1		0	1	6	
Al-Tenant	Test	2	1	5	
almas-sanjose		1	1	2	
AutomationLaksh_Te...	forAutoamtion	0	1	1	

The table on the **Tenants** page displays the following:

- **Tenant Name**
- **Assigned to Sites**
- **Assigned to Users**

- **Assigned to Schemas**
- **Actions**

The features and functionality on this page include the following:

- **Name**: click a tenant name to access the **Tenant Details** settings page. On the **Tenant Details** page you can edit or update the following sections:
  - **General Settings**: change the Display Name and Description as required.
  - **Associated Sites**: view the sites associated with the subject tenant.
  - **Associated Users**: view the users associated with the subject tenant - you can associate a user with the subject tenant by checking the empty box next to the user name.
- **Associated Schemas**: click the **Associated Schema** listing to view the schemas associated with the subject tenant.
- **Actions**: click the **Actions** listing to edit the subject tenant's details sites or to create a new network mapping.

  Note You can delete the Tenant object by selecting **Delete** on the **Actions** drop down menu.
- **Add Tenant:** click **Add Tenant** button to add an existing tenant to your implementation. On the proceeding Tenant Details page, you can add the tenant name, description, security domain, and associated users.

Audit Logs

Click the **Audit Log** icon next to the **Add Tenant** tab to list the log details for the Tenants page. The **Audit Logs: Tenants List** page is displayed.

The table on the page displays the following details:

- **Date**
- **Action**
- **Details**
- **User**

Click the **Most Recent** tab to select the audit logs during a particular time period. For example, when you select the range from November 10, 2017 to November 14, 2017 and click **Apply**, the audit log details for this time period are displayed on the **Audit Logs** page.

Click the **Filter** icon next to the **Most Recent** tab to filter the log details using the following criteria:

- User: Select one user name or all users and click **Apply** to filter the log details using the user name.
- Action: Select the action, for example, created, updated, or deleted, and click **Apply** to filter the log details according to the action.

For more information about adding tenants, refer to Tenant Management, on page 49.

Users Page

The Multi-Site Orchestrator **Users** page displays all of the users. An example of the **Users** page is as follows:

Figure 5: Users Page

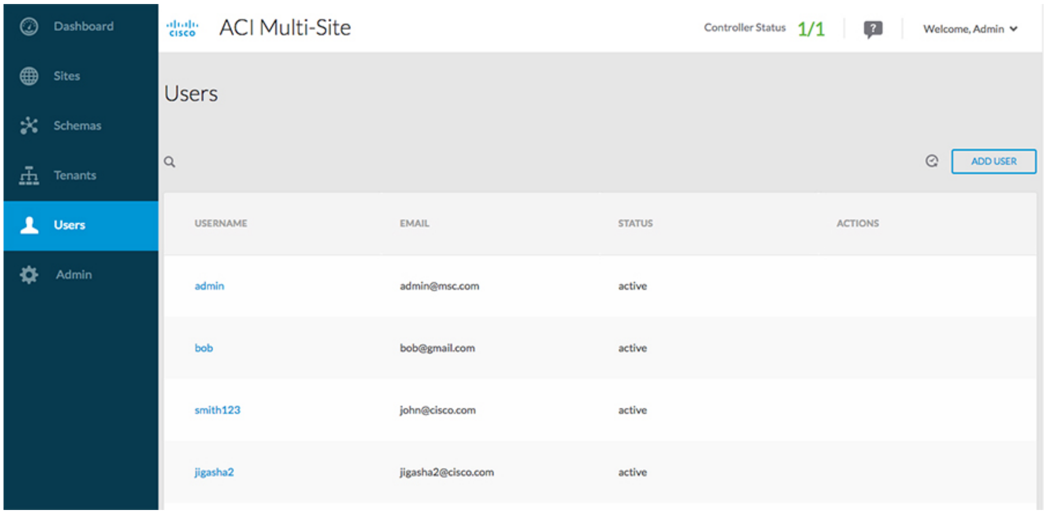

USERNAME	EMAIL	STATUS	ACTIONS
admin	admin@msc.com	active	
bob	bob@gmail.com	active	
smith123	john@cisco.com	active	
jigasha2	jigasha2@cisco.com	active	

The **Users** page features a table containing all of the identified users by username and associated email and current activity status. If you click a selected **Username,** you can access the **General Setting** page attributable to the subject user. On the **General Setting** page, you can edit the details associated with the subject user such as username, password, email, and switch-on user roles.

Click **Add User** to add a new user to your Multi-Site implementation. The **General Setting** page display enables you to assign username, password, email, and switch-on user roles associated with your Multi-Site implementation.

Audit Logs

Click the **Audit Log** icon next to the **Add User** tab to list the log details for the Users page. The **Audit Logs: Users List** page is displayed.

The table on the page displays the following details:

- **Date**
- **Action**
- **Details**
- **User**

Click the **Most Recent** tab to select the audit logs during a particular time period. For example, when you select the range from November 10, 2017 to November 14, 2017 and click **Apply**, the audit log details for this time period are displayed on the **Audit Logs** page.

Click the **Filter** icon next to the **Most Recent** tab to filter the log details using the following criteria:

- User: Select one user name or all users and click **Apply** to filter the log details using the user name.
- Action: Select the action, for example, created, updated, or deleted, and click **Apply** to filter the log details according to the action.

Policies Page

The Multi-Site Orchestrator **Policies** page displays all policies you have configured for your fabrics.

The **Policies** page contains a table of all policies along with the summary of their types, tenants they're associated with, descriptions, and usage. You can use this page to add new policies or edit existing ones.

Detailed information about available policy types and how to configure them is available in the Administrative Operations, on page 87 chapter.

Admin Page

When you select the Admin tab from the Cisco ACI Multi-Site Orchestrator navigation bar, it expands the following additional selection of administrative pages:

- **Providers**
- **Login Domains**
- **Backups**
- **Audit Logs**
- **Security**
- **Remote Locations**
- **System Configuration**

Providers

Figure 6: Cisco ACI Multi-Site Orchestrator Providers Page

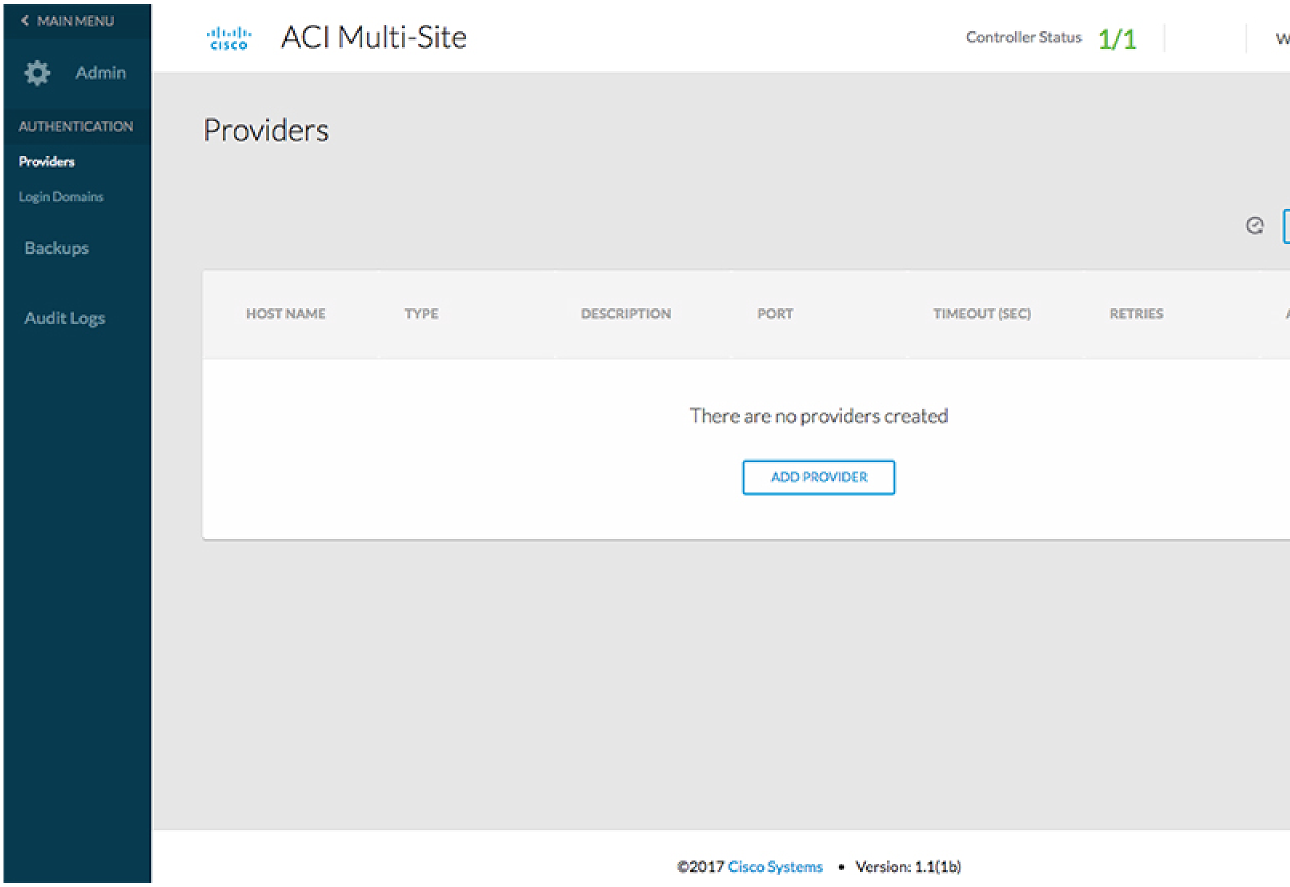

The **Providers** page under the **Admin** heading displays information about any configured external authentication providers. The following details are shown for each provider:

- **Host Name**
- **Type**
- **Description**
- **Port**
- **Timeout (Sec)**
- **Retries**

Working with external authentication providers is described in Administrative Operations, on page 87.

Login Domains

Figure 7: Cisco ACI Multi-Site Orchestrator Login Domains Page

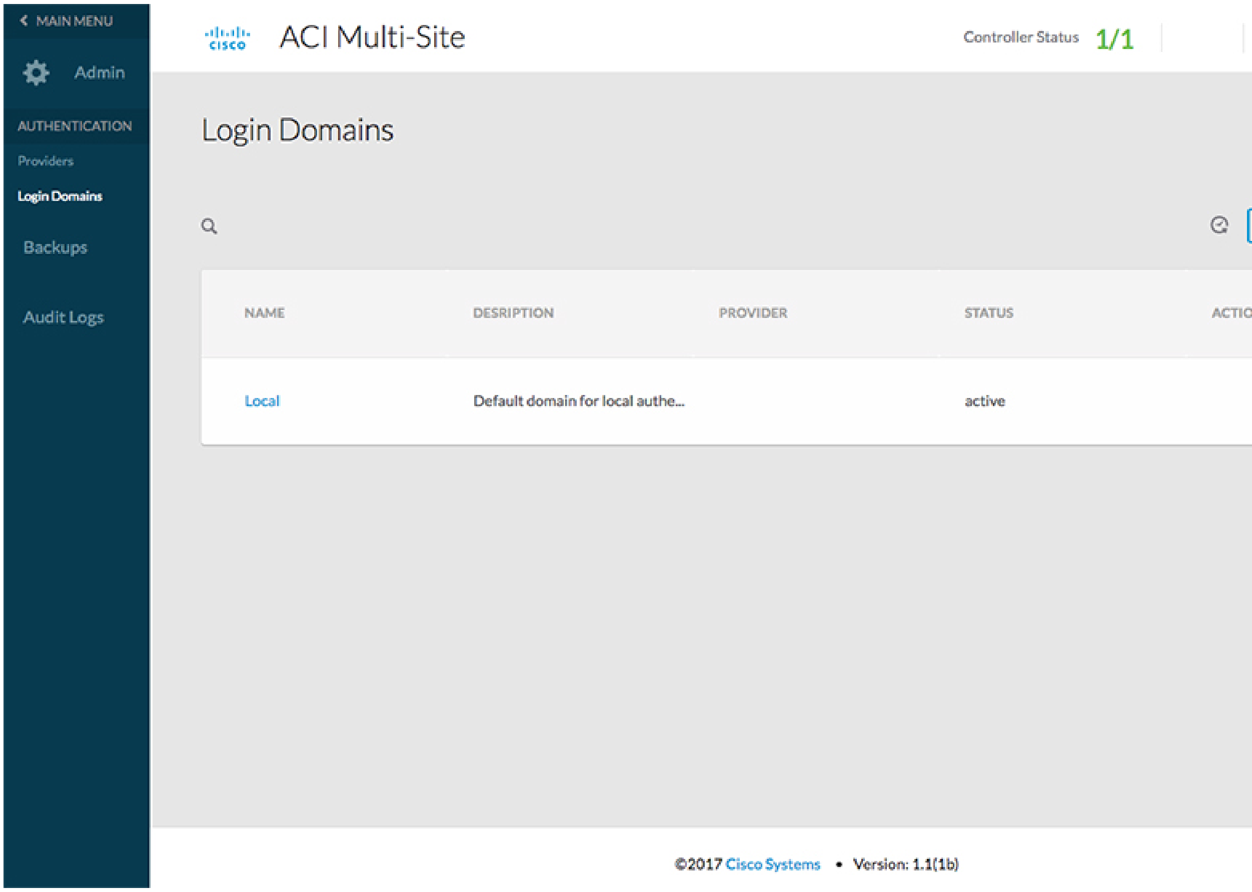

Login Domains

NAME	DESRIPTION	PROVIDER	STATUS	ACTIO
Local	Default domain for local authe...		active	

©2017 Cisco Systems • Version: 1.1(1b)

The **Login Domains** page under the **Admin** heading displays information about the available login domains. The following details are shown for each domain:

- **Name**
- **Description**
- **Provider**
- **Status**
- **Default**

Working with login domains is described in Administrative Operations, on page 87.

Backups

Figure 8: Cisco ACI Multi-Site Orchestrator Backups Page

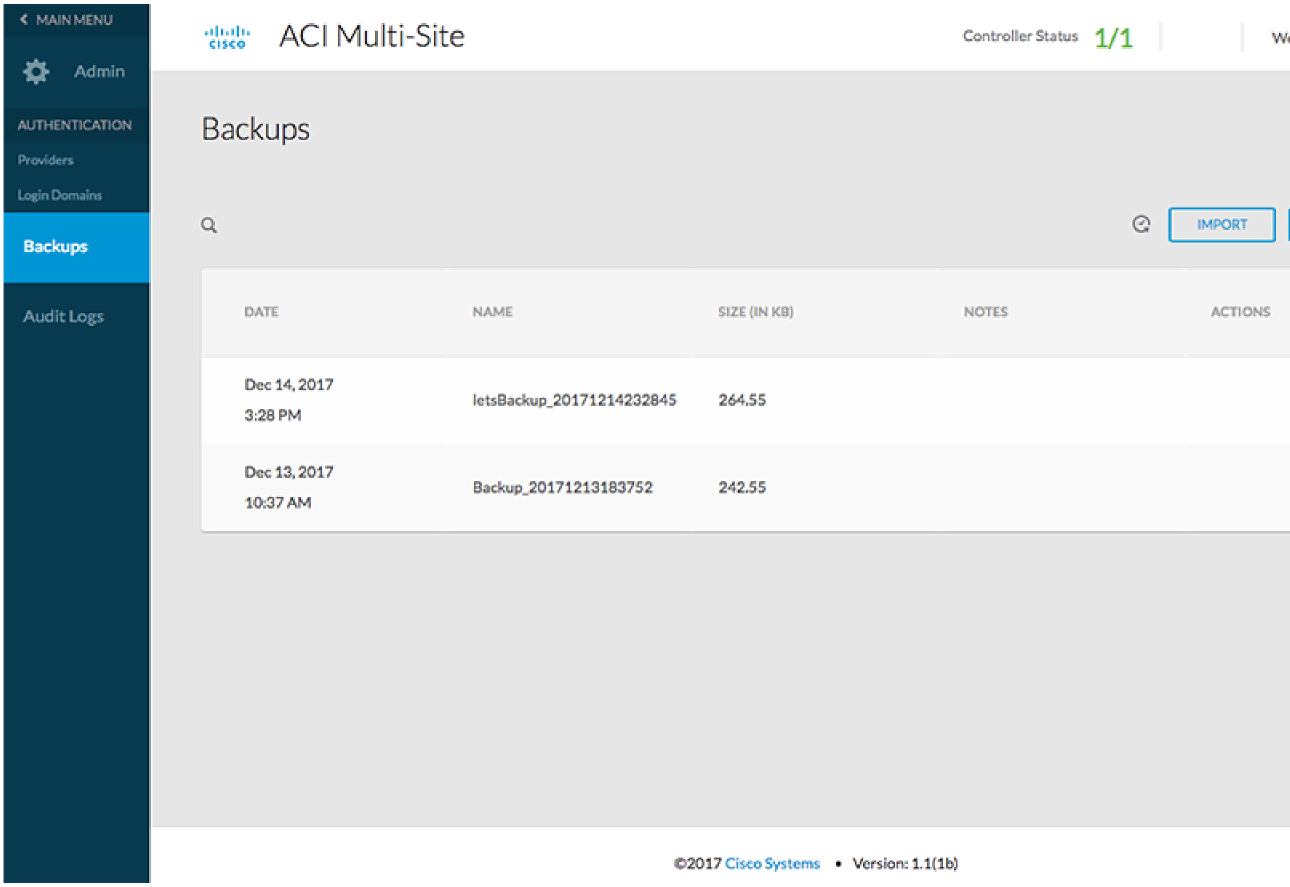

DATE	NAME	SIZE (IN KB)	NOTES	ACTIONS
Dec 14, 2017 3:28 PM	letsBackup_20171214232845	264.55		
Dec 13, 2017 10:37 AM	Backup_20171213183752	242.55		

©2017 Cisco Systems • Version: 1.1(1b)

The **Backups** page under the **Admin** heading displays information about any backups that have been created. The following details are shown for each domain:

- **Date**
- **Name**
- **Size**
- **Notes**

Working with backups is described in Administrative Operations, on page 87.

Figure 9: Cisco ACI Multi-Site Orchestrator Audit Logs Page

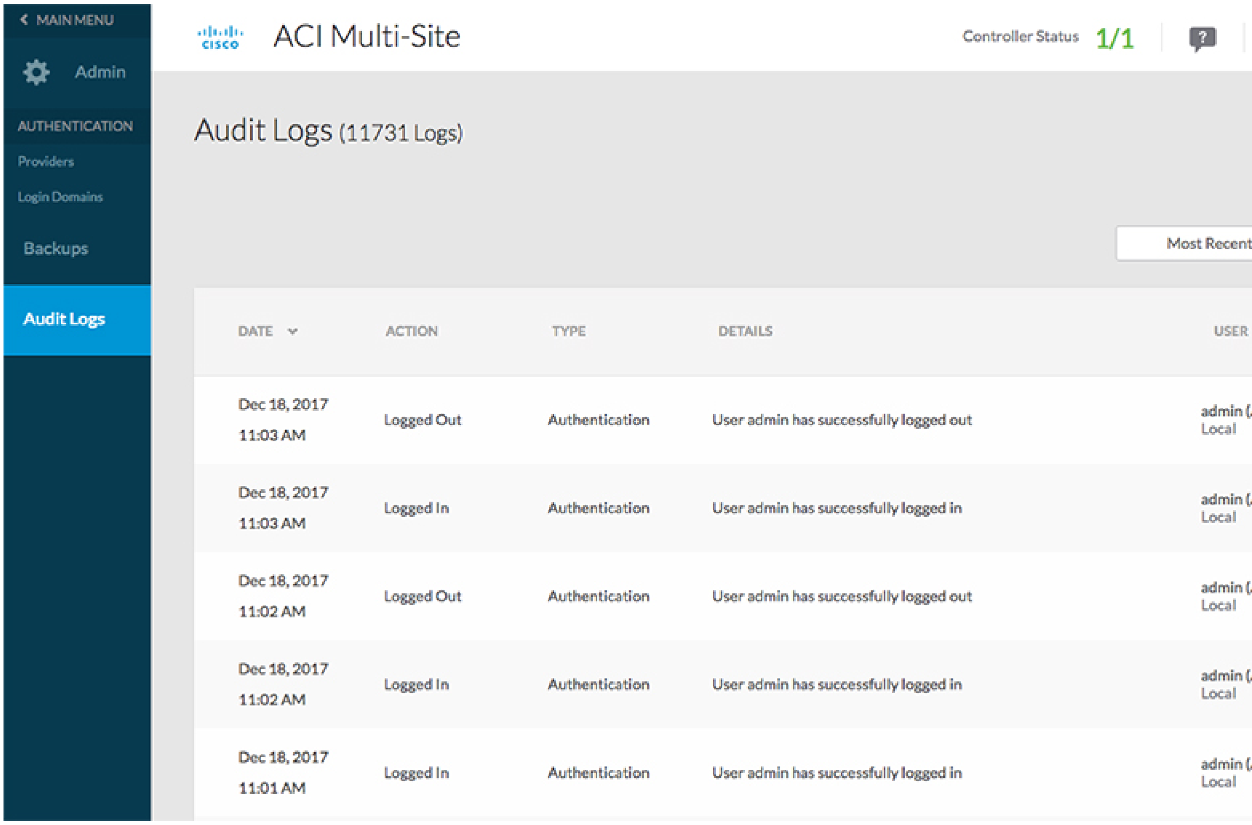

DATE	ACTION	TYPE	DETAILS	USER
Dec 18, 2017 11:03 AM	Logged Out	Authentication	User admin has successfully logged out	admin (Local
Dec 18, 2017 11:03 AM	Logged In	Authentication	User admin has successfully logged in	admin (Local
Dec 18, 2017 11:02 AM	Logged Out	Authentication	User admin has successfully logged out	admin (Local
Dec 18, 2017 11:02 AM	Logged In	Authentication	User admin has successfully logged in	admin (Local
Dec 18, 2017 11:01 AM	Logged In	Authentication	User admin has successfully logged in	admin (Local

Audit Logs

The **Audit Logs** page under the **Admin** heading displays information about the audit logs and records. The following details are shown:

- **Date**
- **Action**
- **Type**
- **Details**
- **User**

Working with logs is described in Administrative Operations, on page 87.

Security

Figure 10: Cisco ACI Multi-Site Orchestrator Security Page

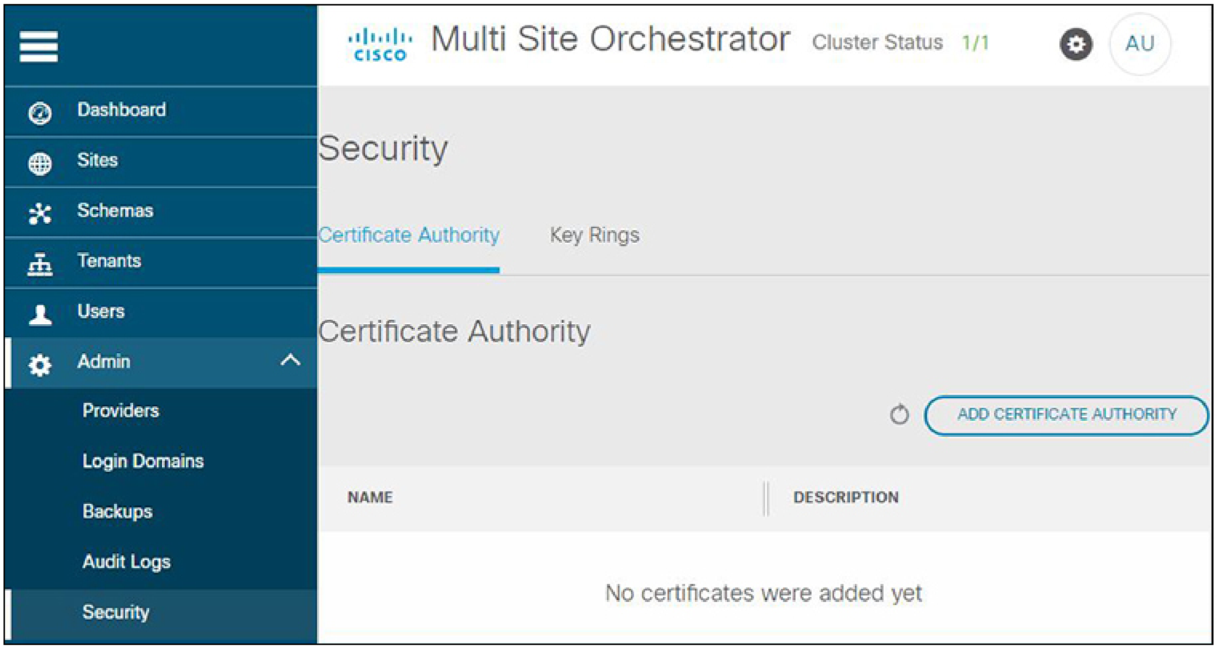

The **Security** page under the **Admin** heading displays information about the custom certificates and key rings you have configured for use by the Orchestrator. The following details are shown:

- **Certificate Authority**
  - **Name**
  - **Description**
- **Key Rings**
  - **Name**
  - **Description**
  - **Trustpoint**
  - **State**

Working with certificates is described in Administrative Operations, on page 87.

Remote Locations

The **Remote Locations** page under the **Admin** heading displays information about any remote backup locations you have configured for use by the Orchestrator. The following details are shown:

- **Name**
- **Host**
- **Protocol**

- **Username**
- **Remote Path**

Working remote backups is described in Administrative Operations, on page 87.

System Configuration

The **System Configuration** page under the **Admin** heading allows you to configure a number of system settings that define how the Orchestrator GUI behaves. For example, you can change how failed login attempts are treated or if a warning banner should be displayed at the top of the GUI.

The available system settings are described in more detail in Administrative Operations, on page 87.

CHAPTER 3

Infrastructure Management

- Cisco ACI Multi-Site and Cisco APIC Interoperability Support, on page 21
- Multi-Site Orchestrator Communication Ports, on page 22
- Configuring Fabric Access Policies for All APIC Sites, on page 23
- Configuring Sites That Contain Remote Leaf Switches, on page 26
- Adding Sites, on page 27
- Configuring Infra Prerequisites and Guidelines, on page 28
- Deleting Sites Using Multi-Site Orchestrator GUI, on page 33
- Cisco ACI CloudSec Encryption, on page 33
- Multi-Site Cross Launch to Cisco APIC, on page 45

Cisco ACI Multi-Site and Cisco APIC Interoperability Support

Prior to Release 2.2(1), you were required to run the same APIC versions in all sites and the version of the Orchestrator that corresponded to that APIC release. During fabric upgrade you were also required to upgrade all the APIC sites first before upgrading the Multi-Site Orchestrator. For example, if you were upgrading the fabrics from APIC Release 4.0(1) to Release 4.1(1), you had to remain on Release 2.0(1) of the Orchestrator until all sites were on APIC Release 4.1(1).

Starting with Release 2.2(1), Multi-Site Orchestrator releases have been decoupled from the APIC releases. The APIC clusters in each site as well as the Orchestrator itself can now be upgraded independently of each other and run in mixed operation mode.

Mixed operation mode is supported for sites running any of the following APIC releases:

- 3.2(6) or later
- 4.0(1) or later
- 4.1(1) or later
- 4.2(1) or later

However, keep in mind that if you upgrade the Orchestrator before upgrading the APIC clusters in one or more sites, the new Orchestrator features may not yet be supported by an earlier APIC release. In that case a check is performed on each template to ensure that every configured option is supported by the target sites. The check is performed when you save a template or deploy a template. If the template is already assigned to a site, any unsupported configuration options will not be saved; if the template is not yet assigned, you will be able to assign it to a site, but not be able to save or deploy the schema if it contains configuration unsupported by the site. In case an unsupported configuration is detected, an error message will show, for example: `This APIC site version <site-version> is not supported by MSO. The minimum version required for this <feature> is <required-version> or above.`

The following table lists the features and the minimum required APIC release for each one:

Feature	Minimum APIC Version
ACI Multi-Pod Support	Release 3.2(6)
Service Graphs (L4-L7 Services)	Release 3.2(6)
External EPGs	Release 3.2(6)
ACI Virtual Edge VMM Support	Release 3.2(6)
DHCP Support	Release 3.2(6)
Consistency Checker	Release 3.2(6)
CloudSec Encryption	Release 4.0(1)
Layer 3 Multicast	Release 4.0(1)
MD5 Authentication for OSPF	Release 4.0(1)
EPG Preferred Group	Release 4.0(2)
Host Based Routing	Release 4.1(1)
Intersite L3Out	Release 4.2(1)

Multi-Site Orchestrator Communication Ports

There are three types of network communication to or from the Multi-Site Orchestrator cluster:

- Client traffic to the Multi-Site Orchestrator cluster.

  Multi-Site Orchestrator uses TCP port 433 (`https`) to allow user access via GUI or REST API for creating, managing, and deploying policy configurations.

- REST API traffic from the Multi-Site Orchestrator to the APIC controllers of the ACI fabrics that are part of the Multi-Site domain

  Multi-Site Orchestrator uses TCP port 433 for REST API traffic to deploy policies to each site.

- Intra-cluster communication.

  All control-plane and data-plane traffic between Cisco ACI Multi-Site Orchestrator nodes (including intra-cluster communication and container overlay network traffic) is encrypted with IPSec's Encapsulating Security Payload (ESP) using IP protocol number 50 to provide security and allow the cluster deployments over a round-trip time distance of up to 150ms. If there is firewall between any Orchestrator nodes, proper rules must be added to allow this traffic.

  If your Multi-Site Orchestrator cluster is deployed directly in VMware ESX without the Application Services Engine, the following ports are used for Docker communications between the cluster nodes:

Note The following TCP and UDP ports are listed for educational perspective only as no traffic is ever sent in clear text across the network leveraging these ports.

- TCP port 2377 for Cluster Management Communications
- TCP and UDP port 7946 for Inter-Manager Communication
- UDP port 4789 for Docker Overlay Network Traffic

Configuring Fabric Access Policies for All APIC Sites

Before your APIC fabrics can be added to and managed by the Multi-Site Orchestrator, there is a number of fabric-specific access policies that you must configure on each site.

Configuring Fabric Access Global Policies

This section describes the global fabric access policy configurations that must be created for each APIC site before it can be added to and managed by the Multi-Site Orchestrator.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

You must configure a number of fabric policies before the site can be added to the Multi-Site Orchestrator. From the APIC's perspective, this is something you do just like you would if you were connecting a bare-metal host, where you would configure domains, AEPs, policy groups, and interface selectors; you must configure the same options for connecting the spine switch interfaces to the inter-site network for all the sites that will be part of the same Multi-Site domain.

Step 3 Specify the VLAN pool.

The first thing you configure is the VLAN pool. We use Layer 3 sub-interfaces tagging traffic with VLAN-4 to connect the spine switches to the inter-site network.

a) In the left navigation tree, browse to **Pools** > **VLAN**.
b) Right-click the **VLAN** category and choose **Create VLAN Pool**.

In the **Create VLAN Pool** window, specify the following:

- For the **Name** field, specify the name for the VLAN pool, for example `msite`.
- For **Allocation Mode**, specify `Static Allocation`.
- And for the **Encap Blocks**, specify just the single VLAN 4. You can specify a single VLAN by entering the same number in both **Range** fields.

Step 4 Configure Attachable Access Entity Profiles (AEP).

a) In the left navigation tree, browse to **Global Policies** > **Attachable Access Entity Profiles**.
b) Right-click the **Attachable Access Entity Profiles** category and choose **Create Attachable Access Entity Profiles**.

In the **Create Attachable Access Entity Profiles** window, specify the name for the AEP, for example `msite-aep`.

c) Click **Next** and **Submit**

No additional changes, such as interfaces, are required.

Step 5 Configure domain.

The domain you configure is what you will select from the Multi-Site Orchestrator when adding this site.

a) In the left navigation tree, browse to **Physical and External Domains** > **External Routed Domains**.
b) Right-click the **External Routed Domains** category and choose **Create Layer 3 Domain**.

In the **Create Layer 3 Domain** window, specify the following:

- For the **Name** field, specify the name the domain, for example `msite-l3`.
- For **Associated Attachable Entity Profile**, select the AEP you created in Step 4.
- For the **VLAN Pool**, select the VLAN pool you created in Step 3.

c) Click **Submit**.

No additional changes, such as security domains, are required.

What to do next

After you have configured the global access policies, you must still add interfaces policies as described in Configure Fabric Access Interface Policies on Each APIC, on page 24.

Configure Fabric Access Interface Policies on Each APIC

This section describes the fabric access interface configurations that must be done for the Multi-Site Orchestrator on each APIC site.

Before you begin

Configure the global fabric access policies, such as VLAN Pool, AEP, and domain, in the site's APIC, as described in Configuring Fabric Access Global Policies, on page 23.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the main navigation menu, select **Fabric** > **Access Policies**.

In addition to the VLAN, AEP, and domain you have configured in previous section, you must also create the interface policies for the fabric's spine switch interfaces that connect to the Inter-Site Network (ISN).

Step 3 Configure a spine policy group.

a) In the left navigation tree, browse to **Interface Policies** > **Policy Groups** > **Spine Policy Groups**.

This is similar to how you would add a bare-metal server, except instead of a Leaf Policy Group, you are creating a Spine Policy Group.

b) Right-click the **Spine Policy Groups** category and choose **Create Spine Access Port Policy Group**.

In the **Create Spine Access Port Policy Group** window, specify the following:

- For the **Name** field, specify the name for the policy group, for example `Spine1-PolGrp`.
- For the **Link Level Policy** field, specify the link policy used between your spine switch and the ISN.
- For **CDP Policy**, choose whether you want to enable CDP.
- For the **Attached Entity Profile**, select the AEP you have configured in previous section, for example `msite-aep`.

Then click **Submit**. No additional changes, such as security domains, are required.

Step 4 Configure a spine profile.

a) In the left navigation tree, browse to **Interface Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Interface Profile**.

In the **Create Spine Interface Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1-ISN`.
- For **Interface Selectors**, click the **+** sign to add the port on the spine switch that connects to the ISN. Then in the **Create Spine Access Port Selector** window, provide the following:
  - For the **Name** field, specify the name for the port selector, for example `Spine1-ISN`.
  - For the **Interface IDs**, specify the switch port that connects to the ISN, for example `5/32`.
  - For the **Interface Policy Group**, choose the policy group you created in the previous step, for example `Spine1-PolGrp`.

  Then click **OK** to save the port selector.

Then click **Submit** to save the spine interface profile.

Step 5 Configure a spine switch selector policy.

a) In the left navigation tree, browse to **Switch Policies** > **Profiles** > **Spine Profiles**.

b) Right-click the **Spine Profiles** category and choose **Create Spine Profile**.

In the **Create Spine Profile** window, specify the following:

- For the **Name** field, specify the name for the profile, for example `Spine1`.
- For **Spine Selectors**, click the **+**to add the spine and provide the following:
  - For the **Name** field, specify the name for the selector, for example `Spine1`.
  - For the **Blocks** field, specify the spine node, for example `201`.

  Then click **Update** to save the selector.

Then click **Next** and on the next screen select the interface profile you have created in the previous step, for example `Spine1-ISN`.

Finally, click **Finish** to save the spine profile.

What to do next

If your fabrics contain Remote Leaf switches, you will need to make additional fabric-specific configuration changes as described in Configuring Sites That Contain Remote Leaf Switches, on page 26

Configuring Sites That Contain Remote Leaf Switches

Starting with Release 2.1(2), the Multi-Site architecture supports APIC sites with Remote Leaf switches. The following sections describe guidelines, limitations, and configuration steps required to allow Multi-Site Orchestrator to manage these sites.

Multi-Site and Remote Leaf Guidelines and Limitations

If you want to add an APIC site with a Remote Leaf to be managed by the Multi-Site Orchestrator, the following restrictions apply:

- You must upgrade your Cisco APIC to Release 4.1(2) or later.
- You must upgrade your Multi-Site Orchestrator to Release 2.1(2) or later.
- Only physical Remote Leaf switches are supported in this release
- Only -EX and -FX or later switches are supported as Remote Leaf switches for use with Multi-Site:
- Remote Leaf is not supported with back-to-back connected sites without IPN switches
- Remote Leaf switches in one site cannot use another site's L3out
- Stretching a bridge domain between one site and a Remote Leaf in another site is not supported

You must also perform the following tasks before the site can be added to and managed by the Multi-Site Orchestrator:

- You must enable Remote Leaf direct communication and configure routable subnets directly in the site's APIC, as described in the following sections.
- You must add the routable IP addresses of Cisco APIC nodes in the DHCP-Relay configuration applied on the interfaces of the Layer 3 routers connecting to the Remote Leaf switches.

  The routable IP address of each APIC node is listed in the **Routable IP** field of the **System** > **Controllers** > **<controller-name>** screen of the APIC GUI.

Configuring Routable Subnets for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Multi-Site Orchestrator, you must configure routable subnets for the pod with which the Remote Leaf nodes are associated.

Step 1 Log in directly to the site's APIC GUI.

Step 2 From the menu bar, select **Fabric** > **Inventory**.

Step 3 In the Navigation pane, click **Pod Fabric Setup Policy**.

Step 4 In the main pane, double-click the pod where you want to configure the subnets.

Step 5 In the **Routable Subnets** area, click the **+** sign to add a subnet.

Step 6 Enter the **IP** and **Reserve Address Count**, set the state to `Active` or `Inactive`, then click **Update** to save the subnet.

When configuring routable subnets, you must provide a netmask between `/22` and `/29`.

Step 7 Click **Submit** to save the configuration.

Enabling Direct Communication for Remote Leaf Switches

Before you can add a site that contains one or more Remote Leaf switches to the Multi-Site Orchestrator, you must configure direct remote leaf communication for that site. Additional information about remote leaf direct communication feature is available in the *Cisco APIC Layer 3 Networking Configuration Guide*. This section outlines the steps and guidelines specific to the integration with Multi-Site.

Note Once you enable Remote Leaf switch direct communication, the switches will function in the new mode only

Step 1 Log in directly to the site's APIC.

Step 2 Enable direct traffic forwarding for Remote Leaf switches.

a) From the menu bar, navigate to **System** > **System Settings**.
b) From the left side bar, select **Fabric Wide Setting**.
c) Check the **Enable Remote Leaf Direct Traffic Forwarding** checkbox.

Note You cannot disable this option after you enable it.

d) Click **Submit** to save the changes.

Adding Sites

This section describes how to add sites using the Cisco ACI Multi-Site Orchestrator GUI.

Before you begin

You must have completed the site-specific configurations in each site's APIC, as decribed in previous sections in this chapter.

Step 1 Log in to the Multi-Site GUI, in the **Main menu**, click **Sites**.

If you are logging in for the first time, log in as the **admin** user with the default password **We1come2msc!**, you will then be prompted to change that default password. The new password requirements are:

- At least 12 characters
- At least 1 letter

- At least 1 number
- At least 1 special character apart from * and space

Step 2 In the **Main menu**, select **Infrastructure** > **Sites**.

Step 3 In the top right of the main pane, click **Add Site**.

Step 4 In the **Add Site** screen, provide the site's details.

a) In the **Name** field, enter the site name.

b) In the **Labels** field, choose or create a label.

You can choose to provide multiple labels for the site.

c) In the **APIC Controller URL** field, enter the Cisco APIC URL.

For the APIC URL, you can use the `http` or `https` protocol and the IP address or the DNS hostname, such as`https://`*`<ip-address>`* or `https://`*`<dns-hostname>`*.

d) If you have a cluster of APICs in the fabric, click **+APIC Controller URL** and provide the additional URLs.

e) In the **Username** field, enter the admin user's username for the site's APIC.

f) In the **Password** field, enter the user's password.

g) You can turn on the **Specify Login Domain for Site** switch, if you want to specify a domain to be used for authenticating the user you provided.

If you turn on this option, enter the domain name in the **Domain Name** field.

h) In the **APIC Site ID** field, enter a unique site ID.

The site ID must be a unique identifier of the Cisco APIC site, ranged between `1` and `127`. Once specified, the site ID cannot be changed without factory resetting Cisco APIC.

Step 5 Click **Save** to add the site.

Step 6 If prompted, confirm proxy configuration update.

If you have configured the Orchestrator to use a proxy server and are adding an on-premises site that is not already part of the "no proxy" list, the Orchestrator will inform you of the proxy settings update.

For additional information on proxy configuration, see the "Administrative Operations" chapter in *Cisco ACI Multi-Site Configuration Guide*.

Step 7 Repeat these steps to add any additional sites.

Configuring Infra Prerequisites and Guidelines

The following sections describe the steps necessary to configure the general as well as site-specific fabric Infra settings.

Before you proceed with Infra configuration, you must have configured and added the sites as described in previous sections, which includes:

- Configuring each site's fabric access policies.
- Configuring direct communication and routable subnets for sites with remote leaf switches.

In addition, keep in mind the following:

- Any infrastructure changes such as adding and removing spine switches or spine node ID changes require a Multi-Site fabric connectivity information refresh described in the Refreshing Site Connectivity Information, on page 29 as part of the general Infra configuration procedures.
- The Overlay Unicast TEP, Overlay Multicast TEP, and BGP-EVPN Router-IDs IP addresses assigned on the Orchestrator should not be taken from the address space of the original fabric's `Infra` TEP pool or from the `0.x.x.x` range.

Configuring Infra: General Settings

This section describes how to configure general Infra settings for all the sites.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, click **Sites**.

Step 3 In the **Sites** view, click **Configure Infra**.

Step 4 In the left pane, under **Settings**, click **General Settings**.

Step 5 From the **BGP Peering Type** dropdown, choose either `full-mesh` or `route-reflector`.

The `route-reflector` option is effective only when all sites are part of the same BGP Autonomous System (AS).

Step 6 In the **Keepalive Interval (Seconds)** field, enter the keep alive interval seconds.

We recommend keeping the default value.

Step 7 In the **Hold Interval (Seconds)** field, enter the hold interval seconds.

We recommend keeping the default value.

Step 8 In the **Stale Interval (Seconds)** field, enter stale interval seconds.

We recommend keeping the default value.

Step 9 Choose whether you want to turn on the **Graceful Helper** option.

Step 10 In the **Maximum AS Limit** field, enter the maximum AS limit.

Step 11 In the **BGP TTL Between Peers** field, enter the BGP TTL between peers.

Refreshing Site Connectivity Information

Any infrastructure changes, such as adding and removing spines or changing spine node IDs, require a Multi-Site fabric connectivity site refresh. This section describes how to pull up-to-date connectivity information directly from each site's APIC.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, select **Infrastructure** > **Infra Configuration**.

Step 3 In the top right of the main **Infra Configuration** view, click the **Configure Infra** button.

Step 4 In the left pane, under **Sites**, select a specific site.

Step 5 In the main window, click the **Reload Site Data** button to pull fabric information from the APIC.

Step 6 (Optional) In the **Confirmation** dialog, check the box if you want to remove configuration for decommissioned spine switch nodes.

If you choose to enable this checkbox, all configuration info for any currently decommissioned spine switches will be removed from the database.

Step 7 Finally, click **Yes** to confirm and load the connectivity information.

This will discover any new or removed spines and all site-related fabric connectivity will be re-imported from the APIC.

Configuring Infra Site-Specific Settings

This section describes how to configure site-specific Infra settings for each site.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, click **Sites**.

Step 3 In the **Sites** view, click **Configure Infra**.

Step 4 In the left pane, under **Sites**, select a specific site.

Step 5 In the right **<*Site*> Settings** pane, enable the site by setting the **ACI Multi-Site** knob to `On`.

Step 6 (Optional) Turn on CloudSec encryption for the site.

CloudSec Encryption provides inter-site traffic encryption. The "Infrastructure Management" chapter in the *Cisco ACI Multi-Site Configuration Guide* covers this feature in detail.

Step 7 Specify the **Overlay Multicast TEP**.

This address is used for the inter-site L2 BUM and L3 multicast traffic. This IP address is deployed on all spine switches that are part of the same fabric, regardless of whether it is a single Pod or Multi-Pod fabric.

Step 8 Specify the **BGP Autonomous System Number**.

Step 9 Specify the **BGP Password**.

Step 10 Specify the **OSPF Area ID**.

When configuring the Multi-Site infra OSPF details, we recommend that you use OSPF Area `0`. If you use an Area ID other than `0`, in the next step configure it as a `regular` OSPF area type and not a `stub` area type.

Step 11 Select the **OSPF Area Type** from the dropdown menu.

The OSPF area type can be one of the following:

- `nssa`
- `regular`
- `stub`

Step 12 Select the external routed domain from the dropdown menu.

Choose an external router domain that you have created in the APIC GUI.

Step 13 Configure OSPF settings for the site.

You can either click an existing policy (for example, `msc-ospf-policy-default`) to modify it or click **+Add Policy** to add a new OSPF policy. Then in the **Add/Update Policy** window, specify the following:

- In the **Policy Name** field, enter the policy name.
- In the **Network Type** field, choose either `broadcast`, `point-to-point`, or `unspecified`.

  The default is `broadcast`.
- In the **Priority** field, enter the priority number.

  The default is `1`.
- In the **Cost of Interface** field, enter the cost of interface.

  The default is `0`.
- From the **Interface Controls** dropdown menu, choose one of the following:
  - **advertise-subnet**
  - **bfd**
  - **mtu-ignore**
  - **passive-participation**
- In the **Hello Interval (Seconds)** field, enter the hello interval in seconds.

  The default is `10`.
- In the **Dead Interval (Seconds)** field, enter the dead interval in seconds.

  The default is `40`.
- In the **Retransmit Interval (Seconds)** field, enter the retransmit interval in seconds.

  The default is `5`.
- In the **Transmit Delay (Seconds)** field, enter the transmit delay in seconds.

  The default is `1`.

Configuring Infra: Pod Settings

This section describes how to configure pod-specific settings in each site.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, click **Sites**.

Step 3 In the **Sites** view, click **Configure Infra**.

Step 4 In the left pane, under **Sites**, select a specific site.

Step 5 In the main window, select a pod.

Step 6 In the right **POD Properties** pane, add the Overlay Unicast TEP for the POD.

This IP address is deployed on all spine switches that are part of the same pod and used for intersite known unicast traffic.

Step 7 Click **+Add TEP Pool** to add a routable TEP pool.

The routable TEP pools are used for public IP addresses for inter-site connectivity.

Step 8 Repeat the procedure for every pod in the site.

Configuring Infra: Spine Switches

This section describes how to configure spine switches in each site for Cisco ACI Multi-Site.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 In the **Main menu**, click **Sites**.

Step 3 In the **Sites** view, click **Configure Infra**.

Step 4 In the left pane, under **Sites**, select a specific site.

Step 5 In the main window, select a spine switch within a pod.

Step 6 In the right **<*Spine*> Settings** pane, click **+Add Port**.

Step 7 In the **Add Port** window, enter the following information:

- In the **Ethernet Port ID** field, enter the port ID, for example `1/29`.
- In the **IP Address** field, enter the IP address/netmask.

  The Orchestrator creates a sub-interface with VLAN 4 with the specified IP ADDRESS under the specified PORT.
- In the **MTU** field, enter the MTU. You can specify either `inherit` or a value between `576` and `9000`.

  MTU of the spine port should match MTU on IPN side.
- In the **OSPF Policy** field, choose the OSPF policy for the switch that you have configured in Configuring Infra Site-Specific Settings, on page 30.

  OSPF settings in the OSPF policy you choose should match on IPN side.
- For **OSPF Authentication**, you can pick either `none` or one of the following:
  - `MD5`
  - `Simple`

Step 8 Enable **BGP Peering** knob.

In a single Pod fabric with more than two spine switches, BGP peering should only be enabled on a pair (for redundancy) of spine switches called **BGP Speakers**. All other spine switches should have BGP peering disabled and will function as **BGP Forwarders**.

In a Multi-Pod fabric BGP peering should only be enabled on a couple of BGP speaker spine switches, each deployed in a different Pod. All other spines switches should have BGP peering disabled and function as BGP forwarders.

Step 9 In the **BGP-EVPN Router-ID** field, provide the IP address used for BGP-eVPN session between sites.

Step 10 Repeat the procedure for every spine switch.

Deploying Infra Configuration

This section describes how to deploy the Infra configuration to each APIC site.

In the top right of the main pane, choose the appropriate **Deploy** option to deploy the configuration.

If you are configuring only on-premises or only cloud sites, simply click **Deploy** to deploy the Infra configuration.

However, if you have both, on-premises and cloud site, the following two additional options become available:

- **Deploy & Download IPN Device config files:** Pushes the configuration to both the on-premises APIC site and the cloud site and enables the end-to-end interconnect between the on-premises and the cloud sites.

  In addition, this option downloads a zip file that contains configuration information that you will use to enable connectivity between the Cisco Cloud Services Router (CSR) deployed in your cloud sites and the on-premises IPsec termination device. A followup screen appears that allows you to select all or some of the configuration files to download.

- **Download IPN Device config files only:** Downloads a zip file that contains configuration information that you will use to enable connectivity between the Cisco Cloud Services Router (CSR) without deploying the configuration.

Deleting Sites Using Multi-Site Orchestrator GUI

This section describes how to delete sites using the Multi-Site GUI.

Step 1 Log in to the Multi-Site GUI.

Step 2 Ensure you unbind the site from any Schema's before trying to delete the site.

Step 3 In the **Main menu**, click **Sites**.

Step 4 In the **Sites List** page, hover over the site you want to delete and choose **Action** > **Delete** .

Step 5 Click **YES**.

Cisco ACI CloudSec Encryption

As most Cisco ACI deployments are adopting the Cisco ACI Multi-Site architecture to address disaster recovery and scale, the current security implementation using MACsec encryption within local site is becoming insufficient to guarantee data security and integrity across multiple sites connected by insecure external IP networks interconnecting separate fabrics. Cisco ACI Multi-Site Orchestrator Release 2.0(1) introduces the CloudSec Encryption feature designed to provide inter-site encryption of traffic.

Cisco ACI Multi-Site topology uses three tunnel end-point (TEP) IP addresses to provide connectivity between sites. These TEP addresses are configured by the admin on Cisco ACI Multi-Site Orchestrator and pushed down to each site's Cisco APIC, which in turn configures them on the spine switches. These three addresses are used to determine when traffic is destined for a remote site, in which case an encrypted CloudSec tunnel is created between the two spine switches that provide physical connectivity between the two sites through the Inter-Site Network (ISN).

The following figure illustrates the overall encryption approach that combines MACsec for local site traffic and CloudSec for inter-site traffic encryption.

Figure 11: CloudSec Encryption

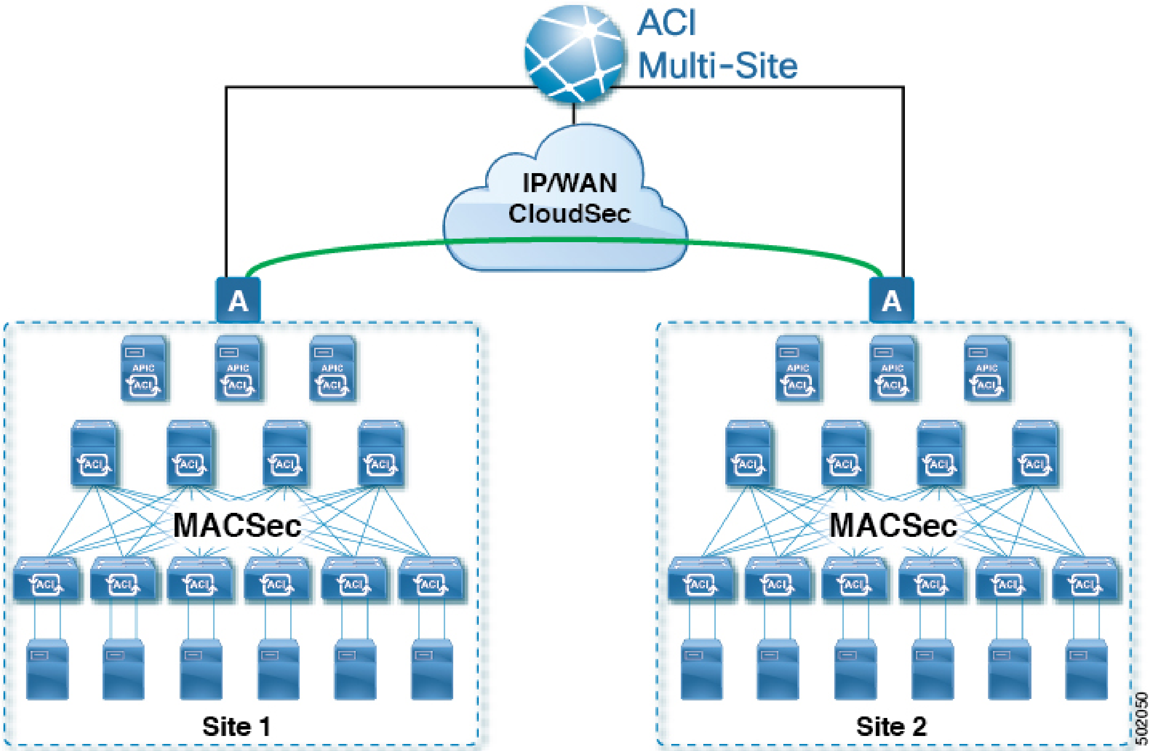

Requirements and Guidelines

When configuring CloudSec encryption, the following guidelines apply:

- CloudSec has been validated using a Nexus 9000 Inter-Site Network (ISN) infrastructure. If your ISN infrastructure is made up of different devices, or the devices are unknown (such as in the case of circuits purchased from a service provider), it is required that an ASR1K router is the first hop device directly connected to the ACI spine, or the Nexus 9000 ISN network. The ASR1K router with padding-fixup enabled allows the CloudSec traffic to traverse any IP network between the sites.

- If one or more spine switches are down when you attempt to disable CloudSec encryption, the disable process will not complete on those switches until the switches are up. This may result in packet drops on the switches when they come back up.

  We recommend you ensure that all spine switches in the fabric are up or completely decommissioned before enabling or disabling CloudSec encryption.

- The CloudSec Encryption feature is not supported with the following features:
  - Remote Leaf Direct
  - Virtual Pod (vPOD)
  - SDA
  - Intersite L3Out
  - Other routable TEP configurations

Requirements

The CloudSec encryption capability requires the following:

- Cisco ACI spine-leaf architecture with a Cisco APIC cluster for each site
- Cisco ACI Multi-Site Orchestrator to manage each site
- One **Advantage** or **Premier** license per each device (leaf only) in the fabric
- An add-on license **ACI-SEC-XF** per device for encryption if the device is a fixed spine
- An add-on license **ACI-SEC-XM** per device for encryption if the device is a modular spine

The following table provides the hardware platforms and the port ranges that are capable of CloudSec encryption.

Hardware Platform	Port Range
N9K-C9364C spine switches	Ports 49-64
N9K-C9332C spine switches	Ports 25-32
N9K-X9736C-FX line cards	Ports 29-36

If CloudSec is enabled for a site, but the encryption is not supported by the ports, a fault is raised with `unsupported-interface` error message.

CloudSec encryption's packet encapsulation is supported if Cisco QSFP-to-SFP Adapters (QSA), such as CVR-QSFP-SFP10G, is used with a supported optic. The full list of supported optics is available from the following link: https://www.cisco.com/c/en/us/support/interfaces-modules/transceiver-modules/products-device-support-tables-list.html.

CloudSec Encryption Terminology

CloudSec Encryption feature provides a secure upstream symmetric key allocation and distribution method for initial key and rekey requirements between sites. The following terminology is used in this chapter:

- `Upstream device` – The device that adds the CloudSec Encryption header and does the encryption of the VXLAN packet payload on transmission to a remote site using a locally generated symmetric cryptography key.
- `Downstream device` – The device that interprets the CloudSec Encryption header and does the decryption of the VXLAN packet payload on reception using the cryptography key generated by the remote site.

- `Upstream site` – The datacenter fabric that originates the encrypted VXLAN packets.
- `Downstream site` – The datacenter fabric that receives the encrypted packets and decrypts them.
- `TX Key` – The cryptography key used to encrypt the clear VXLAN packet payload. In ACI only one TX key can be active for all the remote sites.
- `RX Key` – The cryptography key used to decrypt the encrypted VXLAN packet payload. In ACI two RX keys can be active per remote site.

  Two RX keys can be active at the same time because during the rekey process, the downstream sites will keep the old and the new RX keys after the new key deployment is finished for some duration to ensure that out of order packet deliveries with either key can be properly decrypted.
- `Symmetric Keys` – When the same cryptography key is used to encrypt (`TX Key`) and decrypt (`RX Key`) a packet stream by the upstream and downstream devices respectively.
- `Rekey` – The process initiated by the upstream site to replace its old key with a newer key for all downstream sites after the old key expires.
- `Secure Channel Identifier (SCI)` – A 64-bit identifier that represents a security association between the sites. It is transmitted in encrypted packet in CloudSec header and is used to derive the RX key on the downstream device for packet decryption.
- `Association Number (AN)` – A 2-bit number (`0, 1, 2, 3`) that is sent in the CloudSec header of the encrypted packet and is used to derive the key at the downstream device in conjunction with the SCI for decryption. This allows multiple keys to be active at the downstream device to handle out of order packet arrivals with different keys from the same upstream device following a rekey operation.

  In ACI only two association number values (`0` and `1`) are used for the two active RX keys and only one association number value (`0` or `1`) is used for the TX Key at any point in time.
- `Pre-shared key (PSK)` – One ore more keys must be configured in the Cisco APIC GUI to be used as a random seed for generating the CloudSec TX and RX keys. If multiple PSK are configured, each rekey process will use the next PSK in order of their indexes; if no higher index PSK is available, a PSK with the lowest index will be used. Each PSK must be a hexadecimal string 64 characters long. Cisco APIC supports up to 256 pre-shared keys.

CloudSec Encryption and Decryption Handling

In order to provide a fully integrated, simple, and cost-effective solution that addresses both, data security and integrity, starting with Release 2.0(1), Cisco ACI Multi-Site provides a CloudSec Encryption feature that allows for complete source-to-destination packet encryption between Multi-Site fabrics.

The following figure shows packet diagram before and after CloudSec encapsulation, followed by descriptions of the encryption and decryption processes:

Figure 12: CloudSec Packet

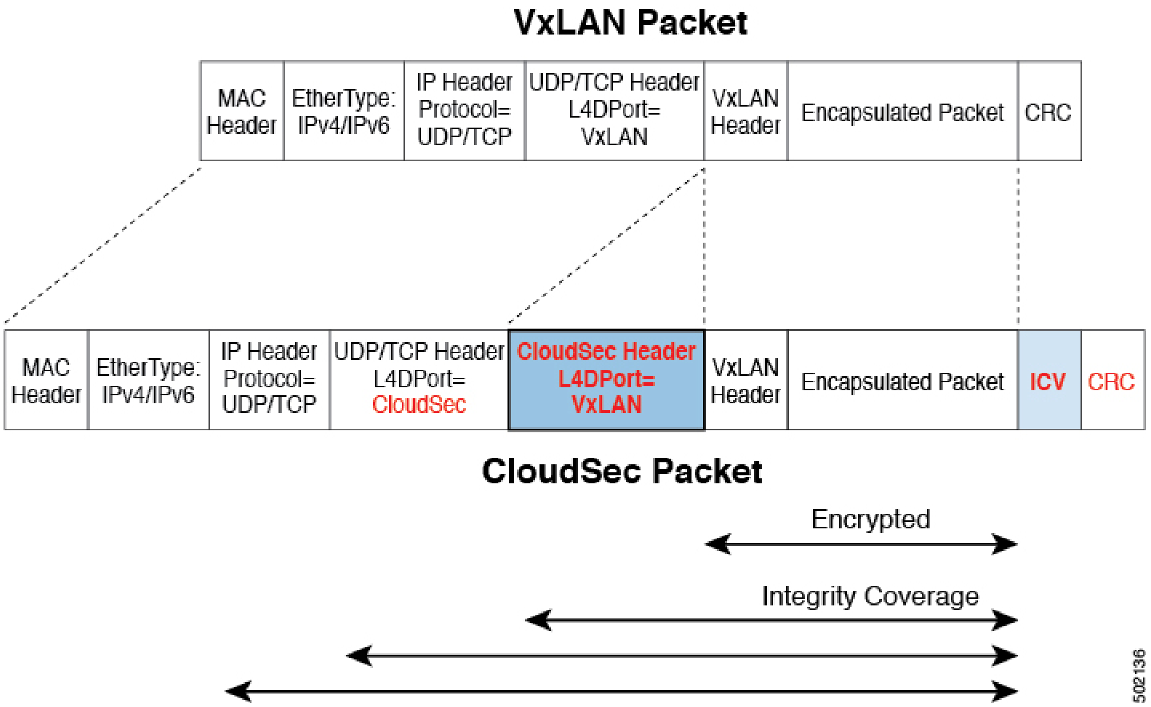

Packet Encryption

The following is a high level overview of how CloudSec handles outgoing traffic packets:

- The packets are filtered using the outer IP header and Layer-4 destination port information and matching packets are marked for encryption.
- The offset to use for encryption is calculated according to the fields of the packet. For example, the offset may vary based on whether there is a 802.1q VLAN or if the packet is an IPv4 or IPv6 packet.
- The encryption keys are programmed in the hardware tables and are looked up from the table using the packet IP header.

Once the packet is marked for encryption, the encryption key is loaded, and the offset from the beginning of the packet where to start the encryption is known, the following additional steps are taken:

- The UDP destination port number is copied from the UDP header into a CloudSec field for recovery when the packet is decrypted.
- The UDP destination port number is overwritten with a Cisco proprietary Layer-4 port number (Port `9999`) indicating that it is a CloudSec packet.
- The UDP length field is updated to reflect the additional bytes that are being added.
- The CloudSec header is inserted directly after the UDP header.
- The Integrity Check Value (ICV) is inserted at the end of the packet, between the payload and the CRC.
- The ICV requires construction of a 128-bit initialization vector. For CloudSec, any use of the source MAC address for ICV purposes is replaced by a programmable value per SCI.

- CRC is updated to reflect the change in the contents of the packet.

Packet Decryption

The way CloudSec handles incoming packets is symmetric to the outgoing packets algorithm described above:

- If the received packet is a CloudSec packet, it is decrypted and the ICV is verified.

  If ICV verification passed, the extra fields are removed, the UDP destination port number is moved from the CloudSec header to the UDP header, the CRC is updated, and the packet is forwarded to destination after decryption and CloudSec header removal. Otherwise the packet is dropped.

- If the key store returns two or more possible decryption keys, the Association Number (AN) field of the CloudSec header is used to select which key to use.

- If the packet is not a CloudSec packet, the packet is left unchanged.

CloudSec Encryption Key Allocation and Distribution

Initial Key Configuration

Figure 13: CloudSec Key Distribution

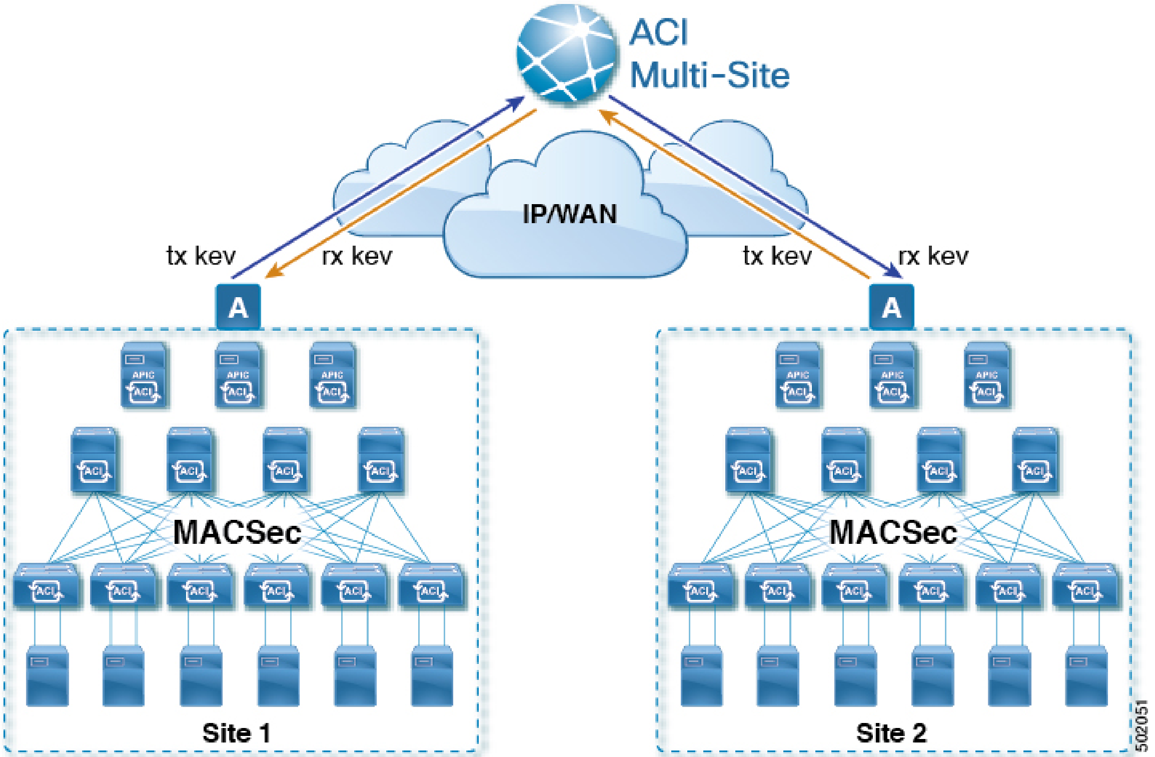

The following is a high level overview of the CloudSec encryption key initial allocation and distribution process illustrated by the figure above:

- The upstream site's Cisco APIC generates a local symmetric key intended to be used for data encryption of VXLAN packets transmitted from its site. The same key that is used by the upstream site for encryption is used for decryption of the packets on the downstream remote receiving sites.

  Every site is an upstream site for the traffic it transmits to other sites. If multiple sites exist, each site generates its own site-to-site key and use that key for encryption before transmitting to the remote site.

- The generated symmetric key is pushed to the Cisco ACI Multi-Site Orchestrator (MSO) by the upstream site's Cisco APIC for distribution to downstream remote sites.
- The MSO acts as a message broker and collects the generated symmetric key from the upstream site's Cisco APIC, then distributes it to downstream remote sites' Cisco APICs.
- Each downstream site's Cisco APIC configures the received key as RX key on the local spine switches which are intended to receive the traffic from the upstream site that generated the key.
- Each downstream site's Cisco APIC also collects the deployment status of the RX Key from the local spine switches and then pushes it to the MSO.
- The MSO relays the key deployment status from all downstream remote sites back to the upstream site's Cisco APIC.
- The upstream site's Cisco APIC checks if the key deployment status received from all downstream remote sites is successful.
    - If the deployment status received from a downstream device is successful, the upstream site deploys the local symmetric key as its TX key on the spine switches to enable encryption of the VXLAN packets that are sent to the downstream site.
    - If the deployment status received from a downstream device is failed, a fault is raised on the Cisco APIC site where it failed and it is handled based on the "secure mode" setting configured on the MSO. In "must secure" mode the packets are dropped and in the "should secure" mode the packets are sent clear (unencrypted) to the destination site.

Note In current release, the mode is always set to “should secure” and cannot be changed.

Rekey Process

Each generated TX/RX key expires after a set amount of time, by default key expiry time is set to 15 minutes. When the initial set of TX/RX keys expires, a rekey process takes place.

The same general key allocation and distribution flow applies for the rekey process. The rekey process follows the "make before break" rule, in other words all the RX keys on the downstream sites are deployed before the new TX key is deployed on the upstream site. To achieve that, the upstream site will wait for the new RX key deployment status from the downstream sites before it configures the new TX key on the local upstream site's devices.

If any downstream site reports a failure status in deploying the new RX key, the rekey process will be terminated and the old key will remain active. The downstream sites will also keep the old and the new RX keys after the new key deployment is finished for some duration to ensure that out of order packet deliveries with either key can be properly decrypted.

Note Special precautions must be taken in regards to rekey process during spine switch maintenance, see Rekey Process During Spine Switch Maintenance, on page 43 for details.

Rekey Process Failure

In case of any downstream site failing to deploy the new encryption key generated by the rekey process, the new key is discarded and the upstream device will continue to use the previous valid key as TX key. This approach keeps the upstream sites from having to maintain multiple TX keys per set of downstream sites. However, this approach may also result in the rekey process being delayed if the rekey deployment failures continue to occur with any one of the downstream sites. It is expected that the Multi-Site administrator will take action to fix the issue of the key deployment failure for the rekey to succeed.

Cisco APIC's Role in Key Management

The Cisco APIC is responsible for key allocation (both, initial key and rekey distribution), collection of the key deployment status messages from the spine switches, and notification of the Cisco ACI Multi-Site Orchestrator about each key's status for distribution to other sites.

Cisco ACI Multi-Site Orchestrator's Role in Key Management

The Cisco ACI Multi-Site Orchestrator is responsible for collecting the TX keys (both, initial key and subsequent rekeys) from the upstream site and distributing it to all downstream sites for deployment as RX keys. The MSO also collects the RX key deployment status information from the downstream sites and notifies the upstream site in order for it to update the TX key on successful RX key deployment status.

Upstream Model

In contrast to other technologies, such as MPLS, that use downstream key allocation, CloudSec's upstream model provides the following advantages:

- The model is simple and operationally easier to deploy in the networks.
- The model is preferred for Cisco ACI Multi-Site use cases.
- It provides advantages for multicast traffic as it can use the same key and CloudSec header for each copy of the replicated packet transmitted to multiple destination sites. In downstream model each copy would have to use a different security key for each site during encryption.
- It provides easier troubleshooting in case of failures and better traceability of packets from the source to destination consistently for both, unicast and multicast replicated packets.

Configuring Cisco APIC for CloudSec Encryption

You must configure one or more Pre-Shared Keys (PSK) to be used by the Cisco APIC for generating the CloudSec encryption and decryption keys. The PSK are used as a random seed during the re-key process. If multiple PSK are configured, each re-key process will use the next PSK in order of their indexes; if no higher index PSK is available, a PSK with the lowest index will be used.

Because PSK is used as a seed for encryption key generation, configuring multiple PSK provides additional security by lowering the over-time vulnerability of the generated encryption keys.

Note If no pre-shared key is configured on the Cisco APIC, CloudSec will not be enabled for that site. In that case, turning on CloudSec setting in Cisco ACI Multi-Site will raise a fault.

If at any time you wish to refresh a previously added PSK with a new one, simply repeat the procedure as if you were adding a new key, but specify an existing index.

You can configure one or more pre-shared keys in one of three ways:

- Using the Cisco APIC GUI, as described in Configuring Cisco APIC for CloudSec Encryption Using GUI, on page 41
- Using the Cisco APIC NX-OS Style CLI, as described in Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI, on page 41
- Using the Cisco APIC REST API, as described in Configuring Cisco APIC for CloudSec Encryption Using REST API, on page 42

Configuring Cisco APIC for CloudSec Encryption Using GUI

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC GUI.

Step 1 Log in to APIC.

Step 2 Navigate to **Tenants** > **infra** > **Policies** > **CloudSec Encryption**

Step 3 Specify the **SA Key Expiry Time**.

This option specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

Step 4 Click the **+** icon in the **Pre-Shared Keys** table.

Step 5 Specify the **Index** of the pre-shared key you are adding and then the **Pre-Shared Key** itself.

The **Index** field specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

Each **Pre-Shared Key** must be a hexadecimal string 64 characters long.

Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC NX-OS Style CLI.

Step 1 Log in to the Cisco APIC NX-OS style CLI.

Step 2 Enter configuration mode.

Example:

```
apic1# configure
apic1 (config)#
```

Step 3 Enter configuration mode for the default CloudSec profile.

Example:

```
apic1(config)# template cloudsec default
apic1(config-cloudsec)#
```

Step 4 Specify the Pre-Shared Keys (PSK) expiration time.

This option specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

Example:

```
apic1(config-cloudsec)# sakexpirytime <duration>
```

Step 5 Specify one or more Pre-Shared Keys.

In the following command, specify the index of the PSK you're configuring and the PSK string itself.

Example:

```
apic1(config-cloudsec)# pskindex <psk-index>
apic1(config-cloudsec)# pskstring <psk-string>
```

The *<psk-index>* parameter specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

The *<psk-string>* parameter specifies the actual PSK, which must be a hexadecimal string 64 characters long.

Step 6 (Optional) View the current PSK configuration.

You can view how many PSK are currently configured and their duration using the following command:

Example:

```
apic1(config-cloudsec)# show cloudsec summary
```

Configuring Cisco APIC for CloudSec Encryption Using REST API

This section describes how to configure one or more pre-shared keys (PSK) using the Cisco APIC REST API.

Configure PSK expiration time, index, and string.

In the following XML POST, replace:

- The value of **`sakExpiryTime`** with the expiration time of each PSK.

  This **`sakExpiryTime`** parameter specifies how long each key is valid (in minutes). Each generated TX/RX key expires after the specified amount of time triggering a re-key process. The expiration time can be between 5 and 1440 minutes.

- The value of **`index`** with the index of the PSK you're configuring.

  The **`index`** parameter specifies the order in which the pre-shared keys are used. After the last (highest index) key is used, the process will continue with the first (lowest index) key. Cisco APIC supports up to 256 pre-shared keys, so the PSK index value must be between 1 and 256.

- The value of **`pskString`** with the index of the PSK you're configuring.

The **pskString** parameter specifies the actual PSK, which must be a hexadecimal string 64 characters long.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">

    <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "false" status=""
 >
         <cloudsecPreSharedKey index="1"
pskString="1234567812345678123456781234567812345678123456781234567812345678" status=""/>
     </cloudsecIfPol>
</fvTenant>
```

Enabling CloudSec Encryption Using Cisco ACI Multi-Site Orchestrator GUI

The CloudSec encryption can be enabled or disabled for each site individually. However, the communications between two sites will be encrypted only if the feature is enabled on both sites.

Before you begin

Before you enable the CloudSec encryption between two or more sites, you must have completed the following tasks:

- Installed and configured the Cisco APIC clusters in multiple sites, as described in *Cisco APIC Installation, Upgrade, and Downgrade Guide*
- Installed and configured Cisco ACI Multi-Site Orchestrator, as described in *Cisco ACI Multi-Site Orchestrator Installation and Upgrade Guide*.
- Added each Cisco APIC site to the Cisco ACI Multi-Site Orchestrator, as described in *Cisco ACI Multi-Site Configuration Guide*.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator.

Step 2 From the left-hand sidebar, select the **Sites** view.

Step 3 Click on the **Configure Infra** button in the top right of the main window.

Step 4 From the left-hand sidebar, select the site for which you want to change the CloudSec configuration.

Step 5 In the right-hand sidebar, toggle the **CloudSec Encryption** setting to enable or disable the CloudSec Encryption feature for the site.

Rekey Process During Spine Switch Maintenance

The following is a summary of the CloudSec rekey process during typical maintenance scenarios for the spine switches where the feature is enabled:

- **Normal Decommissioning** – CloudSec rekey process stops automatically whenever a CloudSec-enabled spine switch is decommissioned. Rekey process will not start again until the decommissioned node is commissioned back or the decommissioned node ID is removed from the Cisco APIC

- **Spine Switch Software Upgrade** – CloudSec rekey process stops automatically if a spine switch is reloaded due to software upgrade. Rekey process will resume after the spine switch comes out of reload.
- **Maintenance (GIR mode)** – CloudSec rekey process must be manually stopped using the instructions provided in Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI, on page 44. Rekey can be enabled back only after the node is ready to forward traffic again.
- **Decommissioning and Removal from Cisco APIC** – CloudSec rekey process must be manually stopped using the instructions provided in Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI, on page 44. Rekey can be enabled back only after the node is removed from Cisco APIC.

Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI

It is possible to manually stop and restart the re-key process. You may be required to manually control the re-key process in certain situations, such as switch decommissioning and maintenance. This section describes how to toggle the setting using Cisco APIC NX-OS Style CLI.

Step 1 Log in to the Cisco APIC NX-OS style CLI.

Step 2 Enter configuration mode.

Example:

```
apic1# configure
apic1(config)#
```

Step 3 Enter configuration mode for the default CloudSec profile.

Example:

```
apic1(config)# template cloudsec default
apic1(config-cloudsec)#
```

Step 4 Stop or restart the re-key process.

To stop the re-key process:

Example:

```
apic1(config-cloudsec)# stoprekey yes
```

To restart the re-key process:

Example:

```
apic1(config-cloudsec)# stoprekey no
```

Disabling and Re-Enabling Re-Key Process Using REST API

It is possible to manually stop and restart the re-key process. You may be required to manually control the re-key process in certain situations, such as switch decommissioning and maintenance. This section describes how to toggle the setting using Cisco APIC REST API.

Step 1 You can disable the rekey process using the following XML message.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">

    <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "true" status=""
 />
</fvTenant>
```

Step 2 You can enable the rekey process using the following XML message.

Example:

```
<fvTenant annotation="" descr="" dn="uni/tn-infra" name="infra" nameAlias="" ownerKey="" ownerTag="">

   <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "false" status=""
 />
</fvTenant>
```

Multi-Site Cross Launch to Cisco APIC

Multi-Site currently supports the basic parameters to choose when creating a Tenant and setting up a site. Multi-Site supports most of the Tenant policies, but in addition to that you can configure some advanced parameters.

Use the Multi-Site GUI to manage the basic properties to configure. If you want to configure advanced properties, the capability to cross launch into Cisco APIC GUI directly from the Multi-Site GUI is provided. You can also configure the additional properties directly in Cisco APIC.

There are three different access points in Multi-Site GUI from where you can cross launch into APIC. From these access points in Multi-Site, you can open a new browser tab with access into Cisco APIC. You will log in to Cisco APIC at that point for the first time, and the associated screen is displayed in the Cisco APIC GUI.

Cross-Launch to Cisco APIC from Sites

Before you begin

- At least one site must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Sites** view.

Step 2 From the **Sites** list, hover over the name of the appropriate site, click the **Actions** icon at the end of the row, and choose **Open in APIC User Interface** to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Cross-Launch to Cisco APIC from Schemas

Before you begin

- At least one site based on a template must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Schemas** view.

Step 2 From the **Schemas** list, click the appropriate **<*schema-name*>**.

Step 3 From the left-hand sidebar **Sites** list, hover over the name of the appropriate site, click the **Actions** icon at the end of the row, and choose **Open in APIC User Interface** to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Cross-Launch to Cisco APIC from the Property Pane

Before you begin

- At least one site must be configured in Multi-Site.
- The site must contain at least one tenant with entities such as VRF and bridge domain configured.

Step 1 From the left-hand sidebar, open the **Schemas** view.

Step 2 From the **Schemas** list, click the appropriate **<*schema-name*>**.

Step 3 From the left-hand sidebar **Sites** list, choose the appropriate site.

Step 4 In the **Canvas**, choose the name of a specific entity.

For example, choose an available VRF, Contract, Bridge Domain, or another entity as appropriate.

The details for the specific entity are displayed in the **Property Pane** on the right.

Step 5 In the top right of the **Property Pane**, click the **Open in APIC User Interface** icon to access the Cisco APIC GUI.

The APIC GUI login screen is displayed for logging in with APIC GUI credentials.

Viewing Cisco ACI Multi-Site-Managed Objects Using the Cisco APIC GUI

When an APIC cluster is managed by Multi-Site, cloud icons indicate the relationships with other sites.

Figure 14: Viewing Multi-Site-Managed Objects Using the APIC GUI

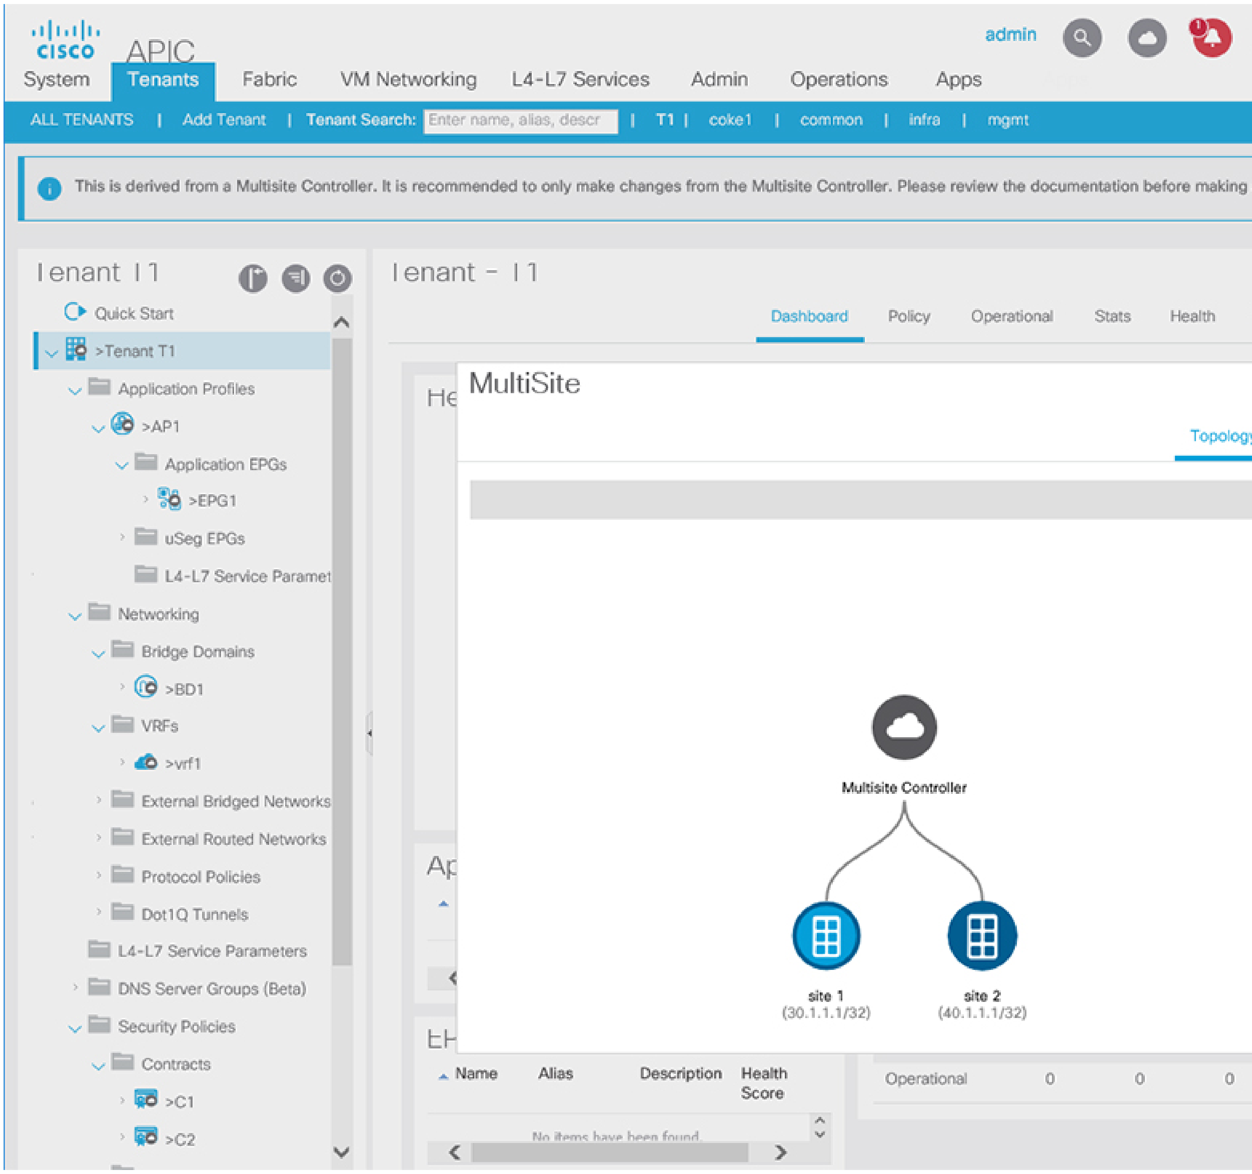

This is derived from a Multisite Controller. It is recommended to only make changes from the Multisite Controller. Please review the documentation before making

Before you begin

The APIC cluster/site must be set up to be managed by Cisco ACI Multi-Site.

Step 1 To view the relationship of the APIC site with other sites, click the cloud icon at the upper right, next to the settings icons.

In the diagram, hover over the light blue site icon to see the local site details, and hover over the dark blue icon to see the remote site details.

In the image, T1 and its Application Profile, EPG, BD, VRF, and contracts are marked with cloud icons. This indicates that they are managed by Multi-Site. We recommend that you only make changes to these objects in the Multi-Site GUI.

Step 2 To view the localized or stretched usage of a VRF, bridge domain, or other objects, where there is a **Show Usage** button on the information page, perform the following steps; for example for Bridge Domain and VRF:

a) On the menu bar, click **Tenants** and double-click on a tenant that is managed by Multi-Site.

b) Click **Networking** > **Bridge Domains** > ***BD-name*** or **Networking** > **VRFs** > ***vrf-name***.

Step 3 Click **Show Usage**.

Here you can view the nodes or policies using the object.

Note It is recommended to make changes to managed policies only in the Multi-Site GUI.

Step 4 To set the scope of deployment notification settings for this BD or VRF, click **Change Deployment Settings**. You can enable warnings to be sent for all deletions and modifications of the object on the **Policy** tab.

Step 5 To enable or disable Global warnings, check or uncheck the **(Global) Show Deployment Warning on Delete/Modify** check box.

Step 6 To enable or disable Local warnings, choose **Yes** or **No** on the **(Local) Show Deployment Warning on Delete/Modify** field.

Step 7 To view any past warnings, click the **History** tab **Events** or **Audit Logs**.

CHAPTER 4

Tenant Management

- Managing Tenants Guidelines, on page 49
- Adding Tenants, on page 50
- Configuring Global Contracts Across Tenants or VRFs, on page 50
- Configuring Intra-EPG Isolation Using the Multi-Site GUI, on page 52
- Configuring Microsegmented EPGs Using the Multi-Site GUI, on page 53
- Associating EPGs with Domains Using the Multi-Site GUI, on page 54
- Displaying All the Tenants in an Aggregated View, on page 55

Managing Tenants Guidelines

To manage tenants, you must have either `Power User` or `Site and Tenant Manager` read-write role.

You can create Tenants and their policies in one of two ways:

- Import a fully configured tenant from an APIC site.
- Create a tenant and configure the policies in the Multi-Site Orchestrator GUI.

The following tenant policies and their associations can be configured in the Multi-Site Orchestrator GUI:

- Application Profiles and EPGs
- VRFs
- Bridge Domains with subnets and stretched or site-local settings
- Contracts and Filters
- L3Outs
- External EPGs
- Physical or VMM domain association with EPGs
- Intra-EPG isolation
- Microsegmented EPGs
- EPGs deployed on a port, PC, or VPC

Adding Tenants

This section describes how to add tenants using the Multi-Site Orchestrator GUI.

Before you begin

You must have a user with either `Power User` or `Site Manager` read-write role to create and manage tenants.

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

Step 2 From the left navigation pane, select **Tenants**.

Step 3 In the main pane, click **Add Tenant**.

Step 4 In the **Display Name** field, provide the tenant's name.

The tenant's **Display Name** is used throughout the Orchestrator's GUI whenever the tenant is shown. However, due to object naming requirements on the Cisco APIC, any invalid characters are removed and the resulting **Internal Name** is used when pushing the tenant to sites. The **Internal Name** that will be used when creating the tenant is displayed below the **Display Name** textbox.

You can change the **Display Name** of the tenant at any time, but the **Internal Name** cannot be changed after the tenant is created.

Step 5 (Optional) In the **Description** field, enter a description of the tenant.

Step 6 In the **Associated Sites** section, add the sites.

a) Check all sites where you plan to deploy templates that use this tenant.

Only the selected sites will be available for any templates using this tenant.

b) From the **Security Domains** drop-down list, choose the site's security domains.

Security domains are created using the Cisco APIC GUI and can be assigned to various Cisco APIC policies and user accounts to control their access. For more information, see the *Cisco APIC Basic Configuration Guide*.

Step 7 In the **Associated Users** section, add Orchestrator users.

Only the selected users will be able to use this tenant when creating templates.

Step 8 (Optional) Enable consistency checker scheduler.

You can choose to enable regular consistency checks. For more information about the consistency checker feature, see *Cisco ACI Multi-Site Troubleshooting Guide*.

Step 9 Click **SAVE** to finish adding the tenant.

Configuring Global Contracts Across Tenants or VRFs

This use case is for a data center that provides services to EPGs in other tenants or VRFs. It provides contracts that enable all the EPGs to consume the services.

For more information, see the *Shared Services with Stretched Provider EPG* use case in the *Cisco ACI Multi-Site Fundamentals Guide*.

Before you begin

Create a schema (for every site that provides and consumes the services) with Tenants, VRFs, bridge domains, application profiles, EPGs, and other contracts.

The tenants, VRFs, BDs, and EPGs do not have to be stretched across the sites.

Step 1 Open the provider schema.

Step 2 Create a filter (essentially an Access Control List) with the following steps:

a) Click the + icon to add a filter.
b) Enter the filter name.
c) Click the + icon to add an entry.
d) Enter the entry name.
e) Enter the rest of the data required for the filter and click **Save**.

Step 3 Create a contract with the following steps:

a) Click the + icon to add a contract.
b) Enter the contract name.
c) Change the contract scope to global.

   This enables the contract to be accessible to EPGs in multiple VRFs.

d) Click the + icon to add a filter and choose the filter you created.
e) Click **Save**.

Step 4 Associate the EPG that provides the services with the contract, with the following actions:

a) Click the EPG.
b) Click the + icon to add a contract.
c) Choose the global contract you previously created.
d) Set the type to **provider**.
e) Click **Save**.
f) Click **DEPLOY TO SITES.**Confirm the sites and click **DEPLOY**.

Step 5 Associate EPGs with the contract as consumers, with the following actions:

a) Open each consumer schema.
b) Click an EPG.
c) Click the + icon to add a contract.
d) In the **Contract** field, start typing the contract name. When the contract appears in the list, choose it.
e) Set the type to **consumer**.
f) Click **Save**.
g) Associate the contract to any other EPGs in the schema.
h) Click **DEPLOY TO SITES.**
i) Confirm the sites and click **DEPLOY**.

Configuring Intra-EPG Isolation Using the Multi-Site GUI

Intra-EPG isolation is allowed between endpoints in an EPG that is operating with isolation enforced. Isolation enforced EPGs reduce the number of EPG encapsulations required when many clients access a common service but are not allowed to communicate with each other. An EPG is isolation enforced for all ACI network domains or none. While the ACI fabric implements isolation directly to connected endpoints, switches connected to the fabric are made aware of isolation rules according to a primary VLAN (PVLAN) tag.

If an EPG is configured with intra-EPG endpoint isolation enforced, these restrictions apply:

- All Layer 2 endpoint communication across an isolation-enforced EPG is dropped within a bridge domain.
- All Layer 3 endpoint communication across an isolation-enforced EPG is dropped within the same subnet.
- Preserving QoS CoS priority settings is not supported when traffic is flowing from an EPG with isolation-enforced to an EPG without isolation enforced.
- In Multi-Site, intra-EPG isolation is not supported in AVS-VLAN mode and DVS-VXLAN mode. Setting Intra-EPG isolation to be enforced may cause the ports to go into a blocked state in these domains.
- Intra-EPG isolation is not supported if the Bridge Domain is configured as "legacy BD mode".

Before you begin

- Create the tenant associated with the EPGs.
- Import the tenant policies or configure a schema containing the tenant's VRF, bridge domain, and the Application Network Profile containing the EPGs that will be subject to intra-EPG isolation.

Step 1 Open the schema and template where the EPGs to be isolated are configured.

Step 2 Click an EPG.

Step 3 Choose **Enforced**, read the warning, and click **OK**.

Step 4 Optional. Configure other EPGs to be isolation-enforced.

Step 5 Push the template containing the EPGs (configured for intra-EPG isolation) to the site where they will be located.

Step 6 Click the deployed site and template and click an EPG.

Step 7 Click **ADD STATIC PORT**.

Step 8 Choose the **PATH TYPE** (Port, Direct Port Channel, or Virtual Port Channel).

Step 9 Choose the **LEAF**.

Step 10 Choose the **PATH**.

Step 11 In the **PORT ENCAP VLAN** field, enter the VLAN number to be used for traffic for the EPG.

Step 12 On the **DEPLOYMENT IMMEDIACY** field, choose **OnDemand** or **Immediate** deployment.

Step 13 On the **MODE** field, choose **Trunk**.

Step 14 Optional, repeat the steps for other EPGs that will have isolation enforced.

What to do next

Push the changes to the site where the EPGs are located.

Configuring Microsegmented EPGs Using the Multi-Site GUI

You can use Cisco ACI Multi-Site to configure Microsegmentation to create an attribute-based EPG using a network-based attribute (IP, MAC, DNS) or VM-based attributes (VM ID, VM Name, VMM domain, and so forth). This enables you to isolate VMs or physical endpoints within a single base EPG or VMs or physical endpoints in different EPGs.

Only the basic options for microsegmented (uSeg) EPGs can be configured in Cisco ACI Multi-Site. For procedures for advanced options and for use cases and detailed information about Microsegmented EPGs, see the *Microsegmentation with Cisco ACI* chapter in *Cisco ACI Virtualization Guide, Release 3.0*.

Note When creating an EPG, if you first create an application EPG and want to change it to a uSeg EPG, you must either assign the EPG a different name or remove the application EPG and add the uSeg EPG, with the following process:

1. Delete the application EPG from the schema.
2. Deploy the schema to the sites.
3. Create the uSeg EPG.
4. Redeploy the schema to the sites.

To configure a microsegmented EPG using Cisco ACI Multi-Site, perform the following steps:

Before you begin

- Create the tenant associated with the EPGs that will be microsegmented.
- Import the tenant policies or configure a schema containing the tenant's VRF, bridge domain, and the Application Network Profile containing the EPGs.
- Create at least one application EPG in the tenant.

Step 1 Open the schema where the EPGs are configured.

Step 2 Click an EPG.

Step 3 Click **USEG EPG.**

Step 4 Click **ADD USEG ATTRIBUTES**.

Step 5 On the DISPLAY NAME field, enter the name for the attribute.

Step 6 Choose the **ATTRIBUTE TYPE; it can be one of the following:**

- **IP**
- **Mac**
- **DNS**

- **VM Name**
- **VM Data Center**
- **VM Hypervisor Identifier**
- **VM Operating System**
- **VM Tag**
- **VM Identifier**
- **VM VMM Domain**
- **VM VNIC DN** (vNIC domain name

Step 7 Save your changes.

What to do next

Associate the USeg EPG with a domain using the Multi-Site GUI.

Associating EPGs with Domains Using the Multi-Site GUI

Before you begin

- Create the tenant associated with the EPGs in Cisco ACI Multi-Site.
- Create the domain profiles (VMM, L2, L3, or Fibre Channel) in APIC.
- Import the tenant policies from Cisco APIC or configure a schema (with template) in Multi-Site, that contains the tenant's VRF, bridge domain, and the Application Network Profile containing the EPGs that will be associated with a domain.

  Associate the template with a site.

Step 1 In the **Sites** list, click the site and template for the site where the EPG and domain are configured, and click the EPG.

Step 2 Click **ADD DOMAINS**.

Step 3 On the **DOMAIN ASSOCIATION TYPE** field, choose the type, which can be:

- **VMM**
- **Fibre Channel**
- **L2 External**
- **L3 External**
- **Physical**

Step 4 On the **DOMAIN PROFILE** field, choose a previously created profile or **phys**.

Step 5 On the **DEPLOYMENT IMMEDIACY** field, choose **OnDemand** or **Immediate**.

Step 6 On the **RESOLUTION IMMEDIACY** field, choose **OnDemand**, **Immediate**, or **Pre-Provision**.

Step 7 Save your changes.

What to do next

Push the template containing the changes to the site.

Displaying All the Tenants in an Aggregated View

Using the Multi-Site GUI **Tenants** tab, you can view the aggregated list of the tenants.

In the **Tenants** panel under the **Tenants** tab, the following fields are displayed in the GUI:

- NAME: Name of the tenant.
- DESCRIPTION: Description of each tenant.
- ASSIGNED TO SITES: The number of the sites that the tenant is assigned to.
- ASSIGNED TO USERS: The number of the users that the tenant is assigned to.
- ASSIGNED TO SCHEMAS: The number of the schemas that the tenant is assigned to.
- ACTIONS: Perform actions for each tenant, for example, **Edit**, **Delete**, or configure **Network Mappings** for the tenant.

Based on the **Tenants** chart, you can determine the resource utilization of the tenants.

CHAPTER 5

User Management

- Users, Roles, and Permissions, on page 57
- Guidelines for User Management, on page 59
- Adding a User, on page 59
- Managing Users, on page 60

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	• Dashboard • Sites • Schemas • Tenants • Users • Troubleshooting Reports	`shell:msc-roles=powerUser`
Site Manager	• Dashboard—Sites • Sites • Tenants	`shell:msc-roles=siteManager`
Schema Manager	• Dashboard—Sites and Schema Health • Schemas	`shell:msc-roles=schemaManager`
Schema Editor	• Dashboard—Sites and Schema Health • Schemas	`shell:msc-roles=schemaEditor`
User Manager	• Users	`shell:msc-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 `We1come2msc!`
- 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.

Guidelines for User Management

- Users authentication and authorization can be local or external. For external authentication, you can use RADIUS, TACACS+, or LDAP servers. For more information about external authentication, see External Authentication, on page 106 in the *Administrative Operations* chapter.
- For both local and external authentication, the username supports a maximum length of 20 characters.
- For both local and external authentication, you must associate at least one role with every user. A user may be associated with more than one role. Associating a user to multiple roles offers a combination of objects that the user may access.
- Users must be associated with tenants before they can use a tenant or a schema.
- Starting with Release 2.1(2), users can be assigned roles 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.

  If you configure any read-only user roles and then downgrade your Multi-Site Orchestrator to an earlier version, which does not support read-only permissions, those roles will be removed from all users. This also means that any user that has **only** the read-only roles will have no roles assigned to them and be deleted. A Power User or User Manager will need to recreate the users and re-assign them new read-write roles.

Adding a User

This section describes how to create a Multi-Site Orchestrator user.

Step 1 Log in to Cisco ACI Multi-Site Orchestrator.

Step 2 From the main menu, select **Users**.

Step 3 In the top right of the main window pane, click **Add User**.

Step 4 In the **Add User** page, specify the following:

a) In the **Username** field, enter the new user's username.

b) In the **Password** and **Confirm Password** fields, provide the user's password.

The password must:

- Be at least 12 characters in length
- Contain at least one letter
- Contain at least one number
- Contain at least one special character apart from `*` and spaces

c) In the **First Name** field, enter the first name of the user.

d) In the **Last Name** field, enter the last name of the user.

e) In the **Email Address** field, enter the email address of the user.

f) (Optional) In the **Phone Number** field, enter the phone number of the user.

g) In the **Account Status** field, choose the account status.

You can set users to either `Active` or `Inactive` status. Only active users can log in to the Multi-Site Orchestrator.

Step 5 In the **User Roles** list, assign one or more user roles for the new user you are adding.

You must associate at least one role with every user. A user may be associated with more than one role. Associating a user to multiple roles offers a combination of features that the user can access.

Each of the available roles can be configured in read-only mode. When a user is assigned a read-only role, they can view any fabric objects available to that role, but cannot make any changes to those objects

Step 6 Click **Save**.

Managing Users

This section describes how to edit or delete existing users.

Step 1 Log in to Cisco ACI Multi-Site Orchestrator.

Step 2 If you want to update your own password...

a) Click the **User** icon in the top right of the screen.
b) Select **Reset Password**

Step 3 If you want to delete a user...

a) From the main menu, select **Users**.
b) Click the actions icon next to the user's name and select **Delete**.

You cannot delete the default `admin` user.

Step 4 If you want to edit an existing user and their permissions...

a) From the main menu, select **Users**.
b) Click the actions icon next to the user's name and select **Edit**.

You cannot change the default `admin` user's name, account status, and roles.

The default `admin` user or a user associated with the **Power User** or **User Manager** roles can update the passwords for other users. On initial log in, the user will be prompted to update their own password.

CHAPTER 6

Schema Management

- Schema Design Considerations, on page 61
- Schema Design for Cisco Cloud APIC Use-Cases, on page 67
- Creating a Schema Template, on page 67
- Migrating Objects Between Templates, on page 72
- Intersite L3Out, on page 73
- EPG Preferred Groups, on page 82
- Layer 3 Multicast, on page 84
- Shadow EPGs and BDs, on page 85

Schema Design Considerations

A schema is a collection of templates, which are used for defining policies, with each template assigned to a specific tenant. There are multiple approaches you can take when it comes to creating schema and template configurations specific to your deployment use case. The following sections describe a few simple design directions you can take when deciding how to define the schemas, templates, and policies in your Multi-Site environment. Keep in mind that when designing schemas, you must consider the supported scalability limits for the number of schemas, number of templates, and number of objects per schema. Detailed information on verified scalability limits is available in the *Verified Scalability Guides for Cisco APIC, Cisco ACI Multi-Site, and Cisco Nexus 9000 Series ACI-Mode Switches* specific to your release.

Single Schema Deployment

The simplest schema design approach is a single schema, single template deployment. You can create a single schema with a single template within it and adds all VRFs, Bridge Domains, EPGs, Contracts and other elements to that template. You can then create a single application profile or multiple application profiles within the template and deploy it to one or more sites.

Figure 15: Single Schema

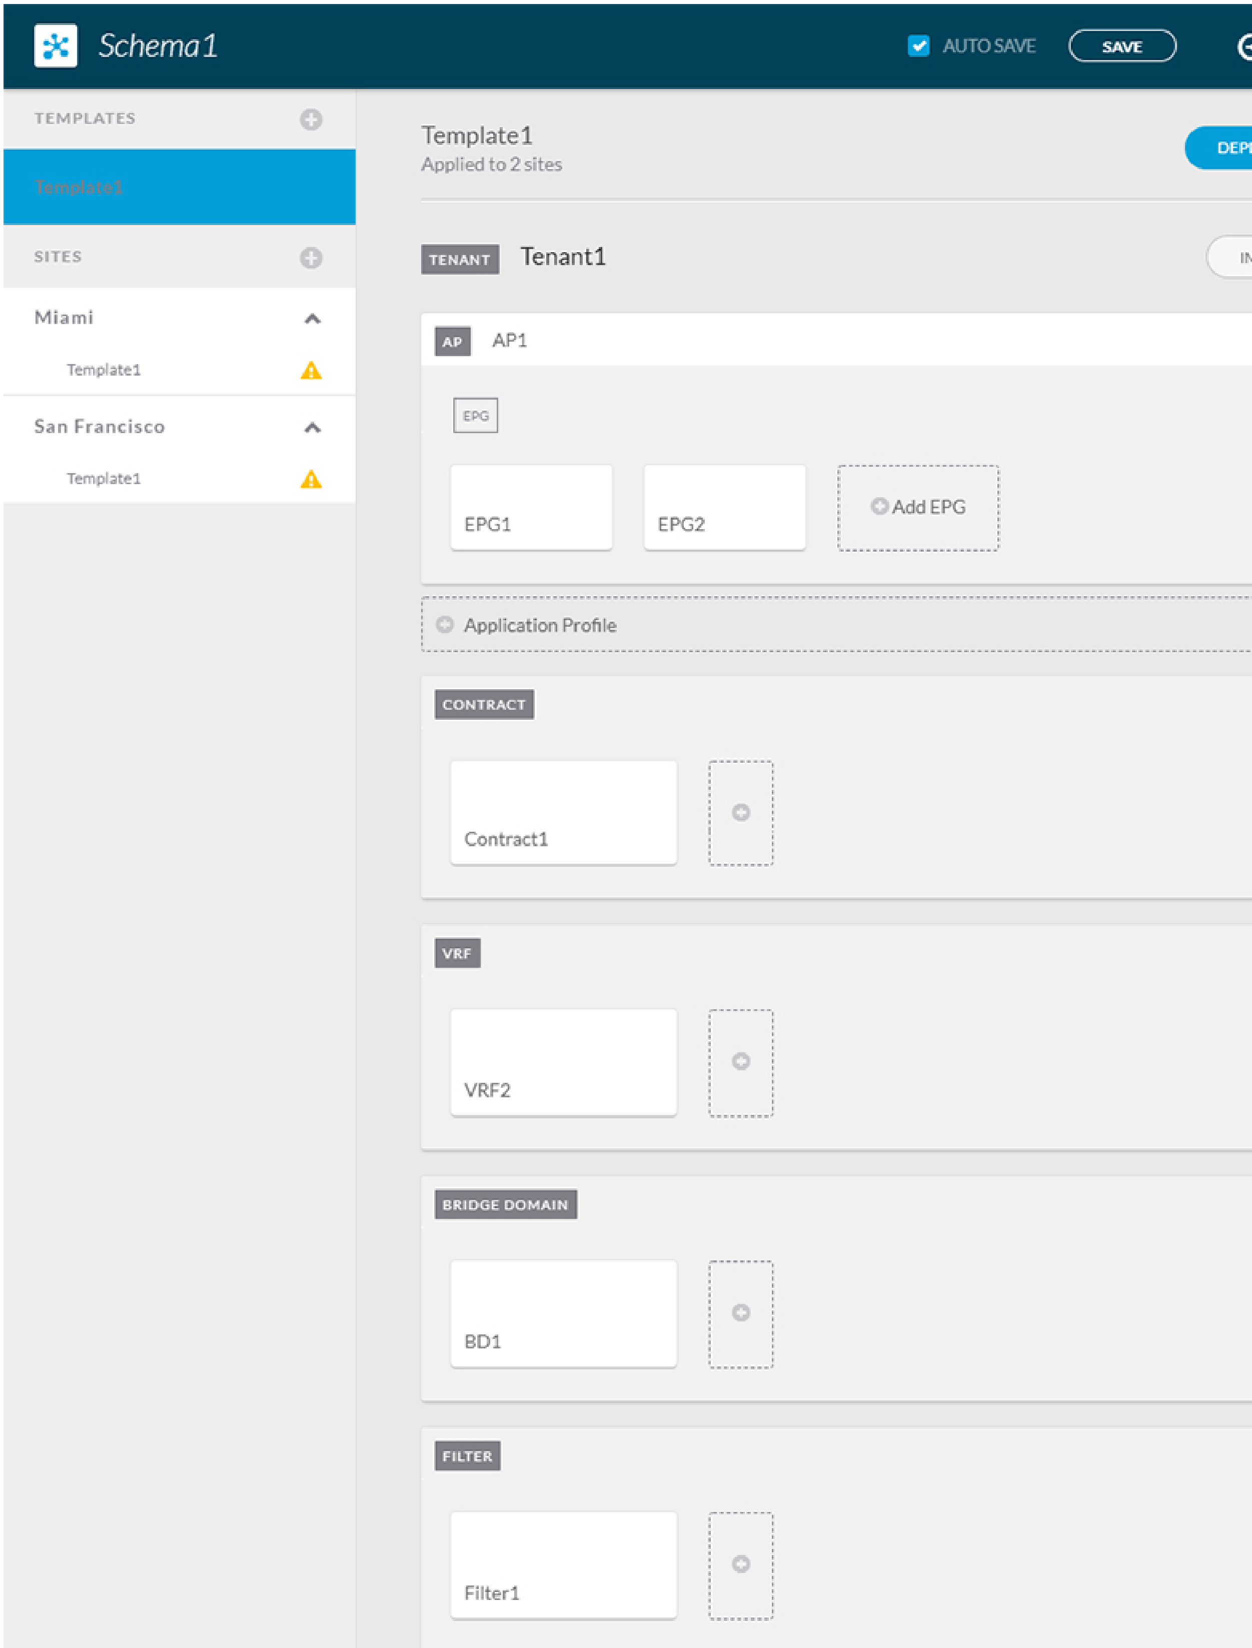
Schema1
AUTO SAVE
SAVE
TEMPLATES
Template1
SITES
Miami
Template1
San Francisco
Template1
Template1
Applied to 2 sites
DEPI
TENANT
Tenant1
IN
AP
AP1
EPG
EPG1
EPG2
Add EPG
Application Profile
CONTRACT
Contract1
VRF
VRF2
BRIDGE DOMAIN
BD1
FILTER
Filter1

This simple approach to Multi-Site schema creation is illustrated in the figure above and allows for all objects to be readily visible within the same schema. However, the supported number of schemas or templates per schema scalability limit may make this approach unsuitable for large scale deployments, which could exceed those limits.

Multiple Schemas with Network Separation

Another approach to schema design is to separate the networking objects from the application policy configuration. Networking objects include VRFs, Bridge Domains, and subnets, while the application policy objects include EPGs, Contracts, Filters, External EPGs, and Service Graphs.

You begin by defining a schema that contains the network elements. You can choose to create a single schema that contains all the network elements or you can split them into multiple schemas based on which applications reference them or which sites the network is stretched to.

The following figure shows a single networking template configuration with VRF, BD, and subnets configured and deployed to two sites:

Figure 16: Network Schema

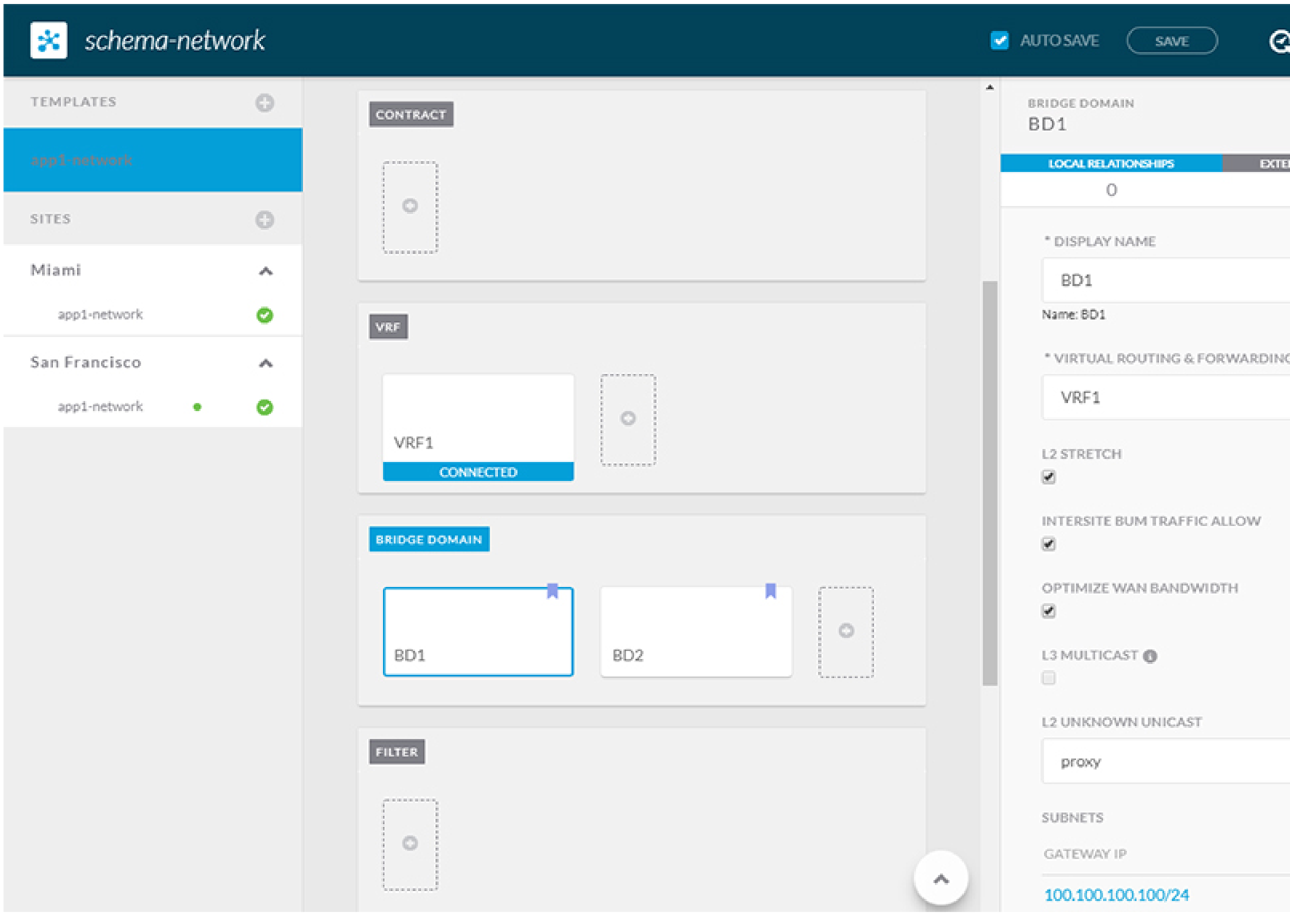

☑ AUTO SAVE

BRIDGE DOMAIN
BD1

* DISPLAY NAME
BD1
Name: BD1

* VIRTUAL ROUTING & FORWARDING
VRF1

L2 STRETCH
☑

INTERSITE BUM TRAFFIC ALLOW
☑

OPTIMIZE WAN BANDWIDTH
☑

L3 MULTICAST
☐

L2 UNKNOWN UNICAST
proxy

SUBNETS
GATEWAY IP
100.100.100.100/24

You can then define one or more separate schemas which contain each application's policy objects. This new schema can reference the network elements, such as bridge domains, defined in the previous schema. The following figure shows a policy schema that contains two application templates both of which reference the networking elements in an external schema. One of the applications is local to one site while the other is stretched across two sites:

Figure 17: Policy Schema

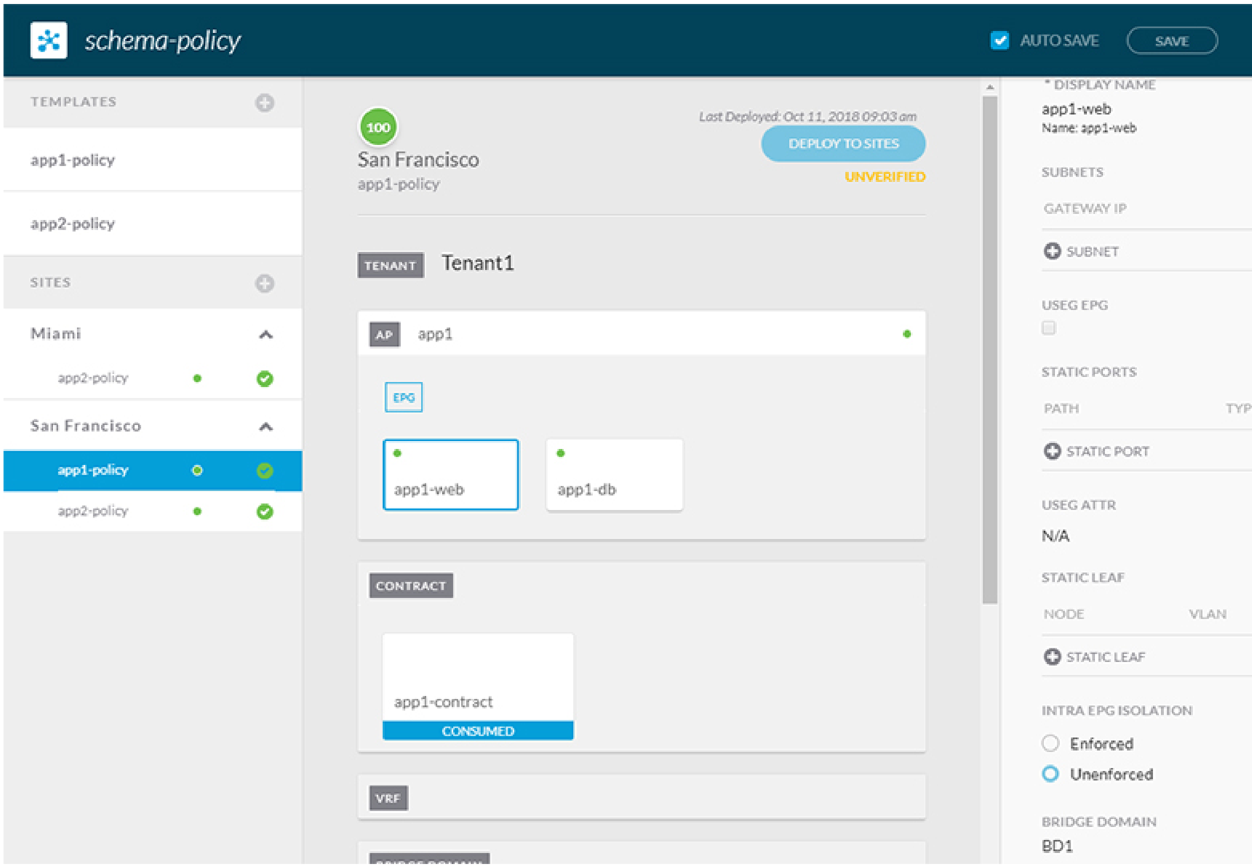

* DISPLAY NAME: app1-web
Name: app1-web

USEG EPG: □

USEG ATTR: N/A

INTRA EPG ISOLATION
○ Enforced
○ Unenforced

BRIDGE DOMAIN: BD1

After creating and deploying the policy schemas and templates, the networking objects in the networking schema will display the number of external references by the policy schema elements. The object with external references will also be denoted by the ribbon icon as shown in the Figure 16: Network Schema figure above.

Schemas designed this way provide logical separation of networking objects from the policy objects. However, this creates additional complexity when it comes to keeping track of externally referenced objects in each schema.

Multiple Schemas Based On Object Relationships

When configuring multiple schemas with shared object references, it is important to be careful when making changes to those objects. For instance, making changes to or deleting a shared networking object can impact applications in one or more sites. Because of that, you may choose to create a template around each individual site that contains only the objects used by that site and its applications, including the VRFs, BDs, EPGs, Contracts, and Filters. And create different templates containing the shared objects.

Figure 18: One Template per Site

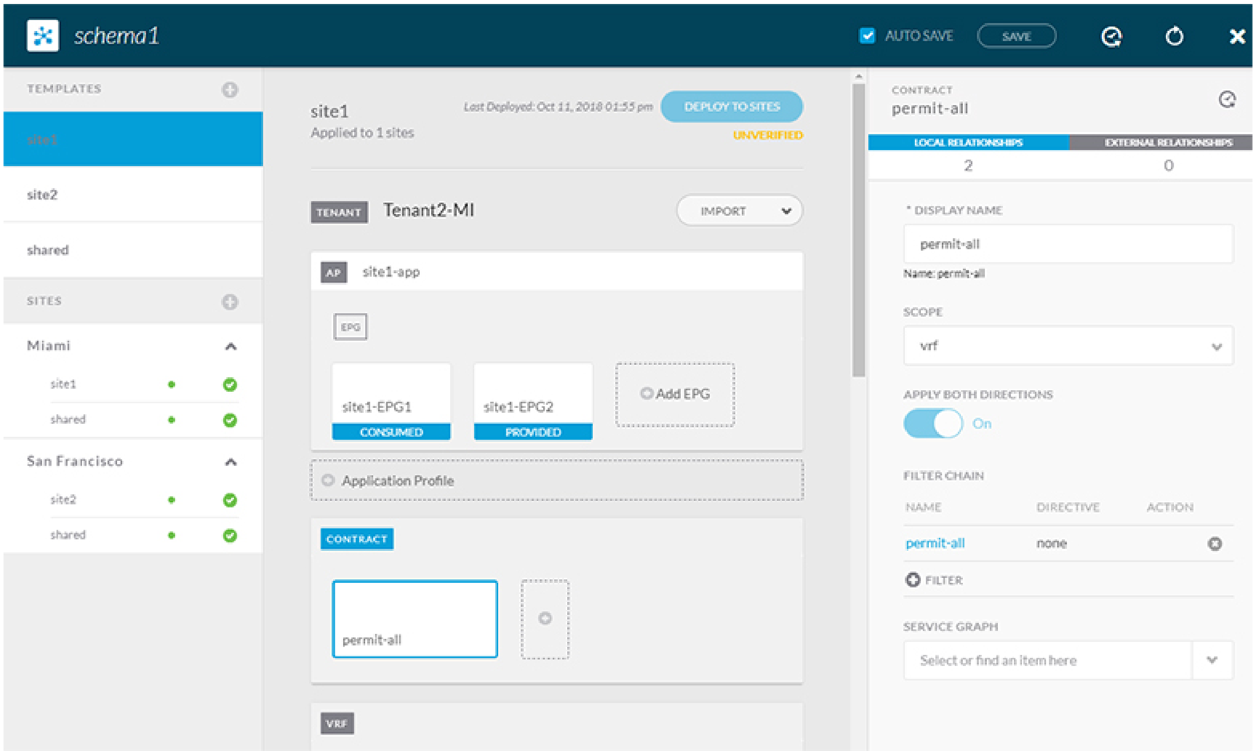

* DISPLAY NAME: permit-all

Name: permit-all

SCOPE: vrf

APPLY BOTH DIRECTIONS: On

FILTER CHAIN

NAME	DIRECTIVE	ACTION
permit-all	none	

SERVICE GRAPH: Select or find an item here

The **site1** template in the above figure contains only the objects that are local to Site1 and the template is deployed to only the Miami site. Similarly, the **site2** template contains only the object relevant to site2 and is deployed to the San Francisco site. Any change made to any object in either of these templates has no effect on the other one. The **shared** template contains any objects that are shared between the sites.

You can extend this scenario for an additional site with the following template layout:

- Site 1 template
- Site 2 template
- Site 3 template
- Site 1 and 2 shared template
- Site 1 and 3 shared template
- Site 2 and 3 shared template
- All shared template

Similarly, rather than separating objects based on which site they are deployed to, you can also choose to create schemas and templates based on individual applications instead. This would allow you to easily identify each application profile and map them to schemas and sites as well as easily configure each application as local or stretched across sites.

However, as this would exceed the 5 templates per schema limit, you would have to create additional schemas to accommodate the multiple combinations. While this creates additional complexity with multiple additional schemas and templates, it provides true separation of objects based on site or application.

Schema Design for Cisco Cloud APIC Use-Cases

Cisco ACI Multi-Site supports Cisco Cloud APIC installed in the Amazon Web Services (AWS) starting with Release 2.1(1) and Microsoft Azure starting with Release 2.2(1). Each cloud deployment can be added to and managed by the Multi-Site Orchestrator as its own APIC site.

While the sections below outline generic steps required to create and manage schemas, specific use-case scenarios supported with Cloud APIC sites are detailed in the configuration examples available from the following Cloud APIC documentation landing page: https://www.cisco.com/c/en/us/support/cloud-systems-management/cloud-application-policy-infrastructure-controller/tsd-products-support-series-home.html.

Creating a Schema Template

Before you begin

- You must have an administrative user account with full read/write privileges.
- You must have a Cisco APIC tenant user account with read/write tenant policy privileges.

  For more information, see the *User Access, Authentication, and Accounting* chapter in the *Cisco APIC Basic Configuration Guide*.
- You must have at least one available tenant that you want to incorporate into your site.

  For more information, refer to Adding Tenants, on page 50.

Step 1 On the **Schema** page, click the **Add Schema** button.

Step 2 On the **Untitled Schema** page, enter a name for the schema you intend to create.

Step 3 Access the **Select A Tenant** dialog box and select a tenant from the drop-down menu.

Note Keep in mind, the user account you're using to create a new schema must be associated with the tenant you are trying to add to it, otherwise the tenant will not be available in the drop-down menu. Associating a user account with a tenant is described in Adding Tenants, on page 50.

Importing Schema Elements From APIC Sites

You can create new objects and push them out to one or more sites or you can import existing site-local objects and manage them using the Multi-Site Orchestrator. This section describes how to import one or more existing objects, while creating new objects is described later on in this document.

Step 1 On the **Schema** page, select the schema where you want to import objects.

Step 2 In the left sidebar, select the template where you want to import objects.

Step 3 In the main pane click the **Import** button.

Step 4 Select the site from which you want to import objects.

Step 5 In the **Import** window that opens, select one or more objects you want to import.

Note The names of the objects imported into the Multi-Site Orchestrator must be unique across all sites. Importing different objects with duplicate names will cause a schema validation error and the import to fail. If you want to import objects that have the same name, you must first rename them.

Configuring an Application Profile

This section describes how to configure an Application Profile and an EPG.

Step 1 In the schema edit view, click **+ Application Profile**.

Step 2 In the properties pane on the right, provide a name for the application profile.

Step 3 In the **AP <*name*>** area, click **+ Add EPG** to add an EPG.

Step 4 In the properties pane on the right, provide a name for the EPG.

Step 5 Add a contract for the EPG.

a) Click **+ Contract**.
b) On the **Add Contract** dialog, enter the contract name and type.
c) Click **SAVE**.

Step 6 From the **Bridge Domain** dropdown, select the bridge domain for this EPG.

If you are configuring an on-premises EPG, you must associate it with a bridge domain.

Step 7 (Optional) Click **+ Subnet** to add a subnet to your EPG.

You may choose to configure a subnet on the EPG level rather than the bridge domain level, for example for a VRF route-leaking use-case.

a) On the **Add Subnet** dialog, enter the **Gateway IP** address and a description for the subnet you plan to add.
b) In the **Scope** field select either **Private to VRF** or **Advertised Externally**.
c) Click the check box for **Shared Between VRFs** if appropriate.
d) Click the check box for **No Default SVI Gateway** if appropriate.
e) Click **OK**.

Step 8 (Optional) Enable microsegmentation.

If you are configuring a microsegmentation EPG (uSeg), you must provide one or more uSeg attributes for matching endpoints to the EPG.

a) Check the **uSeg EPG** checkbox.
b) Click **+uSeg Attribute**.
c) Provide the **Name** and **Type** for the uSeg attribute.
d) Based on the attribute type you have selected, provide the attribute details.

   For example, if you have selected `MAC` for the attribute type, provide the MAC address to identify an endpoint in this EPG.

e) Click **SAVE**.

Step 9 (Optional) Enable intra-EPG isolation.

By default, endpoints in EPG can freely communicate with each other. If you would like to isolate the endpoints from each other, set the isolation mode to **Enforced**.

Step 10 (Optional) Enable Layer 3 multicast for the EPG.

For additional information about Layer 3 multicast, see Layer 3 Multicast, on page 84

Step 11 (Optional) Enable preferred group membership for the EPG.

The Preferred Group feature allows you to include multiple EPGs within a single VRF to allow full communication between them with no need for contracts to be created. For additional information about EPG preferred group, see EPG Preferred Groups, on page 82

Configuring a VRF for the Tenant

This section describes how to configure a VRF.

Step 1 In the schema edit view, scroll down to the **VRF** area and click **+**.

Step 2 In the properties pane on the right, provide a name for the VRF.

Step 3 (Optional) Enable Layer 3 multicast for the VRF.

For additional information about Layer 3 multicast, see Layer 3 Multicast, on page 84

Configuring a Bridge Domain

Step 1 In the schema edit view, scroll down to the **Bridge Domain** area and click **+**.

Step 2 In the properties pane on the right, provide the following bridge domain details:

- The BD name in the **Display Name** field.
- The VRF in the **Virtual Routing and Forwarding** field.
- If appropriate, check the **L2 STRETCH** checkbox.
- If you enabled **L2 STRETCH**, you can choose to also enable **INTERSITE BUM TRAFFIC ALLOW** checkbox.
- If you did not enable **L2 STRETCH**, you can choose either **proxy** or **flood** for the **L2 UNKNOWN UNICAST** field

Step 3 (Optional) You can choose to add one or more subnets to the bridge domain.

a) Click **+ Subnet**.

An **Add Subnet** window opens.

b) Enter the subnet's **Gateway IP** address and a description for the subnet you want to add.
c) In the **Scope** field, select either **Private to VRF** or **Advertised Externally**.
d) If appropriate, check the **Shared Between VRFs** checkbox.

e) If appropriate, check the **No Default SVI Gateway** checkbox.
f) If appropriate, check the **Querier** checkbox.
g) Click **SAVE**.

Configuring a Filter for Contracts

This section describes how to configure a filter for a contract. A filter is similar to an Access Control List (ACL), it is used to filter traffic through contracts associated to EPGs.

Step 1 In the schema edit view, scroll down to the **Filter** area and click **+**.

Step 2 In the properties pane on the right, provide a name for the filter.

Step 3 Click **+ Entry** to add a filter entry.

In the Add Entry window that opens, provide the following information:

a) A name for the filter entry.
b) (Optional) A description for the filter entry..
c) Enter the details as appropriate to filter EPG communication.

For example, to add an entry allowing HTTPS traffic through a filter, choose the following:

- **Ethertype**: `IP`
- **IP Protocol**: `TCP`
- **Destination Port Range From**: `https`

  Destination Port Range To: `https`

d) Click **SAVE**.

Configuring a Contract

This section describes how to configure a Contract.

Step 1 In the schema edit view, scroll down to the **Contracts** area and click **+**.

Step 2 In the properties pane on the right, provide a name for the contract.

Step 3 Choose a value for **Scope** using the drop-down menu.

Contract scope limits the contract's accessibility; the contract will not be applied to any consumer EPG outside the scope of the provider EPG:

- `application-profile`
- `vrf`
- `tenant`
- `global`

Step 4 Enable **Apply Both Directions** to apply the filter specified in the contract to either one direction or both directions.

The default setting is **ON**.

Step 5 Add a contract filter.

a) Click **+ Filter**.
b) On the **Add Filter Chain** dialog, click the **Name** field to choose or find a filter.
c) (Optional) Select the available directives in the **Directives** field.
d) Click **SAVE**.

Step 6 If you disabled the **Apply Both Directions** option, add the second filter chain in the other direction.

Configuring an External EPG

This section describes how to configure an External EPG.

Before you begin

- Create an L3Out in Cisco APIC on all sites where the tenant and VRF are stretched.
- The VRF for each L3Out must be the same for all sites. Changing the VRF in APIC, after the external EPGs are deployed, resets the L3Out and requires reconfiguring and redeploying the external EPG for the site.

Step 1 In the schema edit view, scroll down to the **External EPG** area and click **+**.

Step 2 In the properties pane on the right, select the type of External EPG and provide a name for it.

Cloud External EPGs are described in more detail in the Cisco Cloud APIC documentation.

Step 3 From the **Virtual Routing & Forwarding** dropdown, select the VRF to associate with this External EPG.

Step 4 Add the contracts required for the external EPGs to communicate.

Note If you are associating a contract with the external EPG as provider, choose contracts only from the tenant associated with the external EPG. Do not choose contracts from other tenants.

If you are associating the contract to the external EPG as consumer, you can choose any available contract.

Step 5 In the **On-Prem Properties** area, select an L3Out for this external EPG.

Configuring an L3Out

This section describes how to create an L3Out in the Multi-Site Orchestrator GUI. The Orchestrator then creates the L3Out on the APIC site where you deploy the template. Keep in mind that when creating an L3Out from the Orchestrator, only the L3Out container object is created in the APIC and you must still perform the full L3Out configuration (such as nodes, interfaces, routing protocols, and so on) directly in the site's APIC.

While in most cases the L3Out will be created directly at the APIC level and then associated to an external EPG that you create in the Orchestrator, it may be useful to create both here in order to directly associate the L3Out to a VRF also created in the Orchestrator.

Before you begin

Step 1 In the schema edit view, scroll down to the **L3Out** area and click **+** to add a new L3Out.

Step 2 In the properties pane on the right, provide a display name for the L3Out and the virtual routing and forwarding (VRF) for it.

Viewing Schemas

After you have created one or more schemas, they are displayed both on the Dashboard and the Schemas page.

You can use the functionality available on these two pages to monitor the usage and the health of your schemas when they are deployed. You can also access and edit specific areas of the implemented schema policies using the Multi-Site Orchestrator GUI.

For more information about the functionality of these Multi-Site Orchestrator GUI pages, refer to Overview of the Cisco ACI Multi-Site Orchestrator GUI, on page 3.

Migrating Objects Between Templates

This section describes how to move objects between templates or schemas. When moving one or more objects, the following restrictions apply:

- Only EPG and Bridge Domain (BD) objects can be moved between templates.
- Migrating objects to or from Cloud APIC sites is not supported.

  You can migrate objects between on-premises sites only.
- The source and destination templates can be in different templates and schemas, but the templates must be assigned to the same tenant.
- The destination template must have been created and assigned to at least one site.
- If the destination template is not deployed and has no other objects, the template will be automatically deployed after the objects are migrated.
- Once you initiate one object migration, you cannot perform another migration that involves the same source or target template. The migration is completed when the templates have been deployed to sites.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Schemas**.

Step 3 Click the schema that contains the objects you want to migrate.

Step 4 In the Schema view, select the Template that contains the objects you want to migrate.

Step 5 In the top right of the main pane, click **Select**.

This allows you to select one or more objects to migrate.

Step 6 Click each object that you want to migrate.

Selected objects will display a check mark in their top right corner.

Step 7 In the top right of the main pane, click the actions (**...**) icon and choose **Migrate Objects**.

Step 8 In the **Migrate Objects** window, select the destination Schema and Template where you want to move the objects.

Only the templates with at least one site attached to them will appear in the list. If you don't see your target Template in the dropdown list, cancel the wizard and assign that template to at least one site.

Step 9 Click **OK** and then **YES** to confirm that you want to move the objects.

The objects will be migrated from the source template to the destination template that you selected. When you deploy your configuration, the objects will be removed from any site where the source Template is deployed and added to the site where the destination template is deployed.

Step 10 After the migration is completed, redeploy both, the source and the destination, templates.

If the destination template is not deployed and has no other objects, the template will be automatically deployed after the objects are migrated, so you can skip this step.

Intersite L3Out

Prior to Release 2.2(1), each site managed by the Multi-Site Orchestrator required its own local L3Out configured in order to route traffic out of the fabric, which often resulted in lack of communication between endpoints in one site and a service (such a firewall, server load balancer, or mainframe) connected to the L3Out of another site.

Release 2.2(1) adds a feature that enables a number of scenarios in which endpoints located in one site are able to establish connectivity with entities, such as external network, mainframe, or service nodes, reachable through a remote L3Out.

These include the following:

- L3Out across sites—endpoints in an application EPG in one site using an L3Out in another site.

  The L3Out and the application EPG can be in the same or different VRFs and tenants.

- Transit L3Out across sites—endpoints in an external EPG in one site communicating with endpoints in an external EPG in another site.

  The external EPGs can be in the same or different VRFs and tenants.

- Shared services for intersite L3Out—shared or transit L3Out between different VRFs.

Intersite L3Out Guidelines and Limitations

When configuring an intersite L3Out, you must consider the following:

- Intersite L3Out is supported for IPv4 and IPv6.

- If you are upgrading from a release prior to Release 2.2(1), any existing External EPG to L3Out association at the site-local level will be preserved. In addition, the Orchestrator will now support creation of an L3Out and associating it with an External EPG at the template level.

  If an L3Out is defined in a schema template, it can be used for an existing External EPG:

  - If the L3Out has the **same name** as the L3Out already defined in the APIC, the Orchestrator will take ownership of that L3Out but will not manage the configuration of L3Out node profiles, interface profiles, protocol settings, or route control settings.

    If you then choose to delete this L3Out from the Orchestrator, it will no longer be managed by the Orchestrator, but any previously existing L3Out configuration will be preserved in the APIC.

  - If the L3Out has a **different name** than the APIC defined L3Out the external EPG will be removed from the APIC defined L3Out and added to the L3Out defined in the Orchestrator. If this is the only external EPG under the APIC defined L3Out this can cause the configuration to be removed from the border leaves and can impact traffic.

- If you choose to downgrade to a release prior to Release 2.2(1), the L3Outs created in the Orchestrator MSO will no longer exist in the template so any template-level association between External EPG and L3Out will be removed. In this case, you will need to manually re-configure the External EPG to L3Out association at the site-local level. Any site-local associations will be preserved during the downgrade.

- You can now associate a bridge domain in one site with the L3Out in another site, however they must both be in the same tenant.

- The Policy Control Enforcement direction for the VRF associated to the intersite L3Out must be kept configured in the default ingress mode.

- The following scenarios are not supported with intersite L3Out and remote leaf (RL):

  - Transit routing between L3Outs deployed on RL pairs associated to separate sites

  - Endpoints connected to a RL pair associated to a site communicating with the L3Out deployed on the RL pair associated to a remote site

  - Endpoints connected to the local site communicating with the L3Out deployed on the RL pair associated to a remote site

  - Endpoints connected to a RL pair associated to a site communicating with the L3Out deployed on a remote site

- The following other features are not supported with intersite L3Out in ACI Multi-Site:

  - Multicast receivers in a site receiving multicast from an external source via another site L3Out. Multicast received in a site from an external source is never sent to other sites. When a receiver in a site receives multicast from an external source it must be received on a local L3Out.

  - An internal multicast source sending multicast to an external receiver with PIM-SM any source multicast (ASM). An internal multicast source must be able to reach an external Rendezvous Point (RP) from a local L3Out

  - GOLF

  - Preferred Groups for External EPG

Configuring Routable TEP Addresses

Intersite L3Out requires a routable TEP address for the border leaf switches in each pod. If you already have a routable TEP pool configured, for example for another feature such as Remote Leaf, the same pool can be used. Otherwise, you can add a TEP pool in the Orchestrator GUI, as described in this section. Keep in mind, if you are adding a new TEP pool, it must not overlap with any other TEP pool in the fabric.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Sites**.

Step 3 In the top right of the main pane, click **Configure Infra**.

Step 4 In the left sidebar, select the site you want to configure.

Step 5 In the main window, click a pod in the site.

Step 6 In the right sidebar, click **+Add TEP Pool**.

Step 7 In the **Add TEP Pool** window, specify the routable TEP pool you want to configure for that site.

Note You must ensure that the TEP pool you are adding does not overlap with any other TEP pools or fabric addresses.

Step 8 Repeat the process for each site and pod where you plan to use intersite L3Outs.

Creating or Importing Intersite L3Out and VRF

This section describes how to create an L3Out and associate it to a VRF in the Orchestrator GUI, which will then be pushed out to the APIC site, or import an existing L3Out from one of your APIC sites. You will then associate this L3Out with an external EPG and use that external EPG to configure specific intersite L3Out use cases.

Note The VRF you assign to the L3Out can be in any template or schema, but it must be in the same tenant as the L3Out.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Schemas**.

Step 3 Select the schema and then the template where you want to create or import the VRF and L3Out.

If you create the L3Out in a template that is associated to multiple sites, the L3Out will be created on all of those sites. If you create the L3Out in a template that is associated with a single site, the L3Out will be created in that site only.

Step 4 Create a new VRF and L3Out.

If you want to import an existing L3Out, skip this step.

Note While you can create the L3Out object in the Orchestrator and push it out to the APIC, the physical configuration of the L3Out must be done in the APIC.

a) Scroll down to the **VRF** area and click the **+** icon to add a new VRF.

In the right sidebar, provide the name for the VRF, for example `vrf-l3out`

b) Scroll down to the **L3Out** area and click the **+** icon to add a new L3Out.

In the right sidebar, provide the required information.

c) Provide the name for the L3Out, for example `l3out-intersite`.

d) From the **Virtual Routing & Forwarding** dropdown, select the VRF you created in the previous step.

Step 5 Import an existing L3Out.

If you created a new L3Out in previous step, skip this step.

a) At the top of the main template view, click **Import**.

b) Select the site from which you want to import the L3Out.

c) In the import window's **Policy Type** menu, select **L3Out**.

d) Check the L3Out you want to import.

e) Click **Import**.

Configuring an External EPG to Use the Intersite L3Out

This section describes how to create an external EPG that will be associated to the intersite L3Out. You can then use this external EPG and contracts to configure specific use cases for endpoints in one site to use an L3Out in another site.

Before you begin

Create the L3Out and associate it with a VRF as described in Creating or Importing Intersite L3Out and VRF, on page 75.

Step 1 From the left navigation pane, select **Schemas**.

Step 2 Select the schema and then the template where you want to create the external EPG.

If you create the external EPG in a template that is associated to multiple sites, the external EPG will be created on all of those sites. If you create the external EPG in a template that is associated with a single site, the external EPG will be created in that site only.

Step 3 Scroll down to the **External EPG** area and click the **+** icon to add an external EPG.

In the right sidebar, provide the required information.

a) Provide the name for the external EPG, for example `eepg-intersite-l3out`.

b) From the **Virtual Routing & Forwarding** dropdown, select the VRF you created and used for the L3Out.

Step 4 If you want to assign the L3Out at the template level...

You can choose to configure the L3Out for the external EPG at the template level, in which case, you will not be able to set the L3Outs at the site-local level.

a) In the left sidebar of the schema view, select the template where the external EPG is located

b) Scroll down to the **External EPG** area and select the external EPG.

c) In the right sidebar, scroll down to the **L3Out** dropdown and choose the intersite L3Out you created.

Step 5 If you want to assign the L3Out at the site local level...

Alternatively, you can choose to associate an L3Out with the external EPG at the site-local level.

a) In the left sidebar of the schema view, select the site where the external EPG is deployed.
b) Scroll down to the **External EPG** area and select the external EPG.
c) In the right sidebar, scroll down to the **L3Out** dropdown and choose the intersite L3Out you created.

In this case, both the APIC-managed and the Orchestrator-managed L3Outs will be available for selection. You can select either the L3Out you have created in the previous section specifically for this or pick an L3Out that exists in the site's APIC.

Creating a Contract for Intersite L3Out

This section describes how to create a filter and a contract you will use to enable traffic flow between your application EPG and the external EPG that contains the intersite L3Out.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Schemas**.

Step 3 Select the schema and then the template where you want to create contract and filter.

You can use the same schema and template where you created the L3Out, VRF, and the external EPG or you can choose a different schema and template.

Step 4 Create a filter for the contract.

a) Scroll down to the **Filter** area and click **+** to create a filter.
b) In the right sidebar, provide the **Display Name** for the filter.
c) Under **Entries**, click **+Entry** to provide a filter entry.
d) In the **Add Entry** window provide the details.

The filter you create depends on your deployment and the types of traffic you want to allow.

e) Click **Save** to save the filter.

Step 5 Create a contract.

a) Scroll down to the **Contracts** area and click **+** to create a contract.
b) In the right sidebar, provide the **Display Name** for the contract.
c) From the **Scope** dropdown, select the appropriate scope.

If you plan to configure shared services and the endpoints are in a different VRF from the intersite L3Out, you must select `tenant` for the scope. Otherwise, if both are in the same VRF, you can set the scope to `vrf`.

d) You can leave the **Apply Both Directions** knob on.
e) Click **+Filter**.
f) From the **Name** dropdown menu, select the filter you created in the previous step.
g) Click **Save** to add the filter to the contract.

Configuring Intersite L3Out for Application EPGs

This section describes how to configure an application EPG to use an L3Out in another site.

The figure below shows a stretched External EPG and the associated L3Out which will be created in both sites. An application EPG (`epg1`) is created in Site 1 and has a contract with the external EPG.

Figure 19: Stretched External EPG

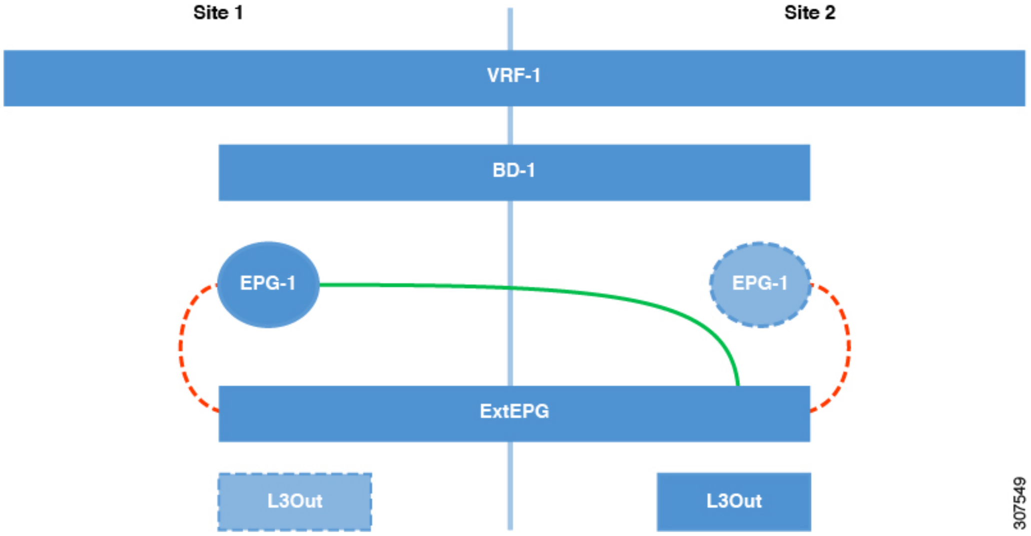

The second figure below shows a similar use case but with the external EPG being deployed to only the site where the physical L3Out is located. The application EPG and the contract are configured in the same exact way to allow the traffic flow between the EPG in one site and the physical L3Out in the other.

Figure 20: Non-Stretched (Site-Local) External EPG

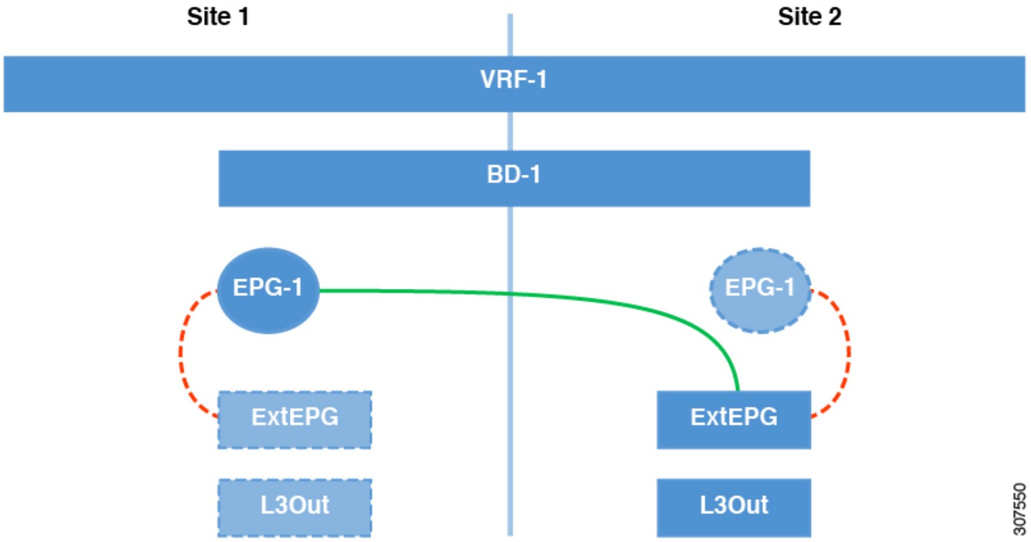

Regardless of whether you choose to stretch the external EPG containing L3Out or not, the communication between the application EPG and external EPG is enabled by the contract. The following steps describe how to create the application EPG and configure the contract between it and the L3Out external EPG you configured previously.

Before you begin

You need to have the following already configured:

- The external EPG for the intersite L3Out, as described in Configuring an External EPG to Use the Intersite L3Out, on page 76.
- The contract you will use between the application EPG and the L3Out external EPG, as described in Creating a Contract for Intersite L3Out, on page 77.
- The application EPG which will use the intersite L3Out.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Schemas**.

Step 3 Select the schema for the application EPG.

Step 4 Configure an application EPG and its bridge domain.

If you already have an EPG that will use the intersite L3Out, you can skip this step.

You can create a new or import and existing EPG and bridge domain as you typically would.

Step 5 Assign the contract to the application EPG.

a) Select the EPG.
b) In the right sidebar, click **+Contract**.
c) Select the contract you created in previous section and its type.

Step 6 Assign the contract to the external EPG that contains the intersite L3Out.

a) Browse to the template where the external EPG is located.
b) Select the external EPG.
c) In the right sidebar, click **+Contract**.
d) Select the contract you created in previous section and its type.

Step 7 Assign the templates to appropriate sites.

If you are configuring the use case shown in the first figure above where the external EPG is stretched, assign the external EPG's template to all sites and the application EPG to one site.

If you are configuring the use case shown in the second figure above where the external EPG and application EPG are local to their sites, assign the external EPG's template to one site and the application EPG's template to the other.

Step 8 Associate the application EPG's bridge domain with the L3Out.

a) In the left sidebar, under **Sites**, select the application EPG's template.
b) Select the bridge domain associated with the application EPG.
c) In the right sidebar, click **+L3Out**.
d) Select the intersite L3Out you created.

Step 9 Deploy the schema.

Configuring Transit L3Out Across Sites

This section describes how to configure communication between endpoints behind an L3Out in one site and endpoints behind an L3Out in another site.

The figure below shows two L3Outs (`l3out1` and `l3out2`) configured in different sites. Each L3Out is associated with a respective external EPG (`ExtEPG1` and `ExtEPG2`). A contract between the two external EPGs allows communication between endpoints behind two different L3Outs in two different sites.

Figure 21: Transit L3Out

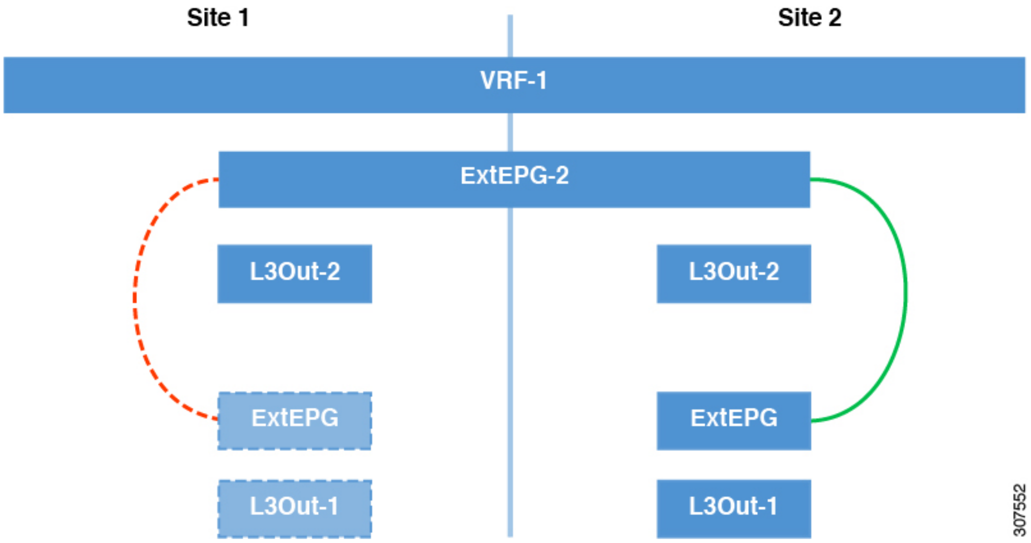

While the figure shows one of the external EPGs stretched and the other as site-local, transit L3Out supports all 3 combinations where neither external EPG is stretched, one of them is stretched, or both are stretched between sites.

Before you begin

You need to have the following already configured:

- Two different external EPGs for two different L3Outs in different sites. You can use the same procedure to create both external EPGs, as described in Configuring an External EPG to Use the Intersite L3Out, on page 76.
- The contract you will use between the application EPG and the L3Out external EPG, as described in Creating a Contract for Intersite L3Out, on page 77.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Schemas**.

Step 3 Assign the contract to one of the external EPGs.

a) Select the schema and template where the external EPG is located.

b) Select the external EPG.

c) In the right sidebar, click **+Contract**.

d) Select the contract you created in previous section and its type.

While you can pick

Step 4 Assign the contract to the other external EPG.

a) Select the schema and template where the external EPG is located.
b) Browse to the template where the external EPG is located.
c) Select the external EPG.
d) In the right sidebar, click **+Contract**.
e) Select the contract you created in previous section and its type.

Step 5 Deploy the templates to appropriate sites.

You can choose to deploy the external EPGs to one site or multiple sites. The figure above shows a example where one external EPG is stretched while the other is deployed to one site only, but you can choose any combination of stretched or site-local for the external EPGs. Since the L3Outs are in different sites, the traffic will flow through the ACI fabrics across sites.

Shared Services with Intersite L3Out

The shared services configuration for shared or transit intersite L3Out is similar to the configurations described in Configuring Intersite L3Out for Application EPGs, on page 78 and Configuring Transit L3Out Across Sites, on page 80 with a couple key differences outlined below.

Inter-VRF Shared L3Out

The figure below shows an example of inter-VRF shared L3Out scenario where an application EPG (`epg-1`) in `site1` and `vrf-1` is using `site2`'s L3Out, which is in `vrf-2`.

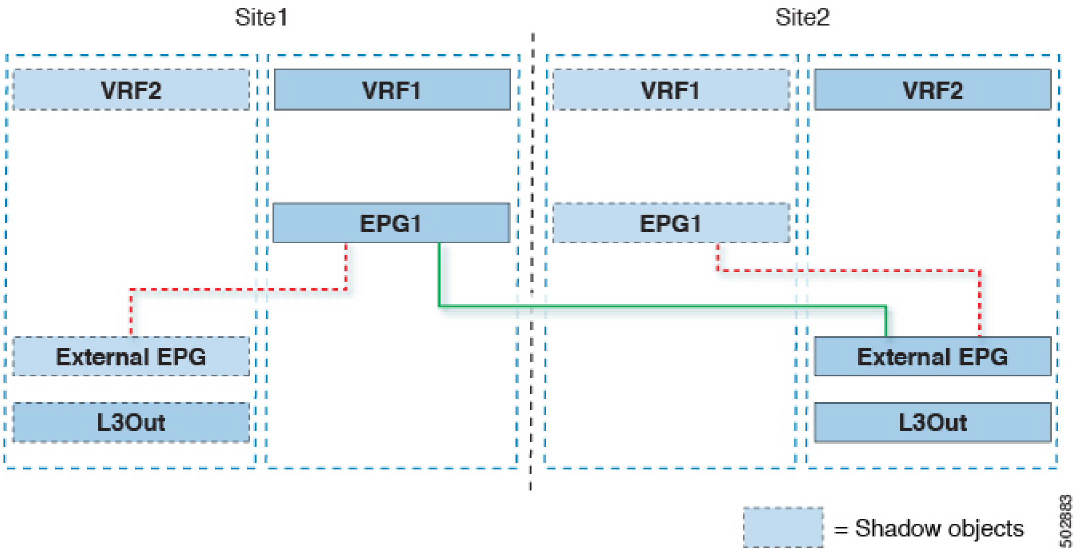

When configuring this inter-VRF use case, you must enabled the **Advertised Externally** and **Shared Between VRFs** flags when configuring the bridge domain subnets for the application EPG.

Inter-VRF Transit L3Out

And the following figure shows an example of inter-VRF transit L3Out scenario where two external EPGs with two different L3Outs located in two different VRFs are configured with a contract.

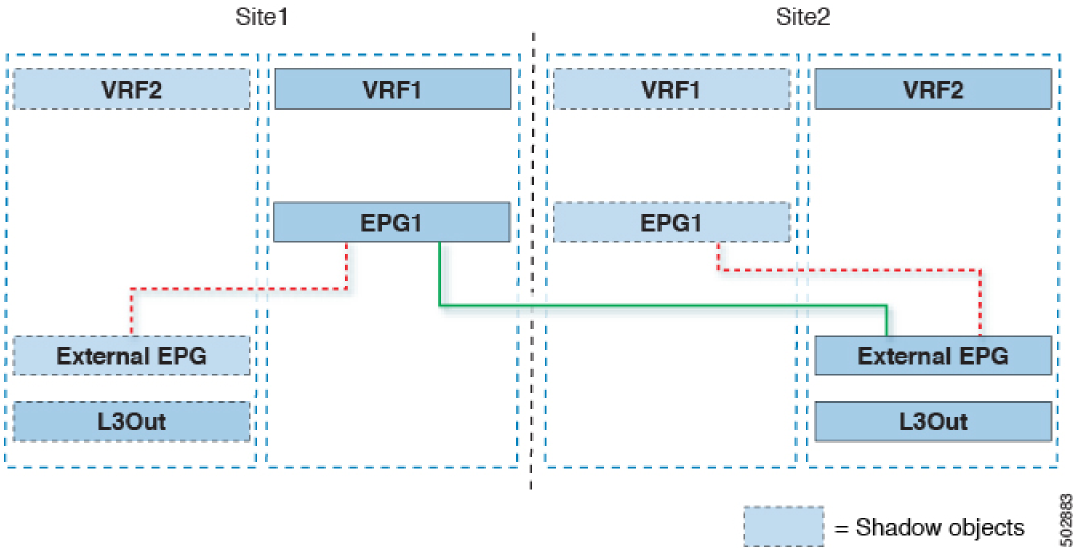

When configuring this inter-VRF use case, you must enabled the **Shared Route Control Subnet**, **Shared Security Import Subnet**, and **Aggregate Shared Routes** flags when configuring the subnets for the external EPG.

EPG Preferred Groups

By default, Multi-Site architecture allows communication between EPGs only if a contract is configured between them. If there is no contract between the EPGs, any inter-EPG communication is explicitly disabled. The Preferred Group feature allows you to specify a set of EPGs that are part of the same VRF to allow full communication between them with no need for contracts to be created.

Preferred Group vs Contracts

There are two types of policy enforcements available for EPGs in a VRF which is stretched to multiple sites with a contract preferred group configured:

- **Included EPGs** – Any EPG that is a member of a preferred group can freely communicate with all other EPGs in the group without any contracts. The communication is based on the `source-any-destination-any-permit` default rule and appropriate Multi-Site translations.
- **Excluded EPGs** – EPGs that are not members of preferred groups continue to require contracts to communicate with each other. Otherwise, the default `source-any-destination-any-deny` rule applies.

The contract preferred group feature allows for greater control and ease of configuration of communication between EPGs across sites in a stretched VRF context. If two or more EPGs in the stretched VRF require open communication while others must have only limited communication, you can configure a combination of a contract preferred group and contracts with filters to control the inter-EPG communication. EPGs that are excluded from the preferred group can only communicate with other EPGs if there is a contract in place to override the `source-any-destination-any-deny` default rule.

Stretched vs Shadowed

If EPGs from multiple sites are configured to be part of the same contract preferred group, the Multi-Site Orchestrator creates shadows of each site's EPGs in the other sites in order to correctly translate and program the inter-site connectivity from the EPGs. Contract preferred group policy construct is then applied in each site between a real and shadow EPG for inter-EPG communication.

For example, consider a web-service EPG1 in Site1 and an app-service EPG2 in Site2 added to the contract preferred group. Then if EPG1 wants to access EPG2, it will first be translated to a shadow EPG1 in Site2 and then be able to communicate with EPG2 using the contract preferred group. Appropriate BDs are also stretched or shadowed if the EPG under it is part of a contract preferred group.

Limitations

Preferred Groups are not supported for intersite L3Out external EPGs.

Configuring EPGs for Preferred Group

Before you begin

You must have one or more EPGs added to a schema template.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select the **Schemas** view.

Step 3 Click the Schema that you want to change.

Step 4 Configure one or more EPGs in the schema to be part of the preferred group.

Note If you have an existing preferred group in any of the APICs and are planning to import the EPGs from that preferred group into Multi-Site Orchestrator, you must import all EPGs in the group. You must not have a preferred group where some EPGs are managed by the Multi-Site Orchestrator and some are managed by the local APIC.

To add or remove a single EPG:

a) Select an EPG.
b) In the right properties bar, check or uncheck the **Include in Preferred Group** checkbox.
c) Click **SAVE** in the top right corner of the main window.

To add or remove multiple EPGs at once:

a) Click **SELECT** in the top-right corner of the **Application Profile** tab.
b) Select one or more EPGs by clicking on each one or click **Select All** to select all EPGs.
c) Click **...** in the top-right corner of the **Application Profile** tab and choose **Add EPGs to Preferred Group** or **Remove EPGs from Preferred Group**.
d) Click **SAVE** in the top right corner of the main window.

What to do next

You can view the full list of EPGs that are configured to be part of the preferred group by selecting a VRF and checking the **PREFERRED GROUP EPGS** list in the properties sidebar on the right.

Layer 3 Multicast

Note Layer 3 Multicast across sites is a limited availability feature. If you plan to enable this feature in your production environment, please consult Cisco for deployment planning and validation.

Cisco Multi-Site Layer 3 multicast is enabled or disabled at three levels, the VRF, the bridge domain (BD), and any EPGs that have multicast sources present.

At the top level, multicast routing must be enabled on the VRF that has any multicast-enabled BDs. On a multicast-enabled VRF, there can be a combination of multicast-enabled BDs and BDs where multicast routing is disabled. Enabling multicast routing on a VRF from the Cisco Multi-Site Orchestrator GUI enables it on the APIC sites where the VRF is stretched.

Once a VRF is enabled for multicast, the individual BDs under that VRF can be enabled for multicast routing. Configuring Layer 3 multicast on a BD enables protocol independent routing (PIM) on that BD. By default, multicast is disabled in all BDs.

When an EPG sends multicast traffic to a remote site where it is not stretched, the Multi-Site Orchestrator creates a shadow EPG on the remote site for each such EPG. This could potentially result in an increased amount of configuration changes, such as subnet routes, being pushed to the remote Top-of-Rack (TOR) switches. To alleviate this, Layer 3 multicast has to also be enabled on the individual EPGs which have multicast sources present, in which case only the configuration necessary for those EPGs is pushed to the remote sites. EPGs with multicast receivers do not require enabling Layer 3 multicast.

Multi-Site supports all of the following Layer 3 multicast source and receiver combinations:

- Multicast sources and receivers inside ACI fabric
- Multicast sources and receivers outside ACI fabric
- Multicast sources inside ACI fabric with external receivers
- Multicast receivers inside ACI fabric with external sources

Enabling Layer 3 Multicast

The following procedure describes how to enable Layer 3 multicast on VRF, BD, and EPG using the Cisco ACI Multi-Site Orchestrator GUI.

Before you begin

Cisco ACI Multi-Site Orchestrator cannot create the required local policies on each site, as such you must configure IGMP related policies, PIM related policies, route-maps, RPs, and L3Outs on each APIC site individually for end-to-end solution to work.

For specific information on how to configure those settings on each site, see the *Cisco APIC Layer 3 Networking Configuration Guide*.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left-hand sidebar, select the **Schemas** view.

Step 3 Click on the Schema you want to change.

Step 4 Enable Layer 3 multicast on a VRF.

First, you enable Layer 3 multicast on a VRF that is stretched between sites.

a) Select the VRF for which you want to enable Layer 3 multicast.
b) In the right-hand sidebar, check the **L3 Multicast** checkbox.

Step 5 Enable Layer 3 multicast on a BD.

Once you have enabled L3 Multicast on a VRF, you can enable L3 multicast on a Bridge Domain (BD) level.

a) Select the BD for which you want to enable Layer 3 multicast.
b) In the right-hand sidebar, check the **L3 Multicast** checkbox.

Step 6 Enable Layer 3 multicast on an EPG.

Once you have enabled L3 Multicast on the BD, you can select EPGs which have multicast sources. You can only do that if the EPG is part of multicast-enabled BD and VRF.

a) Select the EPG for which you want to enable Layer 3 multicast.
b) In the right-hand sidebar, check the **Intersite Multicast Source** checkbox.

Shadow EPGs and BDs

When a contract exists between site-local EPGs in stretched VRF or in Shared Services use-cases where provider and consumer are in different VRFs and communicate through Tenant contracts, the EPGs and bridge domains (BDs) are mirrored on the remote sites. These mirrored objects appear as if they are deployed in each of these sites' APICs, while only actually being deployed in one of the sites. These mirrored objects are called "shadow" EPGs or BDs.

For example, if the provider site group tenant and VRF are stretched across Site 1 and Site 2, and the consumer site group tenant and VRF are stretched across Site 3 and Site 4, in the APIC GUI at Site 1, Site 2, Site 3, and Site 4, you can see both tenants and their policies. They appear with the same names as the ones that were deployed directly to each site.

You can distinguish these shadow EPGs and BDs in the APIC GUI as described below:

Note Shadow objects should not be removed using the APIC GUI.

Step 1 To identify a shadow EPG in a pair of EPGs with the same name, in the APIC GUI, navigate to **Tenants** > ***tenant-name*** > **Application Profiles** > ***ap-name*** > **Application EPGs** > ***epg-name*** > **Static Ports**.

A shadow EPG has no path to the static port.

Step 2 To identify a shadow BD from a pair of BDs with the same name, in the APIC GUI, navigate to **Tenants** > ***tenant-name*** > **Networking** > **Bridge Domains** > ***bd-name*** > **Subnets** > ***subnet-name***.

The subnet for a shadow BD has **No Default SVI Gateway** enabled.

CHAPTER 7

Administrative Operations

- Viewing Site Status, on page 87
- Viewing Schema Health, on page 87
- Viewing Faults for Individual Sites, on page 88
- DHCP Relay Policy, on page 89
- System Logs, on page 95
- Configuration Backup and Restore, on page 97
- Custom SSL Certificates, on page 103
- External Authentication, on page 106
- System Configuration Settings, on page 110

Viewing Site Status

You can use the Multi-Site Orchestrator GUI's **Dashboard** view to see each site's status, number and types of faults, and schema health.

In the **SITE STATUS** panel, the following fields are displayed on the dashboard:

- **SITE NAME**
- **CRITICAL** Alarms
- **MAJOR** Alarms
- **MINOR** Alarms
- **WARNING** Alarms

Viewing Schema Health

Using the schema health functionality in the Multi-Site Orchestrator GUI dashboard, you can view the health of the individual schemas that are associated with different sites. In the **Schema Details** window, you can view the policy types that are associated with each site.

You can perform the following tasks using the **SCHEMA HEALTH** chart in the GUI:

- View the aggregated health score of the entire Multi-Site fabric and all APICs

- View the aggregated fault counts and the fault types for each schema in the **Schema Details** window
- View the health of the inter-site schemas
- View the health of the multi-sites nodes and their components
- View the health of the connected APICs and ACI clusters

You can view the schema health in the GUI using the following different formats:

- Hovering on an Individual Cell: Each cell in the **SCHEMA HEALTH** chart represents the health of the schema. If the cell is color coded as Green and if you hover on the cell, it displays the application health score of the schema.
- Clicking in the Cell: If you click the individual cell in the table, it provides the additional schema details for the template and the faults with the associated with each policy type, for example, ANP, EPG, Contract, VRF, and BD.

  The faults and warnings are displayed in the columns to the right side of each policy. This functionality is used to collect the details and get more information on the issues causing low health.
- Viewing the Health Score Slider: The health score slider at the top of the page provides capabilities to filter the schemas by the minimum or maximum health score. A range in the slider can be adjusted to view the schemas that match the heath score range. For example, you can adjust the health score to display the schemas matching the health score between 0 to 30 range.
- Using the Search Functionality: The search functionality in the schema health view provides the capabilities to find a schema or a policy based on the keywords that are typed in the search area. When the keywords are typed in the search area, only schemas that contain the keywords are displayed. The results are based on the matching keywords as part of the schema name, template name, or any of the contained policies within that schema.

Viewing Faults for Individual Sites

This section describes how to display the faults for the individual sites using the Multi-Site GUI.

Step 1 Log in to Multi-Site Orchestrator GUI.

Step 2 In the **Main Menu**, click **Sites**.

Step 3 In the **Sites list** page, click **CONFIGURE INFRA**.

Step 4 In the **Fabric Connectivity Infra** page, click the appropriate site in the **Master List**. For example, click site1.

The site details with the associated pods and the spines are displayed in the GUI.

The total number of the faults and the fault types, for example, Critical, Major, Minor, and Warning faults are displayed at the top of the panel. Clicking on each fault type displays the fault details with the individual codes and their explanations.

DHCP Relay Policy

Typically, when your DHCP server is located under an EPG, all the endpoints in that EPG have access to it and can obtain the IP addresses via DHCP. However, in many deployment scenarios, the DHCP server may not exist in the same EPG, BD, or VRF as all the clients that require it. In these cases a DHCP relay can be configured to allow endpoints in one EPG to obtain IP addresses via DHCP from a server that is located in another EPG/BD deployed in a different site or even connected externally to the fabric and reachable via an L3Out connection.

You can create the DHCP `Relay` policy in the Orchestrator GUI to configure the relay. Additionally, you can choose to create a DHCP `Option` policy to configure additional options you can use with the relay policy to provide specific configuration details. For all available DHCP options refer to RFC 2132.

When creating a DHCP relay policy, you specify an EPG (for example, `epg1`) or external EPG (for example, `ext-epg1`) where the DHCP server resides. After you create the DHCP policy, you associate it with a bridge domain, which in turn is associated with another EPG (for example, `epg2`) allowing the endpoints in that EPG to reach the DHCP server. Finally, you create a contract between the relay EPG (`epg1` or `ext-epg1`) and application EPG (`epg2`) to allow communication. The DHCP policies you create are pushed to the APIC when the bridge domain to which the policy is associated is deployed to a site.

Guidelines and Limitations

The DHCP relay policies are supported with the following caveats:

- DHCP relay policies are supported for fabrics running Cisco APIC Release 4.2(1) or later.
- The DHCP servers must support DHCP Relay Agent Information Option (Option 82).

  When an ACI fabric acts as a DHCP relay, it inserts the DHCP Relay Agent Information Option in DHCP requests that it proxies on behalf of clients. If a response (DHCP offer) comes back from a DHCP server without Option 82, it is silently dropped by the fabric.
- DHCP relay policies are supported in user tenants or the `common` tenant only. DHCP policies are not supported for the `infra` or `mgmt` tenants.

  When configuring shared resources and services in the ACI fabric, we recommend creating those resources in the `common` tenant, that way they can be used by any user tenant.
- DHCP relay server must be in the same user tenant as the DHCP clients or in the `common` tenant.

  The server and the clients cannot be in different user tenants.
- DHCP relay policies can be configured for the primary SVI interface only.

  If the bridge domain to which you assign a relay policy contains multiple subnets, the first subnet you add becomes the primary IP address on the SVI interface, while additional subnets are configured as secondary IP addresses. In certain scenarios, such as importing a configuration with a bridge domain with multiple subnets, the primary address on the SVI may change to one of the secondary addresses, which would break the DHCP relay for that bridge domain.

  You can use the `show ip interface vrf all` command to verify IP address assignments for the SVI interfaces.

- If you make changes to the DHCP policy after you have assigned it to a bridge domain and deployed the bridge domain to one or more sites, you will need to re-deploy the bridge domain for the DHCP policy changes to be updated on each site's APIC.
- For inter-VRF DHCP relay with the DHCP server reachable via an L3Out, DHCP relay packets must use site-local L3Out to reach the DHCP server. Packets using an L3Out in a different site (Intersite L3Out) to reach the DHCP server is not supported.
- The following DHCP relay configurations are not supported:
  - DHCP relay clients behind an L3Out.
  - Importing existing DHCP policies from APIC.
  - DHCP relay policy configuration in Global Fabric Access Policies is not supported
  - Multiple DHCP servers within the same DHCP relay policy and EPG.

    If you configure multiple providers under the same DHCP relay policy, they must be in different EPGs or external EPGs.

Creating DHCP Relay Policies

This section describes how to create a DHCP relay policy.

Note If you make changes to the DHCP policy after you have assigned it to a bridge domain and deployed the bridge domain to one or more sites, you will need to re-deploy the bridge domain for the DHCP policy changes to update on each site's APIC.

Before you begin

You must have the following:

- A DHCP server set up and configured in your environment.
- If the DHCP server is part of an application EPG, that EPG must be already created in the Multi-Site Orchestrator, as described in the Schema Management, on page 61 chapter.

  If the DHCP server is external to the fabric, the external EPG associated to the L3Out that is used to access the DHCP server must be already created, as described in the Schema Management, on page 61 chapter.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Policies**.

Step 3 In the top right of the main pane, click **Add Policy** and select **DHCP**.

This opens an **Add DHCP** configuration screen.

Step 4 In the **Name** field, specify the name for the policy.

Step 5 From the **Select Tenant** dropdown, select the tenant that contains the DHCP server.

Step 6 (Optional) In the **Description** field, provide a description for the policy.

Step 7 Select `Relay` for the **Type**.

Step 8 Click **+Provider**.

Step 9 Select the provider type.

When adding a relay policy, you can choose one of the following two types:

- `Application EPG`—specifies a specific application EPG that includes the DHCP server you are adding as an endpoint.
- `L3 External Network`—specifies the External EPG associated to the L3Out that is used to access the DHCP server.

Note You can select any EPG or external EPG that has been created in the Orchestrator and assigned to the tenant you specified, even if you have not yet deployed it to sites. If you select an EPG that hasn't been deployed, you can still complete the DHCP relay configuration, but you will need to deploy the EPG before the relay is available for use.

Step 10 From the dropdown menu, pick the EPG or external EPG.

Step 11 In the **DHCP Server Address** field, provide the IP address of the DHCP server.

Step 12 Click **Save** to add the provider.

Step 13 (Optional) Add any additional providers.

Repeat steps 9 through 12 for each additional DHCP server.

Step 14 Click **Save** to save the DHCP relay policy.

Creating DHCP Option Policies

This section describes how to create a DHCP option policy. DHCP options are appended to the end of the messages that DHCP servers and clients exchange and can be used to provide additional configuration information to your DHCP server. Each DHCP option has a specific code that you must provide when adding the option policy. For a complete list of DHCP options and codes, see RFC 2132.

Before you begin

You must have the following already configured:

- A DHCP server set up and configured in your environment.
- An EPG that contains the DHCP server already created in the Multi-Site Orchestrator, as described in the Schema Management, on page 61 chapter.
- A DHCP Relay policy created, as described in Creating DHCP Relay Policies, on page 90.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Policies**.

Step 3 In the top right of the main pane, click **Add Policy** and select **DHCP**.

This opens an **Add DHCP** configuration screen.

Step 4 In the **Name** field, specify the name for the policy.

This is a name for the policy you're creating, not a specific DHCP option name. Each policy can contain multiple DHCP options.

Step 5 From the **Select Tenant** dropdown, select the tenant that contains the DHCP server.

Step 6 (Optional) In the **Description** field, provide a description for the policy.

Step 7 Select `Option` for the **Type**.

Step 8 Click **+Option**.

Step 9 Specify a name of the option.

While not technically required, we recommend using the same name for the option as listed in RFC 2132.

For example, `Name Server`.

Step 10 Specify an ID for the option .

You must provide the option code as listed in RFC 2132.

For example, `5` for Name Server option.

Step 11 Specify the option's data.

Provide the value if the option requires one.

For example, a list of name servers available to the client for the Name Server option.

Step 12 Click the check mark next to the **Data** field to save the option.

Step 13 (Optional) Repeat the steps to add any additional options.

Step 14 Click **Save** to save the DHCP option policy.

Assigning DHCP Policies

This section describes how to assign a DHCP policy to a bridge domain.

Note If you make changes to the DHCP policy after you have assigned it to a bridge domain and deployed the bridge domain to one or more sites, you will need to re-deploy the bridge domain it for the DHCP policy changes to be updated on each site's APIC.

Before you begin

You must have the following already configured:

- A DHCP relay policy, as described in Creating DHCP Relay Policies, on page 90.
- (Optional) A DHCP option policy, as described in Creating DHCP Option Policies, on page 91.
- The bridge domain to which you will assign the DHCP policy, as described in the Schema Management, on page 61 chapter.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Schemas**.

Step 3 Select the schema where the bridge domain is defined.

Step 4 Scroll down to the **Bridge Domain** area and select the bridge domain.

Step 5 In the right sidebar, scroll down and check the **DHCP Policy** option checkbox.

Step 6 From the **DHCP Relay Policy** dropdown, select the DHCP policy you want to assign to this BD.

Step 7 (Optional) From the **DHCP Option Policy** dropdown, select the option policy.

A DHCP option policy provides additional options to be passed to the DHCP relay. For additional details see Creating DHCP Option Policies, on page 91.

Step 8 Assign the bridge domain to any EPG that needs access to the DHCP server via the relay.

Creating DHCP Relay Contract

DHCP packets are not filtered by contracts but contracts are required in many cases to propagate routing information within the VRF and across VRFs. Even though the DHCP packets are not filtered it is recommended to configure contracts between the client EPG and the EPG configured as the provider in the DHCP relay policy.

This section describes how to create a contract between the EPG that contains the DHCP server and the EPG that contains endpoints that need to use the relay. Even though you have already created and assigned the DHCP policy to the bridge domain and the bridge domain to the clients' EPG, you must create and assign the contract to enable programming of routes to allow client to server communication.

Before you begin

You must have the following already configured:

- A DHCP relay policy, as described in Creating DHCP Relay Policies, on page 90.
- (Optional) A DHCP option policy, as described in Creating DHCP Option Policies, on page 91.
- The bridge domain to which you have assigned the DHCP policy, as described in Assigning DHCP Policies, on page 92.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Schemas**.

Step 3 Select the schema where you want to create the contract.

Step 4 Create a contract.

DHCP packets are not filtered by the contract so no specific filter is required, but a valid contract should be created and assigned to ensure proper BD and routes deployment.

a) Scroll down to the **Contracts** area and click **+** to create a contract.

b) In the right sidebar, provide the **Display Name** for the contract.

c) From the **Scope** dropdown, select the appropriate scope.

Because the DHCP server EPG and application EPG must be in the same tenant, you can select one of the following:

- `vrf`, if both EPGs are in the same VRF
- `tenant`, if the EPGs are in different VRFs

d) You can leave the **Apply Both Directions** knob on.

Step 5 Assign the contract to the DHCP relay EPG.

a) Browse to the template where the EPG is located.
b) Select the EPG or external EPG where the DHCP server resides.

This is the same EPG you selected when creating the DHCP relay policy.

c) In the right sidebar, click **+Contract**.
d) Select the contract you created and `provider` for its type.

Step 6 Assign the contract to the application EPG whose endpoints require DHCP relay access.

a) Browse to the template where the application EPG is located.
b) Select the application EPG.
c) In the right sidebar, click **+Contract**.
d) Select the contract you created and `consumer` for its type.

Verifying DHCP Relay Policies in APIC

This section describes how to verify that the DHCP relay policies you have created and deployed using the Multi-Site Orchestrator are correctly pushed to each site's APIC. The DHCP policies you create are pushed to the APIC when the bridge domain to which the policy is associated is deployed to a site.

Step 1 Log in to the site's APIC GUI.

Step 2 From the top navigation bar, select **Tenants** > **<tenant-name>**.

Select the tenant where you deployed the DHCP policy.

Step 3 Verify that the DHCP relay policy is configured in APIC.

In the left tree view, navigate to **<tenant-name>** > **Policies** > **Protocol** > **DHCP** > **Relay Policies**. Then confirm that the DHCP relay policy you configured has been created.

Step 4 Verify that the DHCP option policy is configured in APIC.

If you have not configured any DHCP option policies, you can skip this step.

In the left tree view, navigate to **<tenant-name>** > **Policies** > **Protocol** > **DHCP** > **Option Policies**. Then confirm that the DHCP option policy you configured has been created.

Step 5 Verify that the DHCP policy is correctly associated with the bridge domain.

In the left tree view, navigate to **<tenant-name>** > **Networking** > **Bridge Domains** > **<bridge-domain-name>** > **DHCP Relay Labels**. Verify that the DHCP policy is also associated with the deployed bridge domain.

Editing or Deleting Existing DHCP Policies

This section describes how to edit or delete a DHCP relay or option policy.

Note

- If you make changes to the DHCP policy after you have assigned it to a bridge domain and deployed the bridge domain to one or more sites, you will need to re-deploy it for the DHCP policy changes to update on each site's APIC.
- You cannot deleted policies that are associated with one or more bridge domains, you must first unassign the policy from every bridge domain.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Policies**.

Step 3 Click the actions menu next to the DHCP policy and select **Edit** or **Delete**.

System Logs

Multi-Site Orchestrator system logging is automatically enabled when you first deploy the Orchestrator cluster and captures the events and faults that occur in the environment.

You can view the Multi-Site Orchestrator logs by selecting **Admin** > **Audit Logs** from the main navigation menu.

From the **Audit Logs** page, you can click the **Most Recent** field to select a specific time period for which you want to see the logs. For example, when you select the range from November 14, 2019 to November 17, 2019 and click **Apply**, the audit log details for this time period are displayed on the **Audit Logs** page.

You can also click the **Filter** icon to filter the log details using the following criteria:

- **User**: Select this option to filter the audit logs by the user type, then click **Apply** to apply the filter.
- **Type**: Select this option to filter the audit logs by the policy types, for example, site, user, template, application profile, bridge domain, EPG, external EPG, filter, VRF, BGP config, contract, OSPF policy, pod, node, port, domain, provider, RADIUS, TACACS+ and click **Apply**.
- **Action**: Select this option to filter the audit logs by an action. The available actions are Created, Updated, Deleted, Added, Removed, Associated, Disassociated, Deployed, Undeployed, Downloaded, Uploaded, Restored, Logged in, Logged Out, Login Failed. Select an action and click **Apply** to filter the log details according to the action.

Generating Troubleshooting Report and System Logs

This section describes how to generate a troubleshooting report and infrastructure logs file for all the schemas, sites, tenants, and users that are managed by Cisco ACI Multi-Site Orchestrator.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 In the main menu, select **Operations** > **Tech Support**.

Step 3 In the top right corner of the **System Logs** frame, click the edit button.

A **System Logs** configuration window opens.

Step 4 In the **System Logs** window, check the logs you want to download.

Check the **Database Backup** to download a backup of the Orchestrator database.

Check the **Server Logs** to download the Orchestrator cluster logs.

Step 5 Click **Download**.

An archive of the selected items will be downloaded to your system. The report contains the following information:

- All schemas in JSON format
- All sites definitions in JSON format
- All tenants definitions in JSON format
- All users definitions in JSON format
- All logs of the containers in the `infra_logs.txt` file

Enabling Log Streaming to an External Log Analyzer

Cisco ACI Multi-Site Orchestrator allows you to send the Orchestrator logs to an external log analyzer tool in real time. By streaming any events as they are generated, you can use the additional tools to quickly parse, view, and respond to important events without a delay.

This section describes how to enable Multi-Site Orchestrator to stream its logs to an external analyzer tool, such as Splunk.

Before you begin

- Set up and configure the log analyzer service provider.

  For detailed instructions on how to configure an external log analyzer, consult its documentation.

  Note This release of Cisco ACI Multi-Site Orchestrator only supports Splunk as the service provider.

- Obtain an authentication token for the service provider.

  Obtaining an authentication token for Splunk service is detailed in the Splunk documentation, but in short, you can get the authentication token by logging into the Splunk server, selecting **Settings** > **Data Inputs** > **HTTP Event Collector**, and clicking **New Token**.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 In the top right corner, click the **Options** icon and select **System Logs**.

Step 3 In the **System Logs** window that opens, enable the **EXTERNAL STREAMING** knob.

Step 4 Select which logs you want to stream.

You can select either all logs or audit logs only.

Step 5 From the **SELECT SERVICE** dropdown menu, select the log analyzer service.

This release of Cisco ACI Multi-Site Orchestrator supports only Splunk as the service provider.

Step 6 Choose the **PROTOCOL** for the traffic.

Select **UNSECURE** for HTTP or **SECURE** for HTTPS.

Step 7 Provide the service's information.

In the **HOST** field, enter the host's IP address.

In the **PORT** field, enter the host's port number.

In the **TOKEN** field, enter the authentication token you obtain from the service provider.

Step 8 For each Multi-Site Orchestrator node, provide the node's root password.

Note This is the `root` user password of each Orchestrator node, not the password you use to log in to the Orchestrator GUI.

Step 9 Click **OK** to save the changes.

Configuration Backup and Restore

You can create backups of your Multi-Site Orchestrator configuration that can facilitate in recovery from Orchestrator failures or cluster restarts. We recommend creating a backup of the configuration before every upgrade or downgrade of your Orchestrator and after every configuration change or deployment. We also recommend exporting the backups to an external storage outside of the Orchestrator nodes' VMs.

Note Restoring a backup action restores the database on the Multi-Site Orchestrator, but it does not make any changes to the APIC databases on each site. Therefore, after you restore the Orchestrator database, you must also re-deploy any existing schemas to avoid potentially mismatching policies between the Orchestrator and APIC sites. For information on specific configuration mismatch scenarios and backup restore procedures related to each one, see Backup and Restore Guidelines, on page 97

Backup and Restore Guidelines

When saving and restoring configuration backups, the following guidelines apply:

- When saving a backup, the configuration is saved in the same state in which it was deployed. When restoring a backup, any policies that were deployed will show as "deployed", while any policies that were not deployed will remain in the "undeployed" state.
- Restoring a backup action restores the database on the Multi-Site Orchestrator, but it does not make any changes to the APIC databases on each site. As such, certain precautions and steps must be taken when restoring a previous configuration to avoid potentially mismatching policies between the Orchestrator and the APIC sites, as described below.

No Configuration Changes Since Backup

If there have been no policy changes between when the backup was created and when it is being restored, no additional considerations are required and you can simply restore the configuration as described in Restoring Backups, on page 102.

Objects or Policies Created, Modified, or Deleted Since Backup

If any configuration changes took place between the time when the configuration backup was created and the time it is being restored, consider the following:

- Restoring a backup will not modify any objects or policies on the APIC sites. Any new objects or policies created and deployed since the backup will remain deployed. You will need to manually remove these after restoring the backup to avoid any stale configurations.

  Alternatively, you can choose to undeploy all policies first, which will avoid any potential stale objects after the configuration is restored from backup. However, this would cause a disruption in traffic or services defined by those policies.
- The steps required to restore a configuration backup are described in Restoring Backups, on page 102.
- If the configuration backup you restored was saved before it was deployed to the APIC sites, it will be restored in the "undeployed" state and you can simply deploy it to the APIC sites as necessary.
- If the configuration backup you restored was saved when the configuration was already deployed, it will be restored in the "deployed" state, even though none of the policies will exist in the APIC sites yet. In this case, in order for the configuration to be properly pushed to each site, you will need to make a minor configuration change and re-deploy it to sync the Orchestrator's configuration with the APIC sites.

Remote Backups

Cisco ACI Multi-Site is deployed as a 3-node cluster. When you first deploy the cluster, any backups you create are saved to a default location which is located on each node's local disk in the `/opt/cisco/msc/backups/` directory.

While the backups are available on any one node and can be viewed using the Orchestrator GUI, we recommend exporting all backups to a remote location outside the Orchestrator VMs. There are two approaches to configuring remote locations for all Orchestrator backups:

- Configuring a remote NFS share and mounting it to the default backups directory on each node, in which case the backup files are written directly to the remote NFS share bypassing the Orchestrator VMs' local drives.

  This approach is less flexible in that it allows only a single remote location to be used for all configuration backups created from the Orchestrator GUI.

- Configuring a remote SCP or SFTP location using the Orchestrator GUI and then exporting the backup files there.

  Unlike the remote NFS share approach, configuring one or more remote locations in the Orchestrator GUI allows you to specify multiple destinations and provides additional flexibility for where the backup files can be stored.

  Note When you create a configuration backup and export it to a remote server, the files are first created on the Orchestrators' local drives, then uploaded to the remote location, and finally deleted from the local storage. There is a limit on the local backups disk space usage, which if reached can prevent remote backups from being created.

Configuring a Remote Location for Backups

This section describes how to configure a remote location in Multi-Site Orchestrator to which you can then export your configuration backups.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Operations** > **Remote Locations**.

Step 3 In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

Step 4 Provide the name for the location and an optional description.

Two protocols are currently supported for remote export of configuration backups:

- `SCP`
- `SFTP`

Note `SCP` is supported for non-Windows servers only. If your remote location is a Windows server, you must use the `SFTP` protocol

Step 5 Specify the host name or IP address of the remote server.

Based on your **Protocol** selection, the server you specify must allow `SCP` or `SFTP` connections.

Step 6 Provide the full path to a directory on the remote server where you will save the backups.

The path must start with a slash (`/`) characters and must not contain periods (`.`) or backslashes (`\`). For example, *`/backups/multisite`*.

Note The directory must already exist on the remote server.

Step 7 Specify the port used to connect to the remote server.

By default, port is set to `22`.

Step 8 Specify the authentication type used when connecting to the remote server.

You can configure one of the following two authentication methods:

- `Password`—provide the username and password used to log in to the remote server.
- `SSH Private Files`—provide the username and the SSH Key/Passphrase pair used to log in to the remote server.

Step 9 Click **Save** to add the remote server.

Moving Existing Backups to a Remote Location

This section describes how to move an existing configuration backup you have created in the Multi-Site Orchestrator GUI from the nodes' local drives to a remote location.

Before you begin

You must have completed the following:

- Created a configuration backup as described in Creating Backups, on page 101.
- Added a remote location for exporting backups as described in Configuring a Remote Location for Backups, on page 99.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Admin** > **Backups**.

Step 3 Locate the backup you want to export, then click the actions (⋮) icon next to it, then click **Move to remote location**.

A **Move Backup To Remote Location** window opens.

Step 4 From the **Remote Location** dropdown menu, select the remote location.

Step 5 (Optional) Update the remote location path.

The target directory on the remote server, which you configured when creating the remote backup location, will be displayed in the **Remote Path** field.

You can choose to append additional subdirectories to the path. However, the directories must be under the default configured path and must have been already created on the remote server.

Adding an NFS Share to Store Backups

This section describes how to add an NFS share to the Multi-Site Orchestrator VMs to store configuration backups.

Note While you can configure a single remote NFS share for your configuration backups, we recommend using the remote backup location feature available in the Orchestrator GUI and described in Configuring a Remote Location for Backups, on page 99 instead.

Step 1 Log in directly to your Multi-Site Orchestrator node's VM as the `root` user.

Step 2 Mount the NFS share.

The following command mounts the shared NFS directory to the default Orchestrator backups folder so all future backups are automatically stored to an external storage outside the Orchestrator VMs.

Note If you have any existing backups in this default directory that you want to save, you must manually move them to a different location before mounting the NFS share. After the share is mounted, any existing files in the mount directory will be hidden from view.

```
# mount <nfs-server-ip>:/<nfs-share-path> /opt/cisco/msc/backups/
```

Step 3 Repeat steps 1 through 2 on each Orchestrator VM.

Because each Orchestrator node can create and store its own backup files, you must mount the same NFS share on all nodes.

Step 4 Update the Docker backup services.

You must run the following Docker update command for the newly mounted file system to be usable by the Orchestrator services. However, since the command updates the services across the cluster, you only need to do this once after mounting the shares on each node.

```
# docker service update msc_backupservice --force
```

What to do next

If at any point you want to remove the NFS share and go back to storing the backups locally on each VM, simply unmount the directory on each node and run the `docker service update msc_backupservice --force` command again.

Creating Backups

This section describes how to create a new backup of your Multi-Site Orchestrator configuration.

Before you begin

If you want to create the backup using a remote location, you must first add the remote location as described in Configuring a Remote Location for Backups, on page 99.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left navigation pane, select **Admin** > **Backups**.

Step 3 In the main window, click **New Backup**.

A **New Backup** window opens.

Step 4 In the **Name** field, provide the name for the backup file.

The name can contain up to 10 alphanumeric characters, but no spaces or underscores (_).

Step 5 (Optional) In the **Notes** field, enter any additional information to describe the backup.

Step 6 Choose the **Backup Location**.

You can save the backup file locally on the Orchestrator nodes or export it to a remote location.

If you want to save the backup file locally, choose **Local**.

Otherwise, if you want to save the backup file to a remote location, choose **Remote** and provide the following:

- From the **Remote Location** dropdown menu, select the remote location.
- In the **Remote Path**, either leave the default target directory or you can choose to append additional subdirectories to the path. However, the directories must be under the default configured path and must have been already created on the remote server.

Step 7 Click **Save** to create the backup.

Restoring Backups

This section describes how to restore a Multi-Site Orchestrator configuration to a previous state.

Before you begin

Restoring a backup action restores the database on the Multi-Site Orchestrator, but it does not make any changes to the APIC databases on each site. Therefore, after you restore the Orchestrator database, you must also re-deploy any existing schemas to avoid potentially mismatching policies between the Orchestrator and APIC sites.

For information on specific configuration mismatch scenarios and recommended restore procedures related to each one, see Backup and Restore Guidelines, on page 97.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 If necessary, undeploy existing policies.

We recommend you perform this step if new objects or policies were added to the configuration between when the backup was created and current configuration. Additional context is available in Backup and Restore Guidelines, on page 97.

Step 3 From the left navigation menu, select **Admin** > **Backups**.

Step 4 In the main window, click the actions (⋮) icon next to the backup you want to restore and select **Rollback to this backup**.

If the version of the selected backup is different from the running Multi-Site version, the rollback could cause a removal of the features that are not present in the backup version.

Step 5 Click **Yes** to confirm that you want to restore the backup you selected.

If you click **Yes**, the system terminates the current session and the user is logged out.

Step 6 If necessary, redeploy the configuration.

We recommend you perform this step to sync the restored configuration with the APIC sites. Additional context is available in Backup and Restore Guidelines, on page 97.

Downloading Backups

This section describes how to download you backup from the Multi-Site Orchestrator.

Before you begin

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Admin** > **Backups**.

Step 3 In the main window, click the actions (⋮) icon next to the backup you want to download and select **Download**.

This will download the backup file in `msc-backups-<timestamp>.tar.gz` format to your system. You can then extract the file to view its contents.

Importing Backups

This section describes how to import an existing backup into your Multi-Site Orchestrator.

Before you begin

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Admin** > **Backups**.

Step 3 In the main window, click **Import**.

Step 4 In the **Import from file** window that opens, click **Select File** and choose the backup file you want to import.

Importing a backup will add it to the list of the backups displayed the **Backups** page.

Custom SSL Certificates

Cisco ACI Multi-Site Orchestrator OVA contains a self-signed SSL certificate that is stored in `/data/msc/secrets` directory on each node during the Orchestrator installation. By default, the Orchestrator GUI uses this certificate for its HTTPS connections.

While you could previously update these certificates by logging directly into an Orchestrator node server and changing its web server (`nginx`) configuration, starting with Cisco ACI Multi-Site Orchestrator Release 2.1(1), you can use the GUI to easily add or update custom certificates to be used for the Orchestrator's GUI connection.

When adding custom certificates, you can use one of the following two options:

- **Self-Signed Certificate** provide you with the ability to create your own public and private keys to be used by the Orchestrator's GUI.
- **CA-Issued Certificate** allows you to use a certificate provided by an existing Certificate Authority (CA) along with its keys.

You can add multiple CAs and Keyrings containing the public/private key combinations in the GUI, however only a single keyring can be active at any given time and used to secure the communication between the Orchestrator GUI and your browser.

Adding Custom Certificate Authority

You can add a custom Certificate Authority (CA) to be used for verifying the public key provided by the Orchestrator for HTTPS traffic encryption.

This section describes how to add and configure a custom CA in Multi-Site Orchestrator GUI. Configuring keyrings and keys is described in the next section.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left navigation menu, select **Admin** > **Security**.

Step 3 In the main window, select the **Certificate Authority** tab and click **Add Certificate Authority**.

Step 4 In the **Add Certificate Authority** window that opens, provide the CA details.

In the **Name** field, enter the CA name.

In the **Description** field, enter the CA description.

In the **Certificate Chain** field, enter the CA's certificate chain. You must include both, intermediate and root, certificates. The intermediate certificate must be entered first, followed by the root certificate.

Step 5 Click **SAVE** to save the changes.

Adding Custom Keyring

You can add a custom keyring containing a public and private encryption keys to be used for Orchestrator GUI HTTPS traffic encryption.

This section describes how to add a custom keyring. For instructions on adding a Certificate Authority (CA) that can be used to verify the public key in this keyring, see the previous section.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left-hand navigation menu, select **Admin** > **Security**.

Step 3 In the main window, select the **Key Rings** tab and click **ADD KEY RING**.

Step 4 In the **Create Key Ring** window that opens, provide the key ring details.

From the **SELECT CERTIFICATE AUTHORITY** dropdown menu, select the certificate authority that will contain the key ring.

In the **NAME** field, enter the key ring name.

In the **KEY RING DESCRIPTION** field, enter the key ring description.

In the **PUBLIC KEY** field, enter the ring's public key.

In the **PRIVATE KEY** field, enter the ring's private key

Step 5 Click **SAVE** to save the changes.

Activating Custom Keyring

After you add a keyring, as described in previous section, you need to activate it as the default keyring.

Step 1 Log in to your Multi-Site Orchestrator GUI.

Step 2 From the left-hand navigation menu, select **Admin** > **Security**.

Step 3 In the main window, select the **Key Rings** tab.

Step 4 In the main window, click the **...** icon next to the keyring you want to activate and choose **Make Keyring Active**.

Step 5 Click **ACTIVATE** to activate the keyring.

Activating a key will log you out of the Multi-Site Orchestrator GUI. When the login page is loaded, it will use the new certificate and key.

Custom Certificates Troubleshooting

The following sections describe how to resolve common issues when using custom SSL certificates with Multi-Site Orchestrator.

Unable to Load the Orchestrator GUI

If you are unable to load the Orchestrator GUI page after installing and activating a custom certificate, it is possible that the certificates were not copied correctly to each Orchestrator node. You can resolve this issue by recovering the default certificates and then repeating the new certificate installation procedure again.

To recover the default Orchestrator certificates:

1. Log in to each Orchestrator node directly.

2. Change into the certificates directory:

```
# cd /data/msc/secrets
```

3. Replace the msc.key and msc.cert files with msc.key_backup and msc.cert_backup files respectively.

```
# cp msc.key_backup msc.key
# cp msc.cert_backup msc.cert
```

4. Restart the Orchestrator GUI service

```
# docker service update msc_ui --force
```

5. Re-install and activate the new certificates as described in previous sections.

Adding a New Orchestrator Node to the Cluster

If you add a new node to you Multi-Site Orchestrator cluster:

1. Log in to the Orchestrator GUI.

2. Re-activate the key you are using as described in previous sections.

External Authentication

You can configure external user authentication and authorization using RADIUS, TACACS+, and LDAP servers.

As a Multi-Site Orchestrator administrator, you can:

- Add one or more external authentication providers.

  It is recommended to set up at least 2 authentication providers for redundancy.

- Create login domains and associate them with providers.

  The default domain is the Local domain, for local authentication.

- Assign users to domains.

After you create domains, you can edit, deactivate, or delete them. You cannot delete the Local domain, but you can deactivate it.

Audit logs support external authentication and authorization.

Guidelines for Configuring External Authentication Servers

When configuring external authentication servers for Multi-Site Orchestrator user authentication:

- You must configure each user on the remote authentication servers.
- For both local and external authentication, the username supports a maximum length of 20 characters.
- For each user, you must add a custom attribute-value (AV) pair, specifying the use roles assigned to that user. The roles are documented in Users, Roles, and Permissions, on page 57.

  When specifying the roles, use the following format:

  ```
  cisco-av-pair=shell:msc-roles=role1,role2
  ```

  For example:

  ```
  cisco-av-pair=shell:msc-roles=siteManager,schemaManager.
  ```

- Starting with Release 2.1(2), each of the user roles 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.

  The AV pair string format differs when configuring a read-only or a combination of read-write and read-only roles for a specific user. In the following example, the read-write roles are separated from the read-only roles using the slash (`/`) character, while the individual roles are separated by the pipe (`|`) character:

  ```
  cisco-av-pair=shell:msc-roles=writeRole1|writeRole2/readRole1|readRole2
  ```

  The following example illustrates how to assign the Schema Manager and User Manager roles to a user, while still allowing them to see objects visible to the Site Manager users:

  ```
  shell:msc-roles=schemaManager|userManager/siteManager
  ```

If you want to configure only either the read-only or read-write permissions for a user, you must still include the slash (/) character. The following examples show how to set just the read-write or read-only access to the objects available to Site Manager role.

- Read-only: `shell:msc-roles=/siteManager`
- Read-write: `shell:msc-roles=siteManager/`

Note While either the old (comma-separated) or the new (pipes and a slash) format is supported, you cannot mix them when configuring a single user. Mixed or incorrectly formatted AV strings are not parsed and the user roles are not configured.

- If you configure any read-only user roles and then downgrade your Multi-Site Orchestrator to an earlier version, which does not support read-only permissions, those roles will be removed from all users. This also means that any user that has **only** the read-only roles will have no roles assigned to them and be deleted. A Power User or User Manager will need to recreate the users and re-assign them new read-write roles.
- For LDAP configurations, we recommend using **`CiscoAVPair`** as the attribute string. If, for any reason, you are unable to use an Object ID `1.3.6.1.4.1.9.22.1`, an additional Object IDs `1.3.6.1.4.1.9.2742.1-5` can also be used in the LDAP server.

Adding RADIUS or TACACS+ as Authentication Provider

This section describes how to add one or more RADIUS or TACACS+ servers as external authentication servers for authenticating Cisco ACI Multi-Site Orchestrator users.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator as the `admin` user using the Local domain.

Step 2 From the left-hand navigation pane, select **Admin** > **Providers**.

Step 3 In the main window, click **ADD PROVIDER**.

Step 4 Enter the host name or IP address of the external authentication server.

Step 5 (Optional) Enter a description for the provider you are adding.

Step 6 Select **RADIUS** or **TACACS+** for the provider type you are adding.

Step 7 Enter the **KEY** and confirm it in the **CONFIRM KEY** field.

Step 8 (Optional). Configure additional settings.

a) Expand **Additional Settings** for more settings.

b) You can specify the port used to connect to the authentication server.

The default port is `1812` for **RADIUS** and `49` for **TACACS+**.

c) You can specify the protocol used.

You can choose between **PAP** or **CHAP** protocols.

d) You can specify the timeout and number of attempts for connecting to the authentication server.

Adding LDAP as Authentication Provider

This section describes how to add one or more LDAP servers as external authentication servers for Cisco ACI Multi-Site Orchestrator users.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator as the `admin` user using the Local domain.

Step 2 From the left-hand navigation pane, select **Admin** > **Providers**.

Step 3 In the main window, click **Add Provider**.

Step 4 Enter the host name or IP address of the external authentication server.

Step 5 (Optional) Enter a description for the provider you are adding.

Step 6 Select **LDAP** for the provider type you are adding.

Step 7 Enter the **Base DN**, **Bind DN**, and the **Key** values for the LDAP server.

The Base DN and Bind DN dependent on how your LDAP server is configured. You can get the Base DN and Bind DN values from the distinguished name of the user created on the LDAP server.

Base DN is the point from which the server will search for users. For example, `DC=mso,DC=local`.

Bind DN is the credentials used to authenticate against the server. For example, `CN=admin, CN=Users,DC=mso,DC=local`.

Bind DN comes with a key, which you can provide in the next field.

Step 8 (Optional) Enable SSL for LDAP communication.

a) Check the **Enable SSL** checkbox.
b) Select the certificate you want to use.
c) Select the validation level.

Permissive: Accept a certificate signed by any certificate authority (CA) and use it for encryption.

Restrictive: Verify the entire certificate chain before using it.

Step 9 (Optional). Configure additional settings.

a) Click **Additional Settings** to expand.
b) Specify the port used to connect to the LDAP server.

The default port for **LDAP** is `389`.

c) Specify the timeout and number of attempts for connecting to the authentication server.
d) Specify the filter used.

The filter value depends on the LDAP server configuration. The default LDAP filter is `(cn=username)`. However, if you're using a Microsoft LDAP server, set the filter to `(sAMAccountName={username})` instead.

e) Specify the authentication type.

The authentication type can be:

- **Cisco-AVPair** – uses an attribute-value (AV) pair to configure authorization based on individual user's role. When using this method, set the **Attribute** field to `ciscoAVPair`.

  You must also configure each user individually in your LDAP server using the AV pair string in the following format:

  - Release 2.1(2) and later:

`cisco-av-pair=shell:msc-roles=`*`writeRole1|writeRole2/readRole1|readRole2`*

- Release 2.1(1) and earlier:

  `cisco-av-pair=shell:msc-roles=`*`role1,role2`*

For additional information, see Guidelines for Configuring External Authentication Servers, on page 106.

- **LDAP Group Map Rules** - use an LDAP server group to configure authorization based on the users' group membership. When using this method, set the **Attribute** field to `memberOf`, then click **+LDAP Group Map Rules** to specify the group membership.

  In the **New Group Map Rule**, specify the group DN (for example, `CN=group1,OU=msc-ou,DC=msc,DC=local`) and the user roles to be assigned to that group. You can add multiple roles for the same group map rule. Detailed descriptions of each user role are available in Users, Roles, and Permissions, on page 57.

Creating Login Domains

A login domain defines the authentication domain for a user. Login domains can be set to the Local, RADIUS, TACACS+, or LDAP authentication mechanisms.

When you are logging in to the Cisco ACI Multi-Site Orchestrator using the GUI, the login screen offers a drop-down list of domains for the user to select from. If no domain is specified, the Local domain is used to look up the username.

When you are logging in to the Cisco ACI Multi-Site Orchestrator using the REST API, the login domain is provided along with the login information in the `POST` message, for example:

```
{
    "username":"bob",
    "password":"Welcome2msc!",
    "domainId":"59d5b5978d0000d000909f65"
}
```

To create a login domain using the Cisco ACI Multi-Site Orchestrator GUI:

Before you begin

You must have added one or more authentication providers as described in Adding RADIUS or TACACS+ as Authentication Provider, on page 107 or Adding LDAP as Authentication Provider, on page 108.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left-hand navigation pane, select **Admin** > **Login Domains**.

Step 3 In the main window, click **ADD LOGIN DOMAIN**.

Step 4 Enter the domain's name.

Step 5 (Optional) Enter a description for the domain.

Step 6 Select **REALM** type to specify the authentication provider.

You must have an external authentication provider added before creating login domains.

Step 7 Assign the login domain to one or more providers.

Mark the checkbox next to one or more providers' names to assign the domain.

What to do next

After you create one or more login domains, you can edit, delete, or deactivate them as described in Editing, Deleting, or Deactivating Login Domains, on page 110.

Editing, Deleting, or Deactivating Login Domains

After you have created one or more login domains, you can use the instruction described in this section to edit, delete, or deactivate them. You cannot delete the Local domain, but you can deactivate it.

Before you begin

You must have created one or more Login domains as described in Creating Login Domains, on page 109.

Step 1 Log in to your Cisco ACI Multi-Site Orchestrator.

Step 2 From the left-hand navigation pane, select **Admin** > **Login Domains**.

Step 3 Click the **...** menu next to the login domain you want to edit.

You can choose to **Edit** the domain information, **Deactivate** the domain so that it cannot be used, or **Set as default** so it is automatically selected when logging in using GUI.

Remote User Logon

When external authentication is enabled in Cisco ACI Multi-Site, you can log in to the Multi-Site Orchestrator as follows:

Step 1 Using a browser, navigate to the Multi-Site URL.

Step 2 Choose your assigned domain from the drop down list.

Step 3 Enter your username and password.

Step 4 Click **Submit**.
If you are authorized and pass authentication, the Multi-Site Orchestrator GUI is displayed and you have privileges according to the roles that are assigned to you. The first time you log on, you will be prompted to change your password.

System Configuration Settings

There is a number of global system settings that are available under **Admin** > **System Configuration**, which you can configure for your Multi-Site Orchestrator as described in the following sections.

System Alias and Banner

This section describes how to configure an alias for your Multi-Site Orchestrator as well as enable a custom GUI-wide banner to be displayed at the top of your screen, as shown in the following figure.

Figure 22: System Banner Display

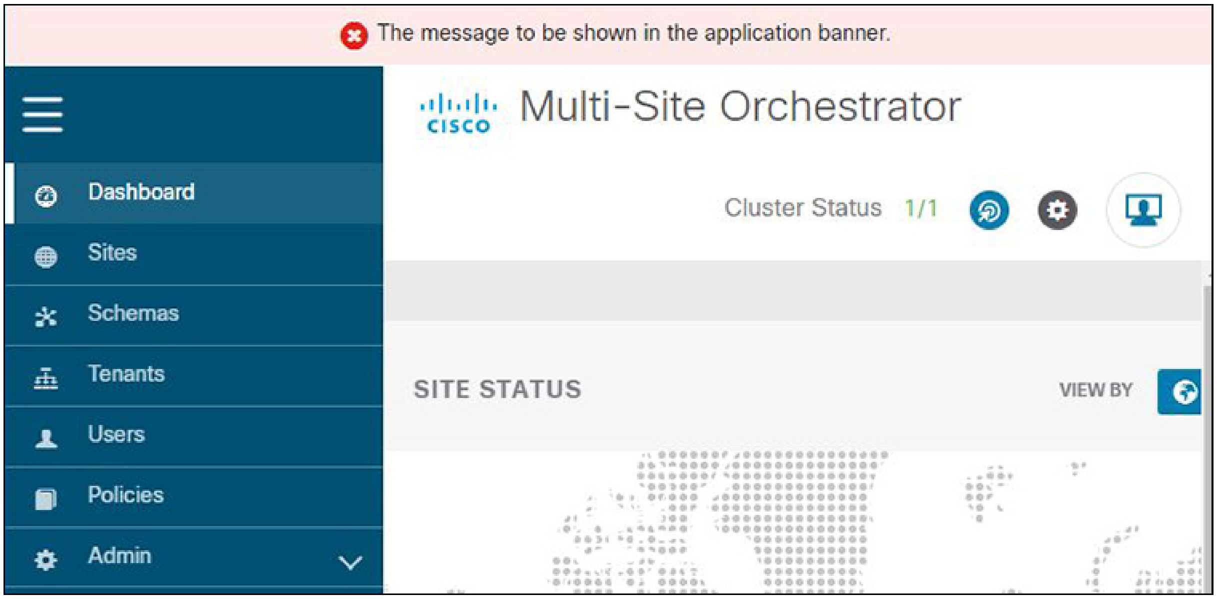

The message to be shown in the application banner.

Step 1 Log in to your Orchestrator.

Step 2 From the left navigation pane, select **Admin** > **System Configuration**.

Step 3 Click the **Edit** icon to the right of the **System Alias & Banners** area.

This opens the **System Alias & Banners** settings window.

Step 4 In the **Alias** field, specify the system alias.

Step 5 Choose whether you want to enable the GUI banner.

Step 6 If you enable the banner, you must provide the message that will be displayed on it.

Step 7 If you enable the banner, you must choose the severity, or color, for the banner.

Step 8 Click **Save** to save the changes.

Login Attempts and Lockout Time

When the Orchestrator detects a significant number of failed consecutive login attempts, the user is locked out of the system to prevent unauthorized access. You can configure how failed log in attempts are treated, for example the number of failed attempts before lockout and the length of the lockout.

Note This feature is enabled by default when you first install or upgrade to Release 2.2(1) or later.

Step 1 Log in to your Orchestrator.

Step 2 From the left navigation pane, select **Admin** > **System Configuration**.

Step 3 Click the **Edit** icon to the right of the **Fail Attempts & Lockout Time** area.

This opens the **Fail Attempts & Lockout Time** settings window.

Step 4 From the **Fail Attempt Settings** dropdown, select the number of attempts before the user is locked out.

Step 5 From the **Lockout Time (Minutes)** dropdown, select the length of the lockout.

This specifies the base lockout duration once it's triggered. The timer is extended up to three times exponentially with every additional consecutive login failure.

Step 6 Click **Save** to save the changes.

Proxy Server

In certain deployment scenarios, such as with a combination of on-premises and cloud sites and the Orchestrator running inside a corporate network, the Orchestrator may have to access the internet and the cloud sites through a proxy. You can configure and enable proxy as described in this section.

When a proxy server is enabled, the Orchestrator will maintain a "no proxy" list of IP addresses and hostnames with which it will communicate directly bypassing the proxy. This list is a combination of user-specified hosts or domains plus all on-premises APIC sites currently added to the Orchestrator. Every time the list is updated with a new address, for example if you add a new site to the Orchestrator, the proxy service is restarted. You can minimize the service restarts by providing a complete list of your on-premises sites in advance, for example by adding an entire domain to the "no proxy" list, while configuring the proxy settings.

Step 1 Log in to your Orchestrator.

Step 2 From the left navigation pane, select **Admin** > **System Configuration**.

Step 3 Click the **Edit** icon to the right of the **Proxy Server** area.

This opens the **Proxy Settings** window.

Step 4 Choose **Enable** to enable the proxy.

Step 5 In the **Proxy Server** field, specify the IP address or the hostname of your proxy server.

Step 6 In the **Proxy Server Port** field, specify the port number used to connect to the proxy server.

Step 7 In the **No Proxy List** field, provide a comma-separated list of hosts and domains that should bypass the proxy.

When specifying the list, you can provide exact IP addresses or hostnames, as well as entire domains using the wildcard (`*`) character. Wildcards cannot be used with IP addresses.

For example, `203.0.113.1, apic1.example.com, *.example.local`.

Step 8 Click **Save** to save the changes.

When you configure and enable proxy, the Orchestrator application will restart.