

Cisco ACI Multi-Site Orchestrator Installation and Upgrade Guide, Release 3.1(x)

First Published: 2020-05-11

Last Modified: 2020-11-25

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.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2020–2021 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

PART I **Cluster Deployments** 3

CHAPTER 2 **Deployment Overview** 5

Deployment Options 5

Multi-Site Orchestrator Communication Ports 6

Multi-Site Orchestrator and Cisco APIC Interoperability Support 7

CHAPTER 3 **Deploying in Cisco Application Services Engine** 9

Prerequisites and Guidelines 9

Deploying Multi-Site Orchestrator in Application Services Engine 11

Migrating Existing Cluster to Application Service Engine 14

CHAPTER 4 **Deploying in VMware ESX** 17

Prerequisites and Guidelines 17

Deploying Cisco ACI Multi-Site Orchestrator Using Python 19

Setting Up Python Environment 19

Sample Deployment Configuration File 21

Deploying Multi-Site Orchestrator Using Python 23

Deploying Orchestrator in vCenter 25

Deploying Orchestrator in ESX Directly 29

PART II **Day-0 Operations** 33

CHAPTER 5 Configuring and Adding Sites 35
Pod Profile and Policy Group 35
Configuring Fabric Access Policies for All APIC Sites 35
Configuring Fabric Access Global Policies 35
Configuring Fabric Access Interface Policies 37
Configuring Sites That Contain Remote Leaf Switches 38
Multi-Site and Remote Leaf Guidelines and Limitations 38
Configuring Routable Subnets for Remote Leaf Switches 39
Enabling Direct Communication for Remote Leaf Switches 39
Cisco Mini ACI Fabrics 40
Adding Sites 41

CHAPTER 6 Configuring Infra 43
Configuring Infra Prerequisites and Guidelines 43
Configuring Infra: General Settings 44
Refreshing Site Connectivity Information 44
Configuring Infra: On-Premises Site Settings 45
Configuring Infra: Cloud Site Settings 47
Configuring Infra: Pod Settings 47
Configuring Infra: Spine Switches 48
Configuring Infra: MPLS L3Out Settings 49
SR-MPLS Infra Guidelines and Limitations 50
Creating SR-MPLS QoS Policy 52
Creating SR-MPLS Infra L3Out 53
Deploying Infra Configuration 56
Enabling Connectivity Between On-Premises and Cloud Sites 57

CHAPTER 7 CloudSec Encryption 63
Cisco ACI CloudSec Encryption 63
Requirements and Guidelines 64
CloudSec Encryption Terminology 66
CloudSec Encryption and Decryption Handling 67
CloudSec Encryption Key Allocation and Distribution 68

Configuring Cisco APIC for CloudSec Encryption 70
Configuring Cisco APIC for CloudSec Encryption Using GUI 71
Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI 71
Configuring Cisco APIC for CloudSec Encryption Using REST API 72
Enabling CloudSec Encryption Using Cisco ACI Multi-Site Orchestrator GUI 73
Rekey Process During Spine Switch Maintenance 74
Disabling and Re-Enabling Re-Key Process Using NX-OS Style CLI 74
Disabling and Re-Enabling Re-Key Process Using REST API 75

CHAPTER 8 **Adding Tenants and Schemas 77**
Adding Tenants 77
Adding Schemas 78

PART III **Cluster Upgrades and Downgrades 81**

CHAPTER 9 **Upgrading or Downgrading Orchestrator Deployments in Application Service Engine 83**
Prerequisites and Guidelines 83
Multi-Site Orchestrator and Cisco APIC Interoperability Support 84
Upgrading Multi-Site Orchestrator, Release 3.0(2) or Later 85
Upgrading Multi-Site Orchestrator, Release 3.0(1) or Earlier 87

CHAPTER 10 **Upgrading Orchestrator Deployments in VMware ESX 91**
Prerequisites and Guidelines 91
Multi-Site Orchestrator and Cisco APIC Interoperability Support 92
Upgrading Cisco ACI Multi-Site Orchestrator Using Python 93
Setting Up Python Environment 93
Sample Upgrade Configuration File 95
Upgrading Multi-Site Orchestrator 96
Upgrading Cisco ACI Multi-Site Orchestrator Using Backup/Restore 98

CHAPTER 11 **Downgrading Orchestrator Deployments in VMware ESX 103**
Downgrading Guidelines and Limitations 103
Downgrading Multi-Site Orchestrator 104

PART IV Single Node Deployments 107

CHAPTER 12 Single Node Overview 109

Overview 109

CHAPTER 13 Installing Single Node Orchestrator 111

Installing Single Node Orchestrator in VMware ESX 111

Installing Single Node Orchestrator in Service Engine 113

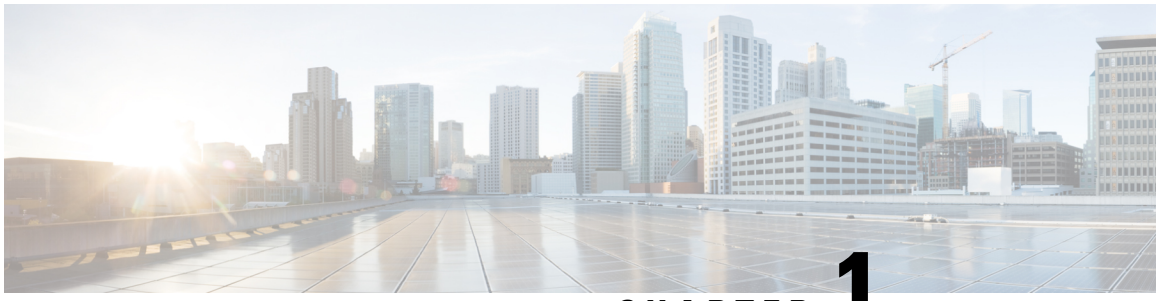
CHAPTER 14 Upgrading Single Node Orchestrator 117

Upgrading Single Node ESX VM 117

Upgrading Single Node Service Engine VM 119

CHAPTER 15 Converting to Production Cluster 123

Converting Single Node to Production Cluster 123



CHAPTER 1

New and Changed Information

This chapter contains the following sections:

- New and Changed Information, on page 1

New and Changed Information

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

Table 1: Latest Updates

Release	New Feature or Update	Where Documented
3.0(2)	Simplified deployment and upgrades of Multi-Site Orchestrator on Cisco Application Services Engine.	For more information, see Deploying in Cisco Application Services Engine, on page 9 or Upgrading or Downgrading Orchestrator Deployments in Application Service Engine, on page 83.
3.0(1)	First release of this document.	--

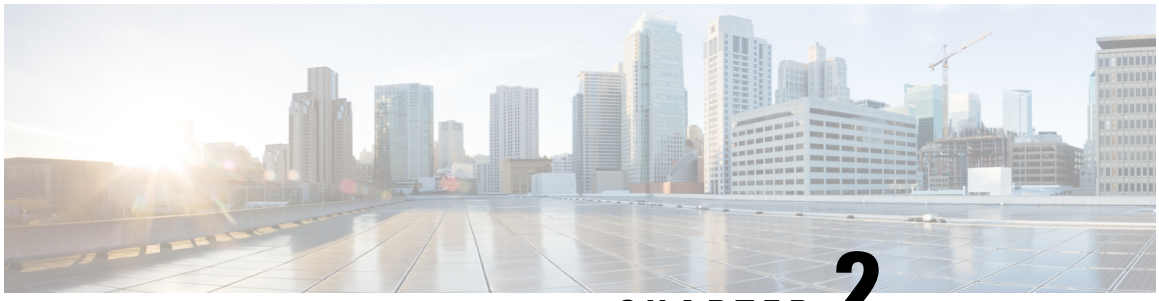
New and Changed Information



PART I

Cluster Deployments

- Deployment Overview, on page 5
- Deploying in Cisco Application Services Engine, on page 9
- Deploying in VMware ESX, on page 17



CHAPTER 2

Deployment Overview

This chapter contains the following sections:

- Deployment Options, on page 5
- Multi-Site Orchestrator Communication Ports, on page 6
- Multi-Site Orchestrator and Cisco APIC Interoperability Support, on page 7

Deployment Options

A typical Cisco ACI Multi-Site deployment requires a 3-node Multi-Site Orchestrator cluster to manage all the sites' fabrics in your ACI Multi-Site environment. You can choose to deploy the Orchestrator cluster in one of the following ways:

- You can deploy the Multi-Site Orchestrator (MSO) cluster in a Cisco Application Services Engine (ASE).

  We recommend this approach for all new ACI Multi-Site deployments, because it provides a common platform to streamline multi-product integrations, additional security through Cisco Secured Development Lifecyle (CSDL) and removal of `root` access to Orchestrator application, file system protection, and single click upgrades to future releases.

  Cisco Application Service Engine itself can be deployed using a number of different form factors, such as a Cisco Application Service physical appliance (`.iso`), in a VMware ESX virtual machine (`.ova`), in Amazon Web Services (`.ami`), or in Linux KVM (`.qcow`), all of which are supported for Multi-Site Orchestrator installations. Keep in mind however, you must use the same form factor Service Engine for all Orchestrator nodes, mixing different form factors within the same Orchestrator cluster is not supported. Installing and configuring the Application Service Engine is outside the scope of this document and is described in *Cisco Application Services Engine User Guide*.

  Installing and configuring the Orchestrator cluster in Application Services Engine is described in the Deploying in Cisco Application Services Engine, on page 9. Upgrading Services Engine MSO deployments is described in the Upgrading or Downgrading Orchestrator Deployments in Application Service Engine, on page 83.

- Alternatively, you can deploy each Orchestrator node directly in VMware ESX VMs.

  When deploying in ESX VMs, you can choose one of the following 2 approaches:

  - Use Cisco-provided Python scripts to deploy the entire Multi-Site Orchestrator cluster. The scripts allow you to execute the deployment and later upgrades remotely, for example from your laptop, as long as you have access to the vCenter where the Orchestrator VMs are to be deployed.

This is the preferred approach when deploying an Orchestrator cluster in ESX VMs as it automates a number of manual steps and allows remote execution of Cisco ACI Multi-Site Orchestrator installation and subsequent software upgrades.

- Using an OVA image to deploy each Orchestrator VM individually. In this case you can also choose to deploy the image either using the vCenter or directly on the ESX server.

Installing and configuring the Orchestrator cluster in VMware ESX VMs is described in the Deploying in VMware ESX, on page 17. Upgrading VMware ESX Orchestrator deployments is described in the Upgrading Orchestrator Deployments in VMware ESX , on page 91.

Single Node Lab Deployments

While production Multi-Site Orchestrator deployments require a 3-node high availability (HA) cluster, single node Orchestrator deployments are supported for lab and testing purposes. The installation and upgrade steps for single node Orchestrator differ slightly from the 3-node cluster deployments and are covered in detail in the Installing Single Node Orchestrator, on page 111.

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 443 (`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 443 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

Multi-Site Orchestrator and Cisco APIC Interoperability Support

Multi-Site Orchestrator (MSO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the MSO itself can be upgraded independently of each other and run in mixed operation mode as long as each fabric is running APIC Release 3.2(6) or later. As such, we recommend that you always upgrade to the latest release of the Multi-Site Orchestrator.

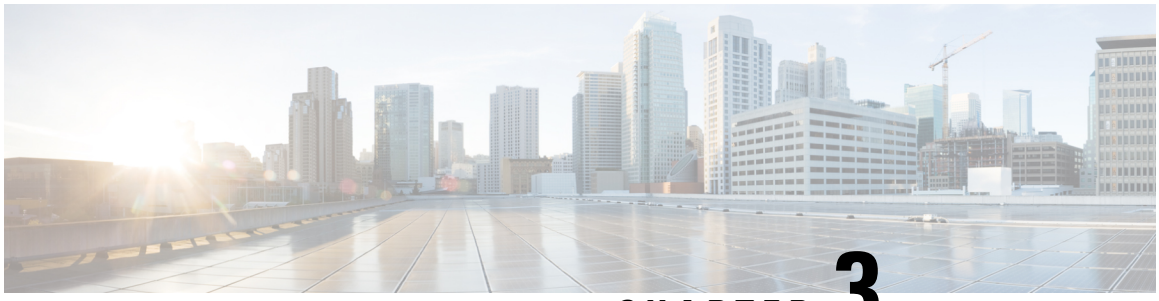
However, keep in mind that if you upgrade the MSO before upgrading the APIC clusters in one or more sites, some of the new MSO features may not yet be supported by an earlier APIC release. In that case a check is performed on each template to ensure that every configured option is supported by the target 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 that 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)
vzAny	Release 3.2(6)
Host Based Routing	Release 4.0(1)
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)
Intersite L3Out	Release 4.2(1)



CHAPTER 3

Deploying in Cisco Application Services Engine

This chapter contains the following sections:

- Prerequisites and Guidelines, on page 9
- Deploying Multi-Site Orchestrator in Application Services Engine, on page 11
- Migrating Existing Cluster to Application Service Engine, on page 14

Prerequisites and Guidelines

This chapter covers production deployment of a 3-node Multi-Site Orchestrator cluster. If you want to set up a single-node Multi-Site Orchestrator (for example, for testing purposes), follow the instruction in the Installing Single Node Orchestrator, on page 111 chapter instead.

Application Services Engine

You must have Cisco Application Services Engine installed and the cluster configured in Fabric External Mode as described in *Cisco Application Services Engine Deployment Guide*.

Cisco Application Service Engine itself can be deployed using a number of different form factors, such as a Cisco Application Service physical appliance (`.iso`), in a VMware ESX virtual machine (`.ova`), in Amazon Web Services (`.ami`), or in Linux KVM (`.qcow`), all of which are supported for Multi-Site Orchestrator installations. Keep in mind however, you must use the same form factor Service Engine for all Orchestrator nodes, mixing different form factors within the same Orchestrator cluster is not supported.

Note AWS deployments of Multi-Site Orchestrator support managing only public cloud sites and not on-premises fabrics.

If you are deploying Services Engine in AWS, by default only PEM-based login is enabled for each node. If you'd like to be able to SSH into the nodes using a password, you will need to explicitly enable password-based logins. You can do that by logging into each node separately using the PEM file the first time, then executing the `# acidiag loginprompt enable` command. After you run the command, you will be logged out and be able to log back in using either the PEM file or password.

Application Services Engine Networks

When first configuring Application Services Engine, two of the parameters that you provide are the Data Network and the Management Network. The data network is used for the nodes' clustering and Cisco ACI

sites traffic. The management network is used to connect to the Cisco Application Services Engine GUI, CLI, or API. You can choose to connect the two interfaces to the same subnet and network, like the Out-of-Band network.

When Multi-Site Orchestrator app is deployed in Application Services Engine, it uses each of the two networks for different purposes as shown in the following table:

Traffic type	Network
• Cisco APIC • Cloud APIC • Any other remote devices or controllers	Data network
Intra-cluster communication	Data network
Audit log streaming (Splunk/syslog)	Management network
Remote authentication (LDAP/TACACS)	Management network
Remote backup	Management network

Network Time Protocol (NTP)

Multi-Site Orchestrator uses NTP for clock synchronization, so you must have an NTP server configured in your environment.

Deployment Requirements

The following table summarizes the Application Services Engine requirements for Cisco ACI Multi-Site Orchestrator.

Orchestrator Version	Requirements
Release 3.0(2) and later* *We do not recommend deploying Release 3.0(1)	Cisco Application Services Engine, Release 1.1.3d. If the Application Services Engine is deployed in an ESX or KVM virtual machine, the following additional requirements apply: • For VMware ESX: ESXi 6.0 or later • For Linux KVM: Linux Kernel `3.10.0-957.el7.x86_64` or later with KVM `libvirt-4.5.0-23.el7_7.1.x86_64` or later • 16 vCPUs 10 GHz CPU reservation is applied automatically • 48 GB of RAM 36 GB reservation is applied automatically • 100 GB disk We recommend thin provisioning with a maximum size of 620 GB with each Application Services Engine VM running on its own disk. • We recommend that each Multi-Site Orchestrator node is deployed in a different ESX or KVM server.

Deploying Multi-Site Orchestrator in Application Services Engine

This section describes how to install Cisco ACI Multi-Site Orchestrator application in Cisco Application Services Engine cluster that was deployed previously. This scenario applies if you have purchased the Cisco Application Services Engine cluster or deployed it separately from an ISO image.

Before you begin

- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 9.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator application.

You can download the required MSO image in one of the following two ways:

a) Browse to the Software Download link:

https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.
c) From the left sidebar, choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site App Image* file (`Cisco-MSO-<version>.aci`) for the release.

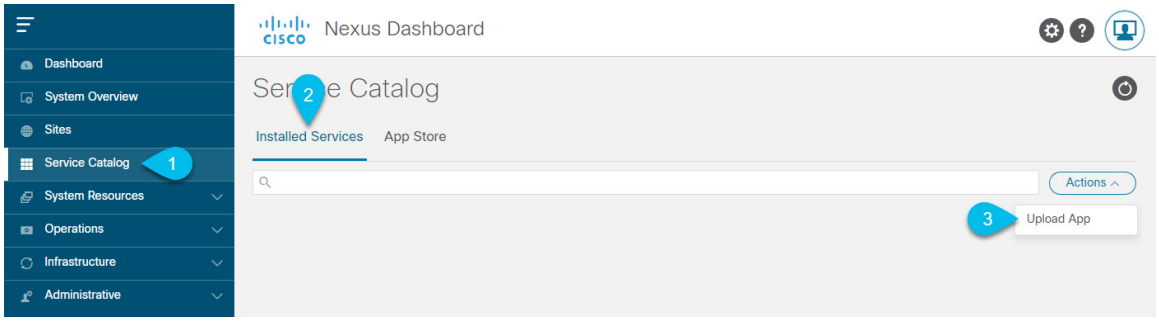
Alternatively, you can download the image from the Cisco DC App Center:

a) Browse to the Multi-Site Orchestrator app page on DC App Center:
https://dcappcenter.cisco.com/multi-site-orchestrator.html
b) From the left sidebar **Version** dropdown, choose the Cisco ACI Multi-Site Orchestrator release version.
c) Click the **Download** button.
d) Click **Agree and download** to accept the license agreement and download the image.

Step 2 Log in to your Cisco Application Services Engine dashboard.

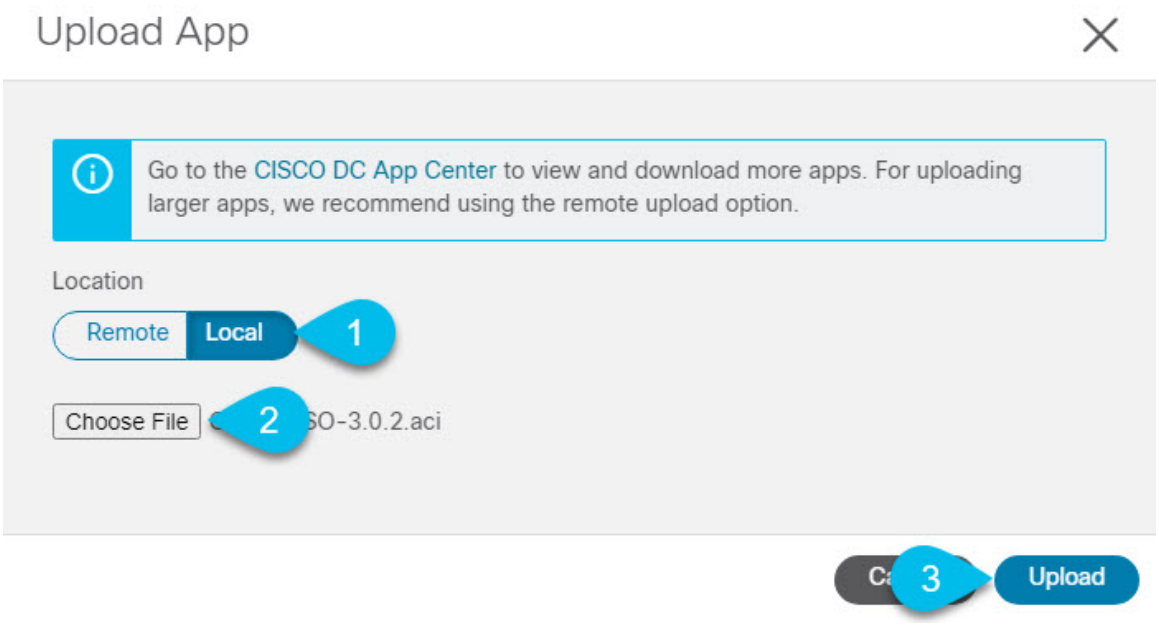
When deploying an app, you need to install it in only one of the Application Services Engine nodes, the application will be replicated to the other nodes in the cluster automatically. So you can log in to any one of your Application Services Engine nodes using its management IP address, for example `https://10.23.237.160`.

Step 3 Add the app.



a) In the left navigation bar, click **Apps**.
b) In the main pane, click **Actions**.
c) Choose **Upload App**.

Step 4 Upload the image file to the Application Services Engine cluster.



a) Choose the location of the image.

If you downloaded the application image to your system, choose **Local**.

If you are hosting the image on a server, choose **Remote**.

b) Choose the file.

If you chose **Local** in the previous step, click **Select File** and select the `Cisco-MSO-<version>.aci` you downloaded.

If you chose **Remote**, provide the full URL to the image file, for example `http://<ip-address>:<port>/<full-path>/Cisco-MSO-<version>.aci`.

c) Click **Upload** to add the app to the cluster.

Step 5 Enable the app.

After installation is complete, the application will remain in the `Disabled` state by default and you must enable it.

To enable the app, click the **...** menu on the app and select **Enable**.

Step 6 Launch the app.

It may take up to 20 minutes for the application to replicate to all nodes and all services to fully deploy.

To launch the app, simply click **Launch App**.

Step 7 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

After you launch the app from the Application Services Engine dashboard, you will be prompted to log in to your new Multi-Site Orchestrator cluster.

The default log in is **admin** and the default password is **We1come2msc!**.

When you first log in, you will be prompted to change the password.

What to do next

For information on migrating your existing Mutli-Site Orchestrator configuration deployed in VMware ESX to Cisco Application Services Engine cluster, see Migrating Existing Cluster to Application Service Engine, on page 14.

For more information about Day-0 Operations, see the *Day-0 Operations* section of the document.

Migrating Existing Cluster to Application Service Engine

This section provides an overview of how to migrate your existing Multi-Site deployment to a new cluster deployed in Cisco Application Service Engine.

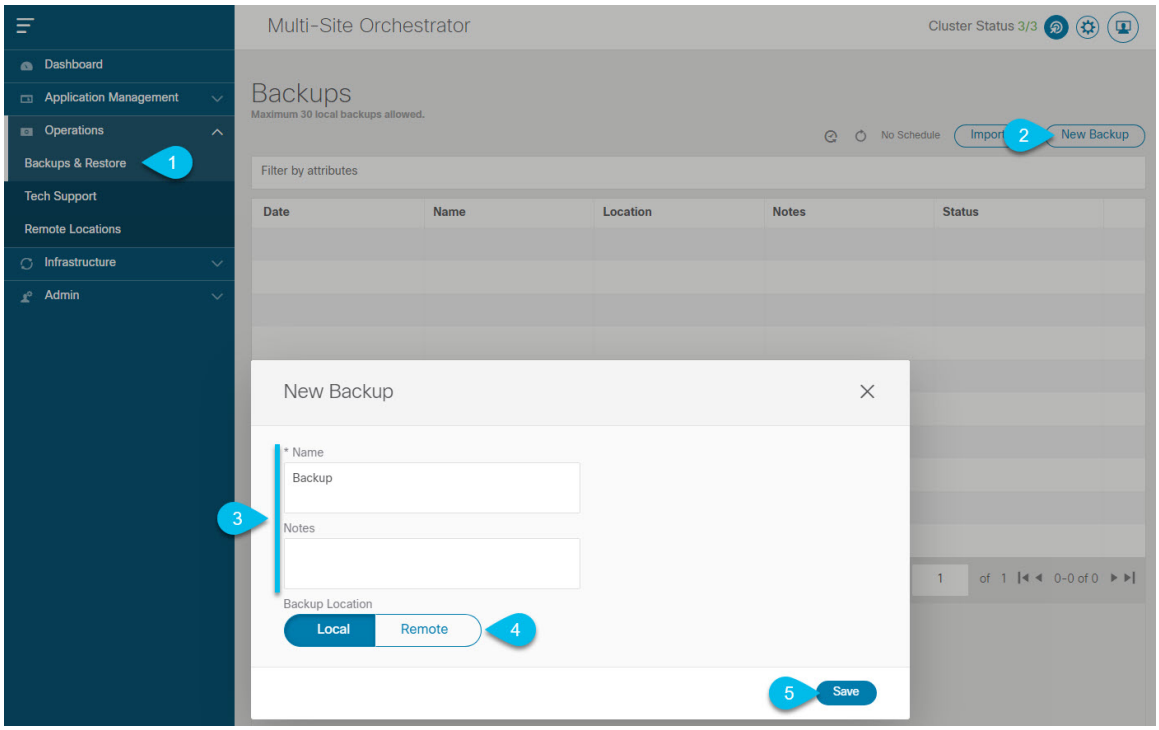
Because the two platforms are vastly different in how they implement clustering and infrastructure, the migration process involves parallel deployment of the new platform and manual transfer of the current configuration database from the existing Orchestrator cluster.

Procedure

Step 1 Deploy a brand new Orchestrator cluster in Application Service Engine.

The procedure is described in the Deploying Multi-Site Orchestrator in Application Services Engine, on page 11 section.

Step 2 Backup existing deployment configuration.



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Operations** > **Backups & Restore**.
c) In the main window, click **New Backup**.

A **New Backup** window opens.

d) 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 (_).

e) 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.

f) Click **Save** to create the backup.

Step 3 Copy the Backup file from the existing Orchestrator.

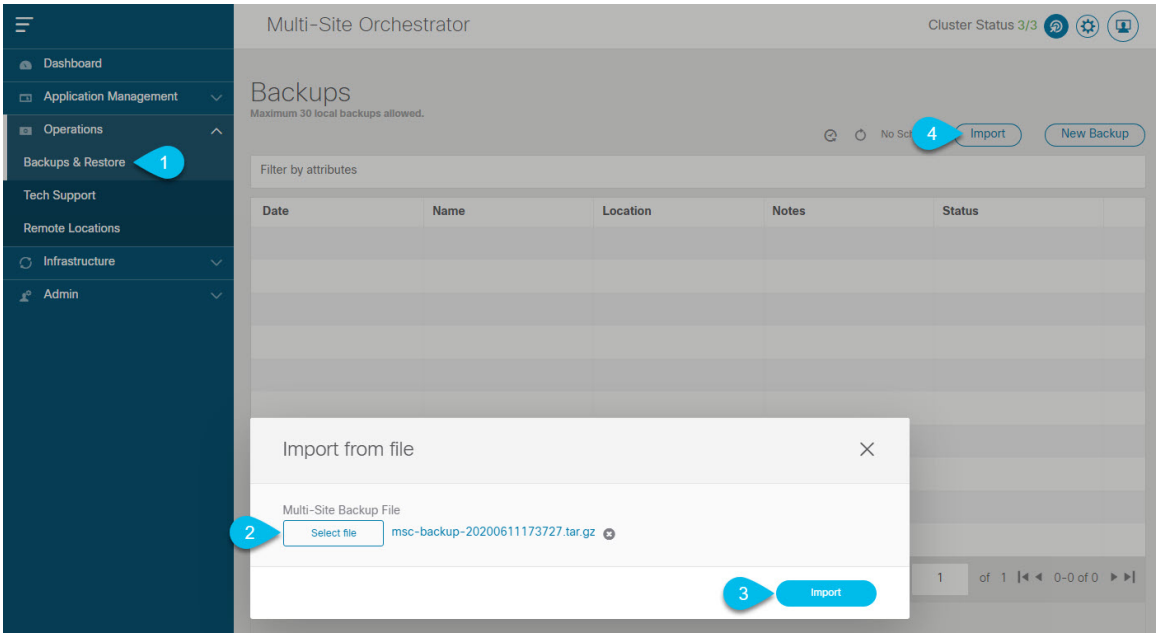
If you created the backup using a remote location, you can skip this step.

Otherwise, in the main window, click the actions ( ) icon next to the backup and select **Download**. This will download the backup file to your system.

Step 4 Bring down the existing Multi-Site Orchestrator cluster VMs.

Step 5 Import the backup file to your new Orchestrator cluster deployed on the Application Service Engine.

If you saved the backup locally, simply import the file:



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.

b) From the left navigation pane, select **Operations** > **Backups & Restore**.

c) In the main window, click **Import**.
d) 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.

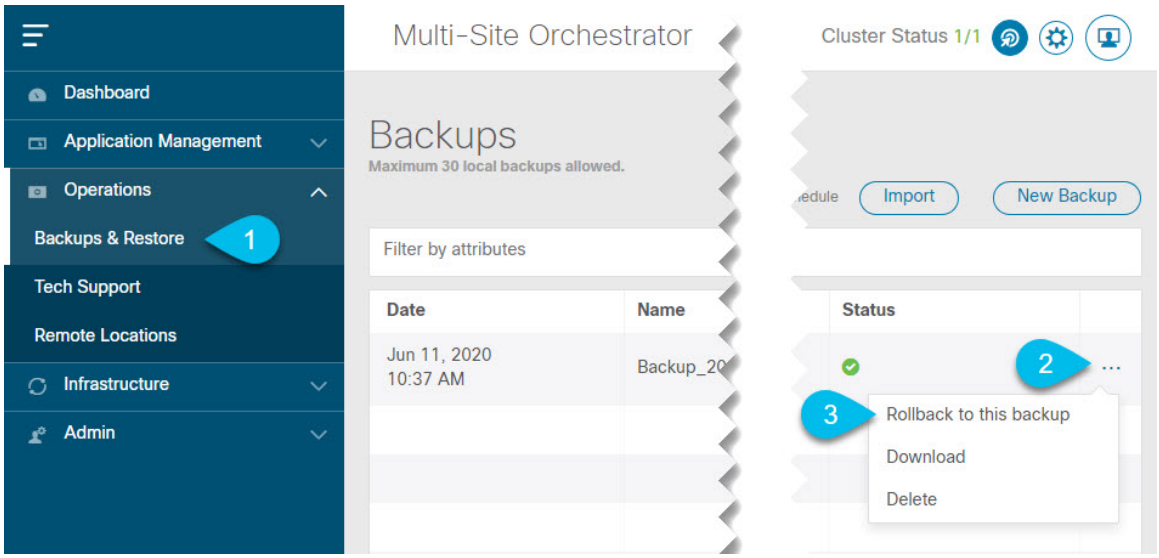
If you saved the backup to a remote location, add the remote location to the new Multi-Site Orchestrator:

a) Log in to your Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Admin** > **Remote Locations**.
c) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

d) Provide the same information for the remote location that you used in your old Orchestrator.
e) Click **Save** to add the remote server.

Step 6 Restore the configuration.

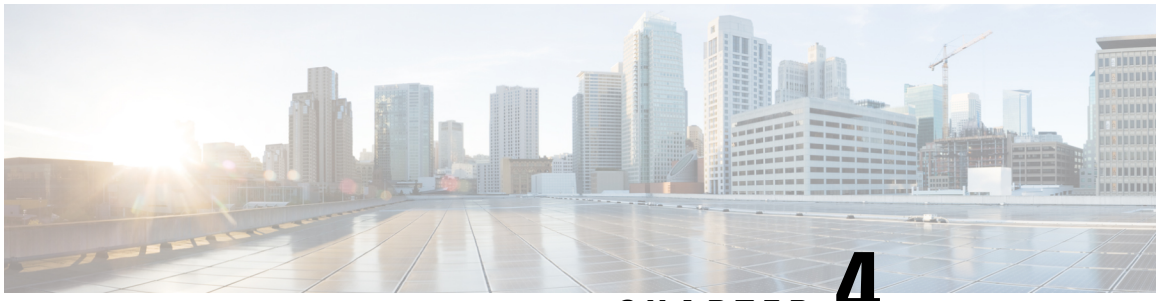


a) From the left navigation menu, select **Admin** > **Backups**.
b) In the main window, click the actions ( ) icon next to the backup you want to restore and select **Rollback to this backup**.

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.

c) 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.



CHAPTER 4

Deploying in VMware ESX

This chapter contains the following sections:

- Prerequisites and Guidelines, on page 17
- Deploying Cisco ACI Multi-Site Orchestrator Using Python, on page 19
- Deploying Orchestrator in vCenter, on page 25
- Deploying Orchestrator in ESX Directly, on page 29

Prerequisites and Guidelines

For all new deployments, we recommend using Cisco Application Services Engine as described in Deploying in Cisco Application Services Engine, on page 9 instead. However, if you still want to deploy the Orchestrator cluster in VMware ESX VMs directly, you can follow the guidelines and procedures in this chapter.

This chapter covers deployment of a 3-node Multi-Site Orchestrator cluster. If you want to set up a single-node Orchestrator (for example, for testing purposes), follow the instruction in the Installing Single Node Orchestrator, on page 111 chapter instead.

Deployment Method

When deploying in ESX VMs, you can choose one of the following 2 approaches:

- Use Cisco-provided Python scripts to deploy the entire Multi-Site Orchestrator cluster. The scripts allow you to execute the deployment and later upgrades remotely, for example from your laptop, as long as you have access to the vCenter where the Orchestrator VMs are to be deployed.

  This is the preferred approach when deploying an Orchestrator cluster in ESX VMs as it automates a number of manual steps and allows remote execution of Cisco ACI Multi-Site Orchestrator installation and subsequent software upgrades.

- Using an OVA image to deploy each Orchestrator VM individually. In this case you can also choose to deploy the image either using the vCenter or directly on the ESX server.

Docker Subnet Considerations

The Multi-Site Orchestrator application services run in Docker containers. When deployed, Docker uses a number of internal networks for its own application services (`bridge`, `docker_gwbridge`) as well as the Orchestrator services (`msc_msc`).

You can configure custom networks for the Docker services during Orchestrator deployment. Two additional parameters are available in the Python configuration file or the OVA template:



Note When configuring these networks, ensure that they are unique and do not overlap with any existing networks in the environment.

- **Application overlay**: The default address pool to be used for Docker internal bridge networks.

  Application overlay must be a `/16` network. Docker then splits this network into two `/24` subnets used for the internal `bridge` and `docker_gwbridge` networks.

  For example, if you set the application overlay pool to `192.168.0.0/16`, Docker will use `192.168.0.0/24` for the `bridge` network and `192.168.1.0/24` for the `docker_gwbridge` network.

- **Service overlay**: The default Docker overlay network IP.

  Service overlay must be a `/24` network and is used for the `msc_msc` Orchestrator Docker service network.

Network Time Protocol (NTP)

Multi-Site Orchestrator uses NTP for clock synchronization, so you must have an NTP server configured in your environment. You provide NTP server information as part of the Orchestrator installation procedure.



Note VMware Tools provides an option to synchronize VMs' time with the host, however you should use only one type of periodic time synchronization in your VMs. Because you will enable NTP during Multi-Site Orchestrator deployment, ensure that VMware Tools periodic time synchronization is disabled for the Orchestrator VMs.

VMware vSphere Requirements

The following table summarizes the VMware vSphere requirements for Multi-Site Orchestrator:

- You must not enable vMotion for Multi-Site Orchestrator VMs.

  vMotion is not supported with docker swarm, which is used by the Multi-Site Orchestrator.

- You must ensure that the following vCPUs, memory, and disk space requirements are reserved for each VM and are not part of a shared resource pool:

Table 2: VMware vSphere Requirements

Orchestrator Version	Requirements
Release 3.0(1) or later	• ESXi 6.0 or later • 6 vCPUs (8 vCPUs recommended) • 48 GB of RAM • 64 GB disk • 10 GHz CPU reservation CPU cycle reservation is automatically applied when first deploying the Orchestrator VMs.

Deploying Cisco ACI Multi-Site Orchestrator Using Python

The following sections describe how to prepare for and deploy Cisco ACI Multi-Site Orchestrator using Python.

Setting Up Python Environment

This section describes how to set up the Python environment for deploying Cisco ACI Multi-Site Orchestrator using Python. You must set up the Python environment on the laptop or server from which you will run the installation scripts.



Note If you have already set up your python environment, for example for another Multi-Site deployment or upgrade, you can skip this section.

Before you begin

you will need:

- A laptop or server from which you will run the scripts.

  You must not use any of the Multi-Site Orchestrator nodes for this purpose.
- Python 3.4 or later already installed on the system from which you will run the scripts.

Procedure

Step 1 Download the **ACI Multi-Site Tools** image from Cisco ACI Multi-Site Software Download link.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site Tools Image* file (`tools-msc-<version>.tar.gz`).

Step 2 Extract the files.

```
# tar -xvzf tools-msc-<version>.tar.gz
```

Step 3 Change to the extracted directory.

```
# cd tools-msc-<version>
```

Step 4 Verify that you are running a correct version of Python.

Depending on your operating system and the versions of Python you have installed previously, the Python 3.x executable may be associated with `python` or `python3` command.

You can use one of the following examples to confirm which executable corresponds to Python 3.4 or later.

- If you have installed Python 2.x first, the `python` command will likely be associated with that version:

  ```
  # python -V
  Python 2.7.18
  ```

  In this case, you may need to use `python3` command instead:

  ```
  # python3 -V
  Python 3.4.5
  ```

- If you have installed only Python 3.x, you can use the `python` command as well:

  ```
  # python -V
  Python 3.7.7
  ```

Note
The following steps use `python` for all commands. If your Python 3.x executable is associated with `python3`, use that command instead.

Step 5 If you plan to use a proxy to access the Internet, make sure to configure the proxy as follows:

```
# export http_proxy=<proxy-ip-address>:<proxy-port>
# export https_proxy=<proxy-ip-address>:<proxy-port>
```

Step 6 Install or update the Python package manager.

```
# python -m ensurepip
```

If the package is already installed, update it to the latest version:

```
# python -m ensurepip --upgrade
```

Step 7 (Optional) Set up Python virtual environment.

We recommend using `virutalenv` to install the packages, so they do not impact the existing packages in the system. The following steps provide a brief overview of how to set up `virutalenv`. For additional information on how to use `virtualenv`, see Installing packages using pip and virtualenv.

a) Install `virtualenv`.

```
# python -m pip install --user virtualenv
```

b) Change into the directory where you want the virtual environment files to be created.
c) Create a virtual environment.

In the following command, provide a name for the virtual environment, for example *mso-deployments*.

```
# python -m venv <env-name>
```

d) Activate the virtual environment.

You need to activate the virtual environment you created before installing the packages required for Orchestrator deployment or upgrade in the next step.

For Windows:

```
# .\<env-name>\Scripts\activate.bat
```

For Linux:

```
# source ./<env-name>/bin/activate
```

Step 8 Install the required packages.

The required packages are listed in the `requirements.txt` file.

```
# python -m pip install -r requirements.txt
```

Note
The Python installation must complete successfully. If you encounter any errors, you must address them before proceeding to the next section or the Cisco ACI Multi-Site Orchestrator Python scripts will not work.

Step 9 If you used virtual Python environment, deactivate it now.

```
# deactivate
```

Sample Deployment Configuration File

When you deploy Multi-Site Orchestrator using Python, several required configuration details are specified in a YAML configuration file. This section provides a sample `msc_cfg.yml` file.

In the following sample configuration file all the VMs are created under the same host. The "host" parameter in the configuration file can be given as a node-level parameter instead if you want to create the Multi-Site VMs in different hosts.

```
# vCenter parameters
vcenter:
  name: 192.168.142.59
  user: administrator@vsphere.local

  # Host under which the Orchestrator VMs will be created
  host: 192.64.142.55

  # Path to the Orchestrator OVA file
  msc_ova_file: ../images/msc-2.1.1h.ova

  # (Optional) If not provided, default library name 'msc-content-lib' will be used
  #library: content-library-name

  # Library datastore name
  library_datastore: datastore1

  # Host datastore name
  host_datastore: datastore1

  # Prefix for Orchestrator VM names, full VM names will be '<vm_name_prefix>-node1',
  # '<vm_name_prefix>-node2', and '<vm_name_prefix>-node3'
```

```
  vm_name_prefix: msc

  # Wait Time in seconds for VMs to come up
  vm_wait_time: 120


# Common parameters for all nodes
common:
  # Network mask
  netmask: 255.255.248.0

  # Gateway' IP address
  gateway: 192.64.136.1

  # Domain Name-Server IP. Leave blank for DHCP
  nameserver: 192.64.136.140

  # Network label of the Management network port-group
  management: "VM Network"

  # Time zone of the node, must be one of the values listed by 'timedatectl list-timezones'
 command
  time_zone: America/Los_Angeles

  # NTP (Network Time Protocol) servers, multiple servers can be listed separated by commas

  ntp_servers: ntp.company.com

  # Application Overlay IP for docker bridge type networks
  # Docker's bridge and docker_gwbridge networks are assigned addresses from this pool
  application_overlay: 192.168.0.0/16

  # Service Overlay IP for docker overlay type networks
  # Docker's msc_msc overlay network created at the time of deployment are assigned this
network address
  service_overlay: 2.1.1.0/24

# Node specific parameters over-ride the vCenter and common parameters
node1:
  # To use static IP, specify a valid IP address for the "ip" attribute
  # To obtain IP via DHCP, leave the "ip" field blank
  ip: 192.64.136.204

  # Node specific "netmask" parameter over-rides the common.netmask
  netmask: 255.255.248.0

  # (Optional) If hostname is not specified, the VM name will be used
  hostname: mso-node1

node2:
  # To use static IP, specify a valid IP address for the "ip" attribute
  # To obtain IP via DHCP, leave the "ip" field blank
  ip:

  # (Optional) If hostname is not specified, the VM name will be used
  hostname: mso-node2

node3:
  # To use static IP, specify a valid IP address for the "ip" attribute
  # To obtain IP via DHCP, leave the "ip" field blank
  ip:

  # (Optional) If hostname is not specified, the VM name will be used
  hostname: mso-node3
```

Deploying Multi-Site Orchestrator Using Python

This section describes how to deploy Cisco ACI Multi-Site Orchestrator using Python.

Before you begin

- Ensure that you meet the hardware requirements and compatibility that is listed in the *Cisco ACI Multi-Site Hardware Requirements Guide*.
- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 17.
- Ensure that the NTP server is configured and reachable from the Orchestrator VMs and that VMware Tools periodic time synchronization is disabled.
- Ensure that the vCenter is reachable from the laptop or server where you will extract the tools and run the installation scripts.
- Ensure that your Python environment is set up as described in Setting Up Python Environment, on page 19.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator image and tools.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.

d) Download the *ACI Multi-Site Image* file (`msc-<version>.tar.gz`) for the release.

e) Download the *ACI Multi-Site Tools Image* file (`tools-msc-<version>.tar.gz`) for the release.

Step 2 Extract the `tools-msc-<version>.tar.gz` file to the directory from which you want to run the install scripts.

```
# tar -xvzf tools-msc-<version>.tar.gz
```

Then change into the extracted directory:

```
# cd tools-msc-<version>
```

Step 3 Create a `msc_cfg.yml` configuration file for your install.

You can copy and rename the provided `msc_cfg_example.yml` file or you can create the file using the example provided in Sample Deployment Configuration File, on page 21.

Step 4 Edit the `msc_cfg.yml` configuration file and fill in all the parameters for your environment.

The parameters that must be filled in are in all caps, for example <*VCENTER_NAME*>. You will also need to update <*MSC_TGZ_FILE_PATH*> with the path to the `msc-<version>.tar.gz` image file you downloaded in Step 1.

For a complete list of available parameters, see the sample `msc_cfg.yml` file is provided in Sample Deployment Configuration File, on page 21.

Step 5 Execute the script to deploy the Orchestrator VMs and prepare them:

```
# python msc_vm_util.py -c msc_cfg.yml
```

Step 6 Enter vCenter, `node1`, `node2` and `node3` passwords when prompted.

The script creates three Multi-Site Orchestrator VMs and executes the initial deployment scripts. This process may take several minutes to complete. After successful execution, the Multi-Site Orchestrator cluster is ready for use.

It may take several minutes for the deployment to complete.

Step 7 Verify that the cluster was deployed successfully.

a) Log in to any one of the deployed Orchestrator nodes.

b) Verify that all nodes are up and running.

```
# docker node ls
ID                            HOSTNAME        STATUS      AVAILABILITY    [...]
y90ynithc3cejkeazcqlu1uqs *   node1           Ready       Active          [...]
jt67ag14ug2jgaw4r779882xp     node2           Ready       Active          [...]
hoae55eoute6l5zpqlnxsk8o8     node3           Ready       Active          [...]
```

Confirm the following:

- The `STATUS` field is `Ready` for all nodes.
- The `AVAILABILITY` field is `Active` for all node.
- The `MANAGER STATUS` field is `Leader` for one of the nodes and `Reachable` for the other two.

c) Verify that all replicas are fully up.

```
# docker service ls
ID                  NAME                      MODE            REPLICAS    [...]
p6tw9mflj06u        msc_auditservice          replicated      1/1         [...]
je7s2f7xme6v        msc_authyldapservice      replicated      1/1         [...]
dbd27y76eouq        msc_authytacacsservice    replicated      1/1         [...]
untetoygqn1q        msc_backupservice         global          3/3         [...]
n5eibyw67mbe        msc_cloudsecservice       replicated      1/1         [...]
8inekkof982x        msc_consistencyservice    replicated      1/1         [...]
0qeisrguy7co        msc_endpointservice       replicated      1/1         [...]
e8ji15eni1e0        msc_executionengine       replicated      1/1         [...]
s4gnm2vge0k6        msc_jobschedulerservice   replicated      1/1         [...]
av3bjvb9ukru        msc_kong                  global          3/3         [...]
rqie68m6vf9o        msc_kongdb                replicated      1/1         [...]
51u1g7t6ic33        msc_mongodb1              replicated      1/1         [...]
vrl8xvvx6ky5        msc_mongodb2              replicated      1/1         [...]
0kwk9xw8gu8m        msc_mongodb3              replicated      1/1         [...]
qhejgjn6ctwy        msc_platformservice       global          3/3         [...]
l7co71lneegn        msc_schemaservice         global          3/3         [...]
1t37ew5m7dxi        msc_siteservice           global          3/3         [...]
tu37sw68a1gz        msc_syncengine            global          3/3         [...]
8dr0d7pq6j19        msc_ui                    global          3/3         [...]
swnrzrbcv60h        msc_userservice           global          3/3         [...]
```

d) Log in to the Cisco ACI Multi-Site Orchestrator GUI.

You can access the GUI using any of the 3 nodes' IP addresses.

The default log in is **admin** and the default password is **We1come2msc!**.

When you first log in, you will be prompted to change the password.

What to do next

For more information about Day-0 Operations, see Adding Tenants and Schemas, on page 77.

Deploying Orchestrator in vCenter

This section describes how to deploy Cisco ACI Multi-Site Orchestrator using an OVA in vCenter.

Before you begin

- Ensure that you meet the hardware requirements and compatibility that is listed in the *Cisco ACI Multi-Site Hardware Requirements Guide*.
- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 17.
- Ensure that the NTP server is configured and reachable from the Orchestrator VMs and that VMware Tools periodic time synchronization is disabled.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator Image.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) From the left sidebar, choose the Cisco ACI Multi-Site Orchestrator release version.

d) Download the *ACI Multi-Site Image* file (`msc-<version>.ova`) for the release.

Step 2 Deploy the OVA using the VMware vCenter.

Note
The OVA cannot be deployed directly in ESX, it must be deployed using vCenter. If you want to deploy Cisco ACI Multi-Site Orchestrator directly in ESX, see the "*Deploying Multi-Site Orchestrator in ESX Directly* " sections in this chapter.for instructions on how to extract the OVA and install the Orchestrator without vCenter.

Step 3 Configure the OVA properties.

In the **Properties** dialog box, enter the appropriate information for each VM:

- In the **Enter password** field, enter the root password for the VM.
- In the **Confirm password** field, enter the password again.
- In the **Hostname** field, enter the hostnames for each Cisco ACI Multi-Site Orchestrator node. You can use any valid Linux hostname.
- In the **Management Address** (network address) field, enter the network address or leave the field blank to obtain it via DHCP.

Note

The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- In the **Management Netmask** (network netmask) field, enter the netmask netmask or leave the field blank to obtain it via DHCP.
- In the **Management Gateway** (network gateway) field, enter the network gateway or leave the field blank to obtain it via DHCP.
- In the **Domain Name System Server** (DNS server) field, enter the DNS server or leave the field blank to obtain it via DHCP.
- In the **Time-zone string (Time-zone)** field, enter a valid time zone string.

  You can find the time zone string for your region in the IANA time zone database or using the `timedatectl list-timezones` Linux command. For example, *`America/Los_Angeles`* .
- In the **NTP-servers** field, enter Network Time Protocol servers separated by commas.
- In the **Application overlay** field, enter the default address pool to be used for Docker internal bridge networks.

  Application overlay must be a `/16` network. Docker then splits this network into two `/24` subnets used for the internal `bridge` and `docker_gwbridge` networks.

  For example, if you set the application overlay pool to `192.168.0.0/16`, Docker will use `192.168.0.0/24` for the `bridge` network and `192.168.1.0/24` for the `docker_gwbridge` network.

  You must ensure that the application overlay network is unique and does not overlap with any existing networks in the environment.

  Note
  The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- In the **Service overlay** field, enter the default Docker overlay network IP.

  Service overlay must be a `/24` network and is used for the `msc_msc` Orchestrator Docker service network.

  You must ensure that the service overlay network is unique and does not overlap with any existing networks in the environment.

  Note
  The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- Click **Next**.
- In the **Deployment settings** pane, check all the information you provided is correct.
- Click **Power on after deployment**.
- Click **Finish**.

In addition to the above parameters, a 10GHz CPU cycle reservation is automatically applied to each Orchestrator VM when deploying the OVA.

Step 4 Repeat the previous two steps to deploy two more VMs.

The three VMs you deploy will join to form the Orchestrator cluster.

Step 5 Ensure that the virtual machines are able to ping each other.

Step 6 Initialize node1.

a) Connect to node1 using SSH.

b) Change to the initialization scripts directory.

```
# cd /opt/cisco/msc/builds/<build_number>/prodha
```

c) Run initialization script and note the generated secret.

```
# ./msc_cfg_init.py
Starting the initialization of the cluster...
.
.
.
Both secrets created successfully.

Join other nodes to the cluster by executing the following on each of the
other nodes:
./msc_cfg_join.py \
SWMTKN-1-4pu9zc9d81gxxw6mxec5tuxdt8nbarq1qnmfw9zcme1w1tljZh-7w3iwsddvd97ieza3ym1s5gj5
 \
<node1-ip-address>
```

You will use the above token and IP address in the following steps to join `node2` and `node3` into the cluster.

d) Note the management IP address of the first node.

```
# ifconfig
inet 10.23.230.151  netmask 255.255.255.0  broadcast 192.168.99.255
```

You will use this IP address in the following steps to join node2 and node3 into the cluster.

Step 7 Join node2 to the cluster.

a) Connect to node2 using SSH.

b) Change to the `/opt/cisco/msc/builds/<build_number>/prodha` directory.

```
# cd /opt/cisco/msc/builds/<build_number>/prodha
```

c) Execute the `msc_cfg_join.py` command using the IP address of the first node.

```
# ./msc_cfg_join.py \
SWMTKN-1-4pu9zc9d81gxxw6mxec5tuxdt8nbarq1qnmfw9zcme1w1tljZh-7w3iwsddvd97ieza3ym1s5gj5 \
10.23.230.151
```

Step 8 Join node3 to the cluster.

a) Connect to node3 using SSH.

b) Change to the `/opt/cisco/msc/builds/<build_number>/prodha` directory.

```
# cd /opt/cisco/msc/builds/<build_number>/prodha
```

c) Execute the `msc_cfg_join.py` command using the IP address of the first node.

```
# ./msc_cfg_join.py \
SWMTKN-1-4pu9zc9d81gxxw6mxec5tuxdt8nbarq1qnmfw9zcme1w1tljZh-7w3iwsddvd97ieza3ym1s5gj5 \
10.23.230.151
```

Step 9 On any node, make sure the nodes are healthy.

```
# docker node ls
ID                            HOSTNAME        STATUS      AVAILABILITY    [...]
y90ynithc3cejkeazcqluluqs *   node1           Ready       Active          [...]
jt67ag14ug2jgaw4r779882xp     node2           Ready       Active          [...]
hoae55eoute6l5zpqlnxsk8o8     node3           Ready       Active          [...]
```

Confirm the following:

- The `STATUS` field is `Ready` for all nodes.
- The `AVAILABILITY` field is `Active` for all node.
- The `MANAGER STATUS` field is `Leader` for one of the nodes and `Reachable` for the other two.

Step 10 On any node, execute the `msc_deploy.py` command:

```
# cd /opt/cisco/msc/builds/<build_number>/prodha
# ./msc_deploy.py
```

Step 11 On any node, make sure that all REPLICAS are up.

```
# docker service ls
ID                  NAME                      MODE            REPLICAS    [...]
p6tw9mflj06u        msc_auditservice          replicated      1/1         [...]
je7s2f7xme6v        msc_authyldapservice      replicated      1/1         [...]
dbd27y76eouq        msc_authytacacsservice    replicated      1/1         [...]
untetoygqn1q        msc_backupservice         global          3/3         [...]
n5eibyw67mbe        msc_cloudsecservice       replicated      1/1         [...]
8inekkof982x        msc_consistencyservice    replicated      1/1         [...]
0qeisrguy7co        msc_endpointservice       replicated      1/1         [...]
e8ji15eni1e0        msc_executionengine       replicated      1/1         [...]
s4gnm2vge0k6        msc_jobschedulerservice   replicated      1/1         [...]
av3bjvb9ukru        msc_kong                  global          3/3         [...]
rqie68m6vf9o        msc_kongdb                replicated      1/1         [...]
51u1g7t6ic33        msc_mongodb1              replicated      1/1         [...]
vrl8xvvx6ky5        msc_mongodb2              replicated      1/1         [...]
0kwk9xw8gu8m        msc_mongodb3              replicated      1/1         [...]
qhejgjn6ctwy        msc_platformservice       global          3/3         [...]
l7co71lneegn        msc_schemaservice         global          3/3         [...]
1t37ew5m7dxi        msc_siteservice           global          3/3         [...]
tu37sw68a1gz        msc_syncengine            global          3/3         [...]
8dr0d7pq6j19        msc_ui                    global          3/3         [...]
swnrzrbcv60h        msc_userservice           global          3/3         [...]
```

Step 12 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

You can access the GUI using any of the 3 nodes' IP addresses.

The default log in is **admin** and the default password is **We1come2msc!**.

When you first log in, you will be prompted to change the password.

What to do next

For more information about Day-0 Operations, see Adding Tenants and Schemas, on page 77.

Deploying Orchestrator in ESX Directly

This section describes how to deploy Cisco ACI Multi-Site Orchestrator directly in ESX without using vCenter.

Before you begin

- Ensure that you meet the hardware requirements and compatibility that is listed in the *Cisco ACI Multi-Site Hardware Requirements Guide*.
- Ensure that you meet the requirements and guidelines described in Prerequisites and Guidelines, on page 17.
- Ensure that the NTP server is configured and reachable from the Orchestrator VMs and that VMware Tools periodic time synchronization is disabled.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator Image.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.

d) Download the *ACI Multi-Site Image (ESX Only)* file (`esx-msc-<version>.ova`) for the release.

Step 2 Untar the OVA file into a temporary directory:

```
# mkdir msc_ova
# cd msc_ova
# tar xvf ../esx-msc-<version>.ova
esx-msc-<version>.cert
esx-msc-<version>.mf
esx-msc-<version>.ovf
esx-msc-<version>-disk1.vmdk
```

Step 3 Use the ESX vSphere client to deploy the OVF.

a) Log in to vSphere.

b) Navigate to **File** > **Deploy OVF Template** > **Browse** and choose the `esx-msc-<version>.ovf` file.

c) Complete rest of the menu options and deploy the VM.

d) Repeat the steps for 2 additional Orchestrator nodes.

Step 4 Configure networking on Node2 and Node3.

After you configure networking on all 3 nodes, you will designate one of the nodes as `Primary` for the Docker swarm and use it to joint all 3 nodes into a cluster. Before you can do that, you must first configure networking on the two secondary nodes.

a) Log in to one of the secondary nodes (for example, `Node2`) as the `root` user.

The default password is `cisco`.

b) Change the default password.

The first time you log in, you will be prompted to change the default `root` password.

c) Run the Orchestrator setup utility.

```
# mso-setup
```

d) When prompted if it is a primary node, enter **n**.

```
  When first deploying, one node must be designated as primary.
  You must configure the other two nodes before configuring the primary node.
        If this is NOT the primary node, simply choose 'no' to proceed.
        If this is the primary node and the other nodes are ready, answer 'yes' to deploy.

Is this the primary node [y/N]? n
```

e) Confirm whether or not you will use a DHCP server to assign IP addresses to the node.

If you choose to use a DHCP server, you will not be prompted for specific IP configuration, otherwise you will enter it in the next step.

```
Is this system going to get it's network configuration from a DHCP server [y/N]? n
```

f) Provide the required information.

The setup utility will prompt for the following information:

- **Management address**, for example `10.195.223.200`

  If you chose to use a DHCP server, this field is skipped.

- **Management netmask**, for example `255.255.255.0`

  If you chose to use a DHCP server, this field is skipped.

- **Management gateway**, for example `10.195.223.1`

  If you chose to use a DHCP server, this field is skipped.

- **DNS server**, for example `171.70.168.183`

  If you chose to use a DHCP server, this field is skipped.

- **Hostname**, for example `mso-node2`

  You can use any valid Linux hostname.

- **Time zone string**, for example `America/Los_Angeles`

  You can find the time zone string for your region in the IANA time zone database or using the `timedatectl list-timezones` Linux command.

- **NTP servers**, for example `ntp.esl.cisco.com`

  You can provide multiple NTP servers separated by commas.

- **Application overlay network**, for example `192.168.0.0/16`

  Application overlay must be a `/16` network. Docker then splits this network into two `/24` subnets used for the internal `bridge` and `docker_gwbridge` networks.

  For example, if you set the application overlay pool to `192.168.0.0/16`, Docker will use `192.168.0.0/24` for the `bridge` network and `192.168.1.0/24` for the `docker_gwbridge` network.

  You must ensure that the application overlay network is unique and does not overlap with any existing networks in the environment.

- **Service overlay network**, for example `1.1.1.0/24`

  Service overlay must be a `/24` network and is used for the `msc_msc` Orchestrator Docker service network.

  You must ensure that the service overlay network is unique and does not overlap with any existing networks in the environment.

g) Verify the provided information.

After you finish entering the information, you will be prompted to verify it. Reply **`y`** to confirm or **`n`** to re-enter the information.

```
== Verify network configuration ==

      Management address: 10.195.223.200
      Management netmask: 255.255.255.0
      Management gateway: 10.195.223.1
      DNS server: 171.70.168.183
      Hostname: msc-node2
      Time zone string: America/Los_Angeles
      NTP servers: ntp.esl.cisco.com
      Application overlay network: 192.168.0.0/16
      Service overlay network: 1.1.1.0/24

Confirm the settings and proceed [Y/n]? y
```

Step 5 Repeat the previous step for the other secondary node (Node3).

Step 6 Configure the primary node (Node1) and deploy the cluster.

a) Log in to the primary node (Node1) as the `root` user.

The default password is `cisco`.

b) Change the default password.

The first time you log in, you will be prompted to change the default `root` password.

c) Run the Orchestrator setup utility.

```
# mso-setup
```

d) When prompted if it is a primary node, enter **`y`**.

```
  When first deploying, one node must be designated as primary.
  You must configure the other two nodes before configuring the primary node.
        If this is NOT the primary node, simply choose 'no' to proceed.
        If this is the primary node and the other nodes are ready, answer 'yes' to deploy.

Is this the primary node [y/N]? y
```

e) Confirm that the other two nodes have been configured.

If you have not configured the other two nodes, you can respond **`n`** and re-run the setup utility at a later time.

```
Are other two nodes network configured [y/N]? y
```

f) Provide the network configuration information like you did for the other two nodes.

g) After you verify and confirm the network settings, provide other two nodes' information.

You will be prompted to enter the IP addresses and `root` passwords for the other 2 nodes.

```
Confirm the settings and proceed [Y/n]? y
== MSO Network configuration done for node1 ==
== MSO Setup begins ==
```

```
Node2 IP address: 10.195.223.200
Node2 root password:
Node3 IP address: 10.195.223.201
Node3 root password:
msc_setup: Start
```

If for any reason the setup does not complete, you can re-run just the deployment part without the full network configuration using the following command:

```
# mso-setup --install-mso
```

Step 7 Wait for the cluster to be deployed.

After you confirm the settings on the primary node, the setup utility

Step 8 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

You can access the GUI using any of the 3 nodes' IP addresses.

The default log in is **admin** and the default password is **We1come2msc!**.

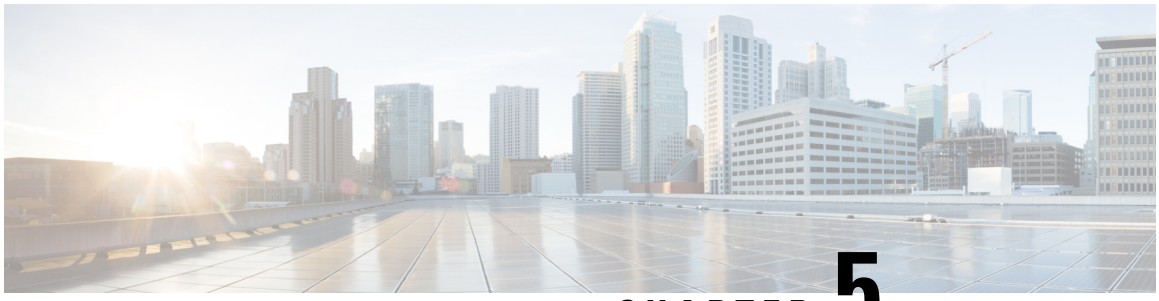
When you first log in, you will be prompted to change the password.



PART II

Day-0 Operations

- Configuring and Adding Sites, on page 35
- Configuring Infra, on page 43
- CloudSec Encryption, on page 63
- Adding Tenants and Schemas, on page 77



CHAPTER 5

Configuring and Adding Sites

This chapter contains the following sections:

- Pod Profile and Policy Group, on page 35
- Configuring Fabric Access Policies for All APIC Sites, on page 35
- Configuring Sites That Contain Remote Leaf Switches, on page 38
- Cisco Mini ACI Fabrics, on page 40
- Adding Sites, on page 41

Pod Profile and Policy Group

In each site's APIC, you must have one Pod profile with a Pod policy group. If you site does not have a Pod policy group you must create one.

To check if the POD profile contains a POD policy group:

- Navigate to the Cisco APIC GUI, **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**.

To create a POD policy group:

- Navigate to the Cisco APIC GUI, **Fabric** > **Fabric Policies** > **Pods** > **Policy Groups**, right-click **Policy Groups** and click **Create Pod Policy Group**. Enter the appropriate information and click **Submit**.

To assign the new pod policy group to the default POD profile:

- Navigate to the Cisco APIC GUI, **Fabric** > **Fabric Policies** > **Pods** > **Profiles** > **Pod Profile default**. Click on the default, choose the new pod policy group and click **Update**.

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.

Procedure

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 Configuring Fabric Access Interface Policies, on page 37.

Configuring Fabric Access Interface Policies

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

Before you begin

You must have configured 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 35.

Procedure

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`.

c) 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.

c) 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`.

c) Click **Update** to save the selector.
d) Click **Next** to proceed to the next screen.
e) Select the interface profile you have created in the previous step

For example `Spine1-ISN`.

f) Click **Finish** to save the spine profile.

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.

Procedure

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

Procedure

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.

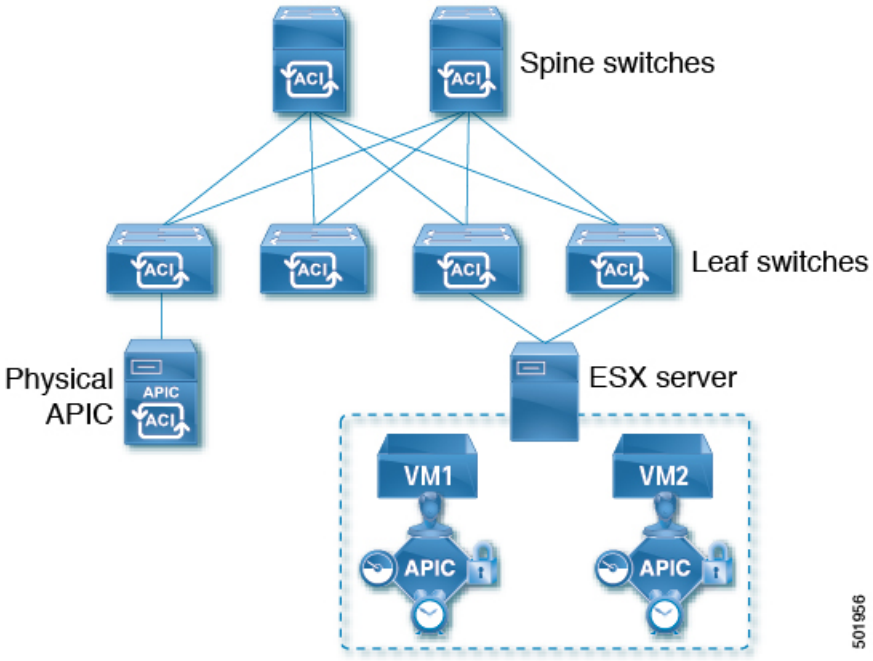
Cisco Mini ACI Fabrics

Cisco ACI Multi-Site supports Cisco Mini ACI fabrics as typical on-premises sites without requiring any additional configuration. This section provides brief overview of Mini ACI fabrics, detailed info on deploying and configuring this type of fabrics is available in *Cisco Mini ACI Fabric and Virtual APICs*.

Cisco ACI, Release 4.0(1) introduced Mini ACI Fabric for small scale deployment. Mini ACI fabric works with Cisco APIC cluster consisting of one physical APIC and two virtual APICs (vAPIC) running in virtual machines. This reduces the physical footprint and cost of the APIC cluster, allowing ACI fabric to be deployed in scenarios with limited rack space or initial budget, such as a colocation facility or a single-room data center, where a full-scale ACI installations may not be practical due to physical footprint or initial cost.

The following diagram shows an example of a mini Cisco ACI fabric with a physical APIC and two virtual APICs (vAPICs):

Figure 1: Cisco Mini ACI Fabric



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.

Procedure

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.

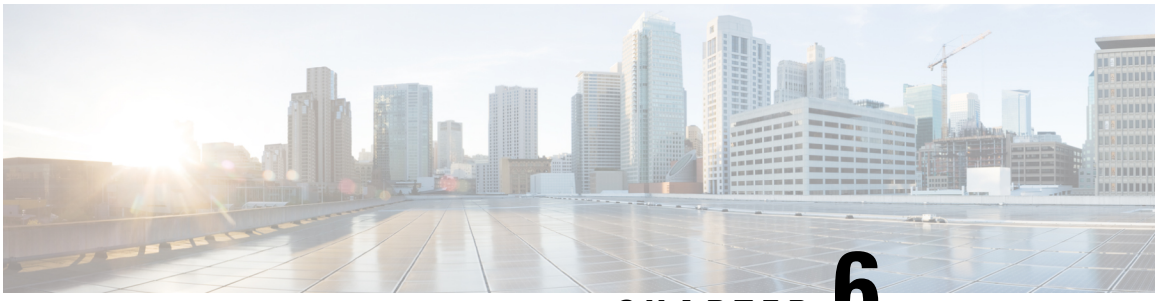
Step 5 Click **Save** to add the site.

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.



CHAPTER 6

Configuring Infra

This chapter contains the following sections:

- Configuring Infra Prerequisites and Guidelines, on page 43
- Configuring Infra: General Settings, on page 44
- Refreshing Site Connectivity Information, on page 44
- Configuring Infra: On-Premises Site Settings, on page 45
- Configuring Infra: Cloud Site Settings, on page 47
- Configuring Infra: Pod Settings, on page 47
- Configuring Infra: Spine Switches, on page 48
- Configuring Infra: MPLS L3Out Settings, on page 49
- Deploying Infra Configuration, on page 56

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 44 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.

Procedure

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.

Procedure

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: On-Premises Site Settings

This section describes how to configure site-specific Infra settings for on-premises sites.

Procedure

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 pane, click **Configure Infra**.

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

Step 5 In the right **<*Site*> Settings** pane, enable the **ACI Multi-Site** knob to manage the site from the Orchestrator.

Step 6 (Optional) Enable the **CloudSec Encryption** knob 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`.

Step 14 (Optional) Configure SR-MPLS settings for the site.

If the site is connected via an MPLS network, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range.

The Segment Routing Global Block (SRGB) is the range of label values reserved for Segment Routing (SR) in the Label Switching Database (LSD). These values are assigned as segment identifiers (SIDs) to SR-enabled nodes and have global significance throughout the domain.

The default range is `16000-23999`.

If you enable MPLS connectivity for the site, you will need to configure additional settings as described in Configuring Infra: MPLS L3Out Settings, on page 49.

Configuring Infra: Cloud Site Settings

This section describes how to configure site-specific Infra settings for Cloud APIC sites.

Procedure

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 pane, click **Configure Infra**.

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

Most of the information in the cloud site area is automatically populated and cannot be changed, with the exception of the BGP password field.

Step 5 In the right **<*Site*> Settings** pane, enable the **ACI Multi-Site** knob to manage the site from the Orchestrator.

Step 6 Specify the **BGP Password**.

Configuring Infra: Pod Settings

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

Procedure

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.

Procedure

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: On-Premises Site Settings, on page 45.

  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.

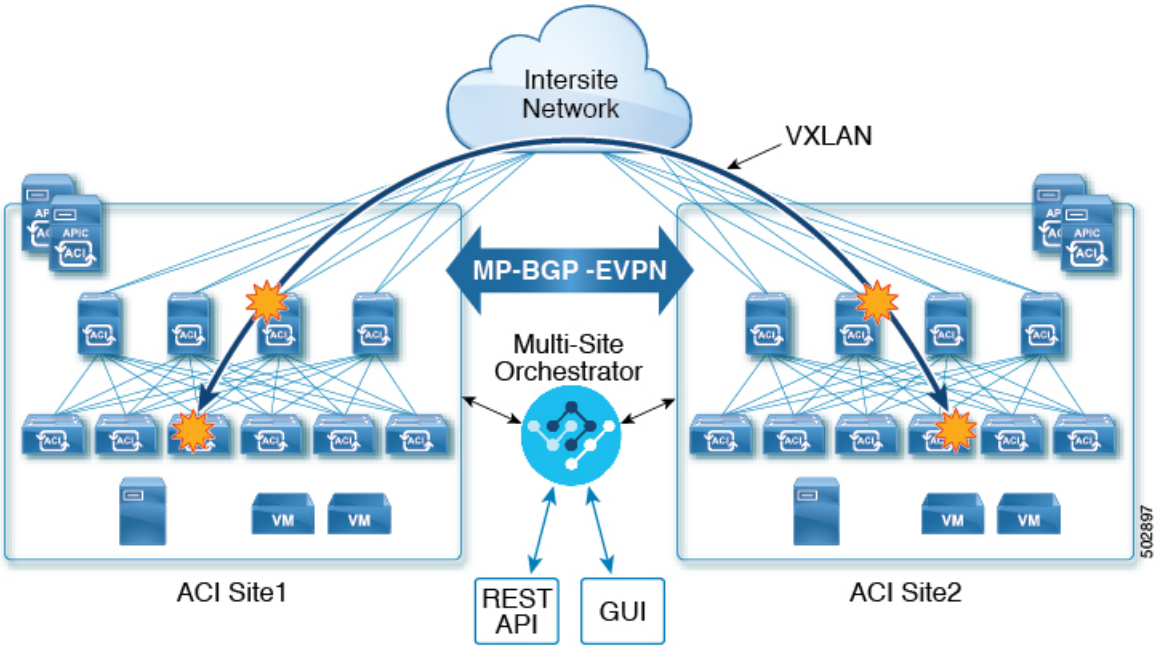
Step 10 Repeat the procedure for every spine switch.

Configuring Infra: MPLS L3Out Settings

Starting with Orchestrator Release 3.0(1) and APIC Release 5.0(1), the Multi-Site architecture supports APIC sites connected via MPLS networks.

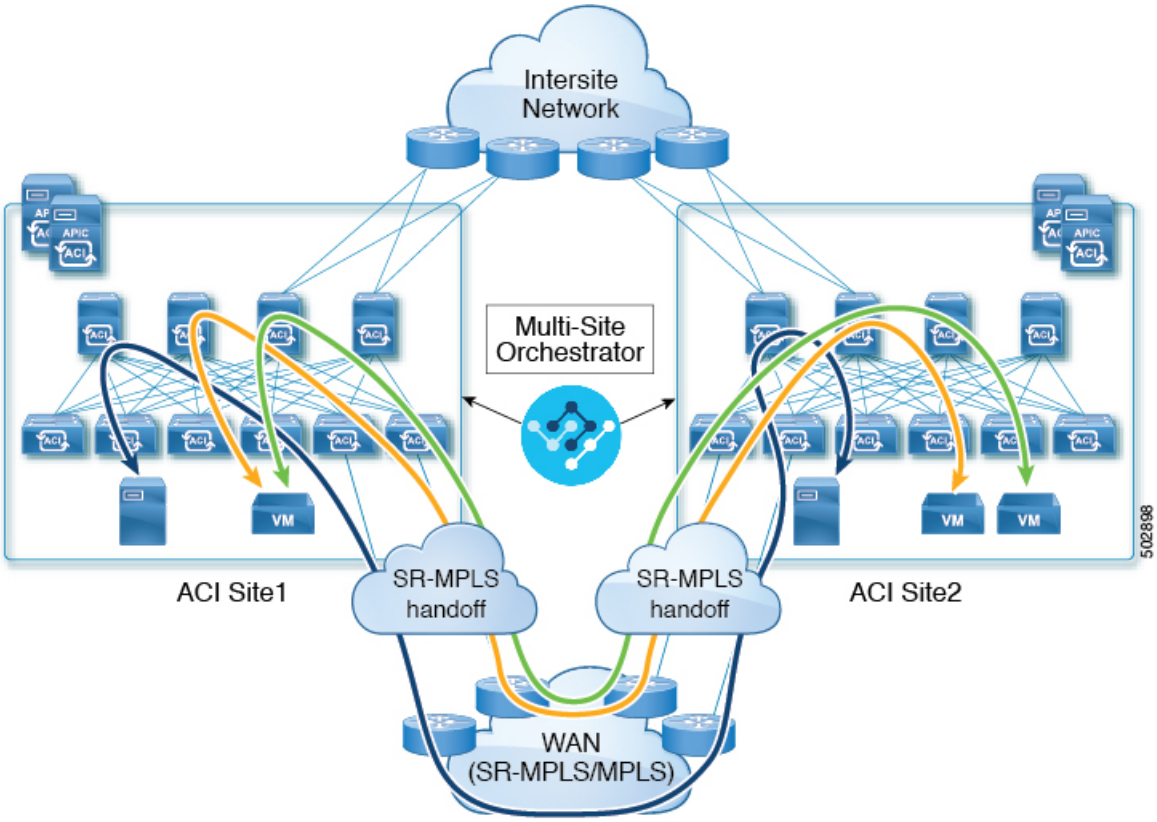
In a typical Multi-Site deployment, traffic between sites is forwarded over an intersite network (ISN) via VXLAN encapsulation:

Figure 2: Multi-Site and ISN



With Release 3.0(1), MPLS network can be used in addition to or instead of the ISN allowing inter-site L3Out communication via WAN:

Figure 3: Multi-Site and MPLS



The following sections describe guidelines, limitations, and configurations specific to managing Schemas that are deployed to these sites from the Multi-Site Orchestrator. Detailed information about MPLS hand off, supported individual site topologies (such as remote leaf support), and policy model is available in the *Cisco APIC Layer 3 Networking Configuration Guide*.

SR-MPLS Infra Guidelines and Limitations

If you want to add an APIC site that is connected to an SR-MPLS network to be managed by the Multi-Site Orchestrator, keep the following in mind:

- Any changes to the topology, such as node updates, are not reflected in the Orchestrator configuration until site configuration is refreshed, as described in Refreshing Site Connectivity Information, on page 44.
- Objects and policies deployed to a site that is connected to an SR-MPLS network cannot be stretched to other sites.

  When you create a template and specify a Tenant, you will need to enable the `SR-MPLS` option on the tenant. You will then be able to map that template only to a single ACI site.
- Tenants deployed to a site that is connected via an SR-MPLS network will have a set of unique configuration options specifically for SR-MPLS configuration. Tenant configuration is described in the "Managing Tenants" chapter of the *Cisco ACI Multi-Site Configuration Guide, Release 3.0(x)*

Supported Hardware

The SR-MPLS connectivity is supported for the following platforms:

- **Leaf switches**: The "FX", "FX2", and "GX" switch models.
- **Spine switches**:
  - Modular spine switch models with "LC-EX", "LC-FX", and "GX" at the end of the linecard names.
  - The Cisco Nexus 9000 series N9K-C9332C and N9K-C9364C fixed spine switches.
- **For sites with remote leaf switch sites, DC-PE routers**:
  - Network Convergence System (NCS) 5500 Series
  - ASR 9000 Series
  - NCS 540 or 560 routers

SR-MPLS Infra L3Out

You will need to create an SR-MPLS Infra L3Out for the fabrics connected to SR-MPLS networks as described in the following sections. When creating an SR-MPLS Infra L3Out, the following restrictions apply:

- Each SR-MPLS Infra L3Out must have a unique name.
- You can have multiple SR-MPLS infra L3Outs connecting to different routing domains, where the same border leaf switch can be in more than one L3Out, and you can have different import and export routing policies for the VRFs toward each routing domain.
- Even though a border leaf switch can be in multiple SR-MPLS infra L3Outs, a border leaf switch/provider edge router combination can only be in one SR-MPLS infra L3Out as there can be only one routing policy for a user VRF/border leaf switch/DC-PE combination.
- If there is a requirement to have SR-MPLS connectivity from multiple pods and remote locations, ensure that you have a different SR-MPLS infra L3Out in each of those pods and remote leaf locations with SR-MPLS connectivity.
- If you have a multi-pod or remote leaf topology where one of the pods is not connected directly to the SR-MPLS network, that pod's traffic destined for the SR-MPLS network will use standard IPN path to another pod, which has an SR-MPLS L3Out. Then the traffic will use the other pod's SR-MPLS L3Out to reach its destination across SR-MPLS network.
- Routes from multiple VRFs can be advertised from one SR-MPLS Infra L3Out to provider edge (PE) routers connected to the nodes in this SR-MPLS Infra L3Out.

  PE routers can be connected to the border leaf directly or through other provider (P) routers.
- The underlay configuration can be different or can be the same across multiple SR-MPLS Infra L3Outs for one location.

  For example, assume the same border leaf switch connects to PE-1 in domain 1 and PE-2 in domain 2, with the underlay connected to another provider router for both. In this case, two SR-MPLS Infra L3Outs will be created: one for PE-1 and one for PE-2. But for the underlay, it’s the same BGP peer to the provider router. Import/export route-maps will be set for EVPN session to PE-1 and PE-2 based on the corresponding route profile configuration in the user VRF.

Guidelines and Limitations for MPLS Custom QoS Policies

Following is the default MPLS QoS behavior:

- All incoming MPLS traffic on the border leaf switch is classified into QoS Level 3 (the default QoS level).
- The border leaf switch will retain the original DSCP values for traffic coming from SR-MPLS without any remarking.
- The border leaf switch will forward packets with the default MPLS EXP (0) to the SR-MPLS network.

Following are the guidelines and limitations for configuring MPLS Custom QoS policies:

- Data Plane Policers (DPP) are not supported at the SR-MPLS L3Out.
- Layer 2 DPP works in the ingress direction on the MPLS interface.
- Layer 2 DPP works in the egress direction on the MPLS interface in the absence of an egress custom MPLS QoS policy.
- VRF level policing is not supported.

Creating SR-MPLS QoS Policy

This section describes how to configure SR-MPLS QoS policy for a site that is connected via an MPLS network. If you have no such sites, you can skip this section.

SR-MPLS Custom QoS policy defines the priority of the packets coming from an SR-MPLS network while they are inside the ACI fabric based on the incoming MPLS EXP values defined in the MPLS QoS ingress policy. It also marks the CoS and MPLS EXP values of the packets leaving the ACI fabric through an MPLS interface based on IPv4 DSCP values defined in MPLS QoS egress policy.

If no custom ingress policy is defined, the default QoS Level (`Level3`) is assigned to packets inside the fabric. If no custom egress policy is defined, the default EXP value of `0` will be marked on packets leaving the fabric.

Procedure

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

Step 2 In the **Main menu**, select **Application Management** > **Policies**.

Step 3 In the main pane, select **Add Policy** > **Create QoS Policy**.

Step 4 In the **Add QoS Policy** screen, provide the name for the policy.

Step 5 Click **Add Ingress Rule** to add an ingress QoS translation rule.

These rules are applied for traffic that is ingressing the ACI fabric from an MPLS network and are used to map incoming packet's experimental bits (EXP) values to ACI QoS levels, as well as to set differentiated services code point (DSCP) values in the VXLAN header for the packet while it's inside the ACI fabric.

The values are derived at the border leaf using a custom QoS translation policy. The original DSCP values for traffic coming from SR-MPLS without any remarking. If a custom policy is not defined or not matched, default QoS Level (`Level3`) is assigned

a) In the **Match EXP From** and **Match EXP To** fields, specify the EXP range of the ingressing MPLS packet you want to match.
b) From the **Queuing Priority** dropdown, select the ACI QoS Level to map.

   This is the QoS Level you want to assign for the traffic within ACI fabric, which ACI uses to prioritize the traffic within the fabric.. The options range from Level1 to Level6. The default value is `Level3`. If you do not make a selection in this field, the traffic will automatically be assigned a `Level3` priority.

c) From the **Set DSCP** dropdown, select the DSCP value to assign to the packet when it's inside the ACI fabric.

   The DSCP value specified is set in the original traffic received from the external network, so it will be re-exposed only when the traffic is VXLAN decapsulated on the destination ACI leaf node.

   If you set the value to `Unspecified`, the original DSCP value of the packet will be retained.

d) From the **Set CoS** dropdown, select the CoS value to assign to the packet when it's inside the ACI fabric.

   The CoS value specified is set in the original traffic received from the external network, so it will be re-exposed only when the traffic is VXLAN decapsulated on the destination ACI leaf node.

   If you set the value to `Unspecified`, the original CoS value of the packet will be retained, but only if the CoS preservation option is enabled in the fabric. For more information about CoS preservation, see *Cisco APIC and QoS*.

e) Click the checkmark icon to save the rule.
f) Repeat this step for any additional ingress QoS policy rules.

Step 6 Click **Add Egress Rule** to add an egress QoS translation rule.

These rules are applied for the traffic that is leaving the ACI fabric via an MPLS L3Out and are used to map the packet's IPv4 DSCP value to the MPLS packet's EXP value as well as the internal ethernet frame's CoS value.

Classification is done at the non-border leaf switch based on existing policies used for EPG and L3Out traffic. If a custom policy is not defined or not matched, the default EXP value of `0` is marked on all labels. EXP values are marked in both, default and custom policy scenarios, and are done on all MPLS labels in the packet.

Custom MPLS egress policy can override existing EPG, L3out, and Contract QoS policies

a) Using the **Match DSCP From** and **Match DSCP To** dropdowns, specify the DSCP range of the ACI fabric packet you want to match for assigning the egressing MPLS packet's priority.
b) From the **Set MPLS EXP** dropdown, select the EXP value you want to assign to the egressing MPLS packet.
c) From the **Set CoS** dropdown, select the CoS value you want to assign to the egressing MPLS packet.
d) Click the checkmark icon to save the rule.
e) Repeat this step for any additional egress QoS policy rules.

Step 7 Click **Save** to save the QoS policy.

What to do next

After you have created the QoS policy, enable MPLS connectivity and configure MPLS L3Out as described in Creating SR-MPLS Infra L3Out, on page 53.

Creating SR-MPLS Infra L3Out

This section describes how to configure SR-MPLS L3Out settings for a site that is connected to an SR-MPLS network.

- The SR-MPLS infra L3Out is configured on the border leaf switch, which is used to set up the underlay BGP-LU and overlay MP-BGP EVPN sessions that are needed for the SR-MPLS handoff.
- An SR-MPLS infra L3Out will be scoped to a pod or a remote leaf switch site.
- Border leaf switches or remote leaf switches in one SR-MPLS infra L3Out can connect to one or more provider edge (PE) routers in one or more routing domains.
- A pod or remote leaf switch site can have one or more SR-MPLS infra L3Outs.

Before you begin

You must have:

- Added a site that is connected via SR-MPLS network as described in Adding Sites, on page 41.
- If necessary, created SR-MPLS QoS policy as described in Creating SR-MPLS QoS Policy, on page 52.

Procedure

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

Step 2 Ensure that SR-MPLS Connectivity is enabled for the site.

a) In the main navigation menu, select **Infrastructure** > **Infra Configuration**.
b) In the **Infra Configuration** view, click **Configure Infra**.
c) In the left pane, under **Sites**, select a specific site.
d) In the right **<*Site*> Settings** pane, enable the **SR-MPLS Connectivity** knob and provide the Segment Routing global block (SRGB) range

The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label.

The Segment Routing Global Block (SRGB) is the range of label values reserved for Segment Routing (SR) in the Label Switching Database (LSD). The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label.

The default range is `16000-23999`.

Step 3 In the main pane, click **+Add SR-MPLS L3Out** within a pod.

Step 4 In the right **Properties** pane, provide a name for the SR-MPLS L3Out.

Step 5 (Optional) From the **QoS Policy** dropdown, select a QoS Policy you created for SR-MPLS traffic.

Select the QoS policy you created in Creating SR-MPLS QoS Policy, on page 52.

Otherwise, if you do not assign a custom QoS policy, the following default values are assigned:

- All incoming MPLS traffic on the border leaf switch is classified into QoS Level 3 (the default QoS level).
- The border leaf switch does the following:
  - Retains the original DSCP values for traffic coming from SR-MPLS without any remarking.
  - Forwards packets to the MPLS network with the original CoS value of the tenant traffic if the CoS preservation is enabled.

- Forwards packets with the default MPLS EXP value (`0`) to the SR-MPLS network.

- In addition, the border leaf switch does not change the original DSCP values of the tenant traffic coming from the application server while forwarding to the SR network.

Step 6 From the **L3 Domain** dropdown, select the Layer 3 domain.

Step 7 Configure BGP settings.

You must provide BGP connectivity details for the BGP EVPN connection between the site's border leaf (BL) switches and the provider edge (PE) router.

a) Click **+Add BGP Connectivity**.

b) In the **Add BGP Connectivity** window, provide the details.

For the **MPLS BGP-EVPN Peer IPv4 Address** field, provide the loopback IP address of the DC-PE router, which is not necessarily the device connected directly to the border leaf.

For the **Remote AS Number**, enter a number that uniquely identifies the neighbor autonomous system of the DC-PE. the Autonomous System Number can be in 4-byte as plain format from 1 to 4294967295. Keep in mind, ACI supports only `asplain` format and not `asdot` or `asdot+` format AS numbers. For more information on ASN formats, see Explaining 4-Byte Autonomous System (AS) ASPLAIN and ASDOT Notation for Cisco IOS document.

For the **TTL** field, specify a number large enough to account for multiple hops between the border leaf and the DC-PE router, for example `10`. The allowed range `2-255` hops.

(Optional) Choose to enable the additional BGP options based on your deployment.

c) Click **Save** to save BGP settings.

d) Repeat this step to for any additional BGP connections.

Typically, you would be connecting to two DC-PE routers, so provide BGP peer information for both connections.

Step 8 Configure settings for border leaf switches and ports connected to the SR-MPLS network.

You need to provide information about the border leaf switches as well as the interface ports which connect to the SR-MPLS network.

a) Click **+Add Leaf** to add a leaf switch.

b) In the **Add Leaf** window, select the leaf switch from the **Leaf Name** dropdown.

c) Provide a valid segment ID (SID) offset.

When configuring the interface ports later in this section, you will be able to choose whether you want to enable segment routing. The SID index is configured on each node for the MPLS transport loopback. The SID index value is advertised using BGP-LU to the peer router, and the peer router uses the SID index to calculate the local label. If you plan to enable segment routing, you must specify the segment ID for this border leaf.

- The value must be within the SRGB range you configured earlier.
- The value must be the same for the selected leaf switch across all SR-MPLS L3Outs in the site.
- The same value cannot be used for more than one leaf across all sites.
- If you need to update the value, you must first delete it from all SR-MPLS L3Outs in the leaf and re-deploy the configuration. Then you can update it with the new value, followed by re-deploying the new configuration.

d) Provide the local **Router ID**.

Unique router identifier within the fabric.

e) Provide the **BGP EVPN Loopback** address.

The BGP-EVPN loopback is used for the BGP-EVPN control plane session. Use this field to configure the MP-BGP EVPN session between the EVPN loopbacks of the border leaf switch and the DC-PE to advertise the overlay prefixes. The MP-BGP EVPN sessions are established between the BP-EVPN loopback and the BGP-EVPN remote peer address (configured in the **MPLS BGP-EVPN Peer IPv4 Address** field in the **BGP Connectivity** step before).

While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

f) Provide the **MPLS Transport Loopback** address.

The MPLS transport loopback is used to build the data plane session between the ACI border leaf switch and the DC-PE, where the MPLS transport loopback becomes the next-hop for the prefixes advertised from the border leaf switches to the DC-PE routers.

While you can use a different IP address for the BGP-EVPN loopback and the MPLS transport loopback, we recommend that you use the same loopback for the BGP-EVPN and the MPLS transport loopback on the ACI border leaf switch.

g) Click **Add Interface** to provide switch interface details.

From the **Interface Type** dropdown, select whether it is a typical interface or a port channel. If you choose to use a port channel interface, it must have been already created on the APIC.

Then provide the interface, its IP address, and MTU size. If you want to use a subinterface, provide the **VLAN ID** for the sub-interface, otherwise leave the VLAN ID field blank.

In the **BGP-Label Unicast Peer IPv4 Address** and **BGP-Label Unicast Remote AS Number**, specify the BGP-LU peer information of the next hop device, which is the device connected directly to the interface. The next hop address must be part of the subnet configured for the interface.

Choose whether you want to enable segment routing (SR) MPLS.

(Optional) Choose to enable the additional BGP options based on your deployment.

Finally, click the checkmark to the right of the **Interface Type** dropdown to save interface port information.

h) Repeat the previous sub-step for all interfaces on the switch that connect to the MPLS network.

i) Click **Save** to save the leaf switch information.

Step 9 Repeat the previous step for all leaf switches connected to MPLS networks.

What to do next

After you have enabled and configured MPLS connectivity, you can create and manage Tenants, route maps, and schemas as described in the *Cisco ACI Multi-Site Configuration Guide, Release 3.0(x)*.

Deploying Infra Configuration

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

Procedure

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.

Enabling Connectivity Between On-Premises and Cloud Sites

If you have only on-premises or only cloud sites, you can skip this section.

This section describes how to enable connectivity between on-premises APIC sites and Cloud APIC sites.

By default, the Cisco Cloud APIC will deploy a pair of redundant Cisco Cloud Services Router 1000Vs. The procedures in this section creates two tunnels, one IPsec tunnel from the on-premises IPsec device to each of these Cisco Cloud Services Router 1000Vs.

The following information provides commands for Cisco Cloud Services Router 1000V as your on-premises IPsec termination device. Use similar commands if you are using a different device or platform.

Procedure

Step 1 Gather the necessary information that you will need to enable connectivity between the CSRs deployed in the cloud site and the on-premises IPsec termination device.

You can get the required configuration details using either the **Deploy & Download IPN Device config files** or the **Download IPN Device config files only** option in ACI Multi-Site Orchestrator as part of the procedures provided in Deploying Infra Configuration, on page 56.

Step 2 Log into the on-premises IPsec device.

Step 3 Configure the tunnel for the *first* CSR.

Details for the first CSR are available in the configuration files for the ISN devices you downloaded from the Multi-Site Orchestrator.

In the following example, replace:

- <*first-CSR-tunnel-ID*> with a unique tunnel ID that you assign to this tunnel.

- <*first-CSR-elastic-IP-address*> with the elastic IP address of the third network interface of the first CSR.
- <*first-CSR-preshared-key*> with the preshared key of the first CSR.
- <*interface*> with the interface that is used for connecting to the Cisco Cloud Services Router 1000V deployed in Amazon Web Services.
- <*peer-tunnel-for-onprem-IPsec-to-first-CSR*> with the peer tunnel IP address for the on-premises IPsec device to the first cloud CSR.
- <*process-id*> with the OSPF process ID.
- <*area-id*> with the OSPF area ID.

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-<first-CSR-tunnel-ID>
    pre-shared-key address <first-CSR-elastic-IP-address> key <first-CSR-preshared-key>
exit

crypto isakmp profile infra:overlay-1-<first-CSR-tunnel-ID>
    local-address <interface>
    match identity address <first-CSR-elastic-IP-address>
    keyring infra:overlay-1-<first-CSR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<first-CSR-tunnel-ID> esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-<first-CSR-tunnel-ID>
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel <first-CSR-tunnel-ID>
    ip address <peer-tunnel-for-onprem-IPsec-to-first-CSR> 255.255.255.252
    ip virtual-reassembly
    tunnel source <interface>
    tunnel destination <first-CSR-elastic-IP-address>
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-<first-CSR-tunnel-ID>
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf <process-id> area <area-id>
    no shut
exit
```

Example:

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
```

```
exit

crypto keyring infra:overlay-1-1000
    pre-shared-key address 192.0.2.20 key 12345678900987654321234567890
exit

crypto isakmp profile infra:overlay-1-1000
    local-address GigabitEthernet1
    match identity address 192.0.2.20
    keyring infra:overlay-1-1000
exit

crypto ipsec transform-set infra:overlay-1-1000 esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-1000
    set pfs group2
    set security-association lifetime seconds 86400
exit

interface tunnel 1000
    ip address 30.29.1.2 255.255.255.252
    ip virtual-reassembly
    tunnel source GigabitEthernet1
    tunnel destination 192.0.2.20
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-1000
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf 1 area 1
    no shut
exit
```

Step 4 Configure the tunnel for the *second* CSR.

Details for the second CSR are also available in the configuration files for the ISN devices you downloaded from the Multi-Site Orchestrator.

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-<second-CSR-tunnel-ID>
    pre-shared-key address <second-CSR-elastic-IP-address> key <second-CSR-preshared-key>
exit

crypto isakmp profile infra:overlay-1-<second-CSR-tunnel-ID>
    local-address <interface>
    match identity address <second-CSR-elastic-IP-address>
    keyring infra:overlay-1-<second-CSR-tunnel-ID>
exit

crypto ipsec transform-set infra:overlay-1-<second-CSR-tunnel-ID> esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-<second-CSR-tunnel-ID>
    set pfs group2
    set security-association lifetime seconds 86400
```

```
exit

interface tunnel <second-CSR-tunnel-ID>
    ip address <peer-tunnel-for-onprem-IPsec-to-second-CSR> 255.255.255.252
    ip virtual-reassembly
    tunnel source <interface>
    tunnel destination <second-CSR-elastic-IP-address>
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-<second-CSR-tunnel-ID>
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf <process-id> area <area-id>
    no shut
exit
```

Example:

```
crypto isakmp policy 1
    encryption  aes
    authentication pre-share
    group  2
    lifetime 86400
    hash sha
exit

crypto keyring infra:overlay-1-1001
    pre-shared-key address 192.0.2.21 key 12345678900987654321234567891
exit

crypto isakmp profile infra:overlay-1-1001
    local-address GigabitEthernet1
    match identity address 192.0.2.21
    keyring infra:overlay-1-1001
exit

crypto ipsec transform-set infra:overlay-1-1001 esp-aes esp-sha-hmac
    mode tunnel
exit

crypto ipsec profile infra:overlay-1-1001
    set pfs group2
    set security-association lifetime seconds 86400
exit

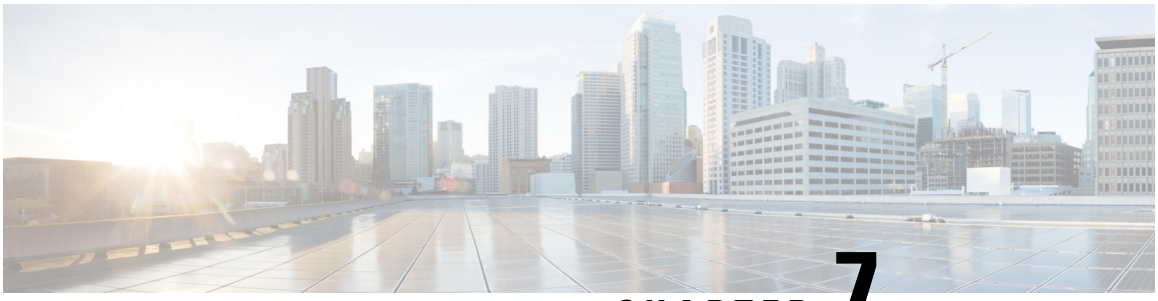
interface tunnel 1001
    ip address 30.29.1.6 255.255.255.252
    ip virtual-reassembly
    tunnel source GigabitEthernet1
    tunnel destination 192.0.2.21
    tunnel mode ipsec ipv4
    tunnel protection ipsec profile infra:overlay-1-1001
    ip mtu 1476
    ip tcp adjust-mss 1460
    ip ospf 1 area 1
    no shut
exit
```

Step 5 Repeat these steps for any additional CSRs that you need to configure.

Step 6 Verify that the tunnels are up on your on-premises IPsec device.

Use the following command to display the status. If you do not see that both tunnels are shown as up, verify the information that you entered in the steps in this section to determine where you might have an issue. Do not proceed to the next section until you see that both tunnels are shown as up.

```
ISN_CSR# show ip interface brief | include Tunnel
Interface              IP-Address      OK? Method Status                Protocol
Tunnel1000             30.29.1.2         YES manual up                    up
Tunnel1001             30.29.1.4         YES manual up                    up
```



CHAPTER 7

CloudSec Encryption

- Cisco ACI CloudSec Encryption, on page 63
- Requirements and Guidelines, on page 64
- CloudSec Encryption Terminology, on page 66
- CloudSec Encryption and Decryption Handling, on page 67
- CloudSec Encryption Key Allocation and Distribution, on page 68
- Configuring Cisco APIC for CloudSec Encryption, on page 70
- Enabling CloudSec Encryption Using Cisco ACI Multi-Site Orchestrator GUI, on page 73
- Rekey Process During Spine Switch Maintenance, on page 74

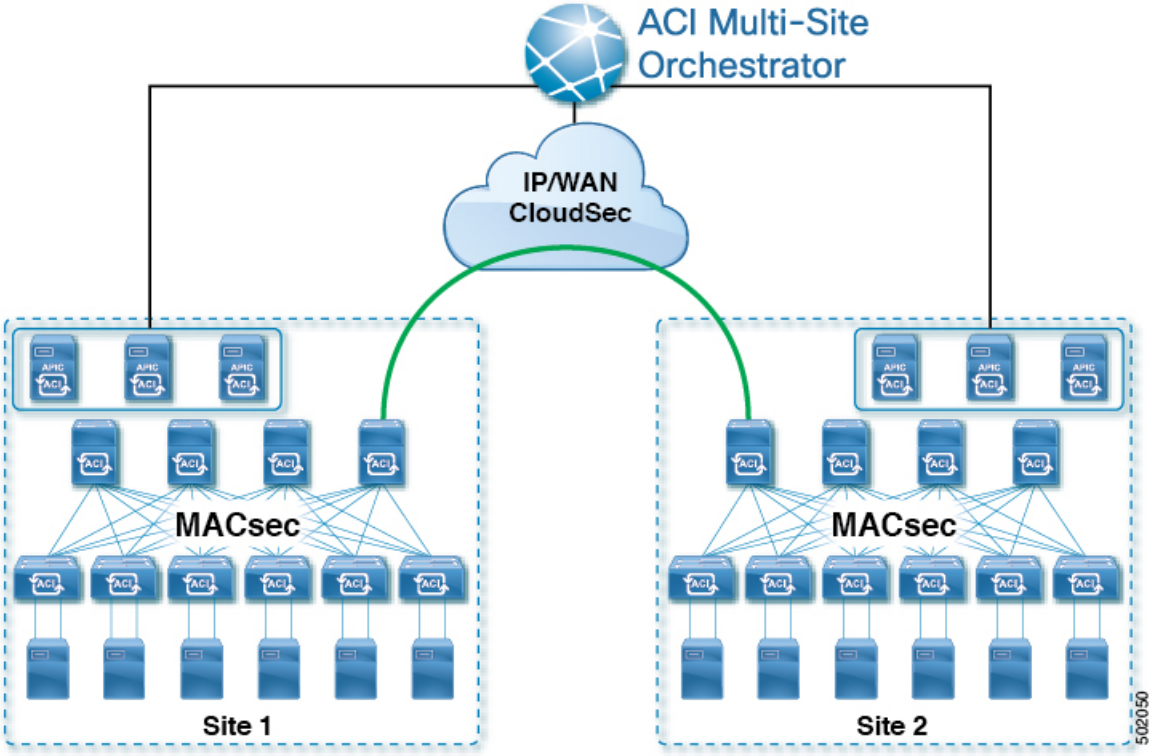
Cisco ACI CloudSec Encryption

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 4: 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.

  To configure an ASR1K router:

  1. Log in to the device.
  2. Configure the UDP ports.

     ```
     ASR1K(config)# platform cloudsec padding-fixup dst-udp-port 9999
     ```

  3. Verify the configuration.

     ```
     ASR1K# show platform software ip rp active cloudsec
       CloudSec Debug: disabled
       CloudSec UDP destination port: enabled
       1st UDP destination port: 9999
       2nd UDP destination port: 0
       3rd UDP destination port: 0

     ASR1K# show platform software ip fp active cloudsec
     ```

```
CloudSec Debug: disabled
CloudSec UDP destination port: enabled
1st UDP destination port: 9999
2nd UDP destination port: 0
3rd UDP destination port: 0
```

- 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:
  - Precision Time Protocol (PTP)
  - 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.

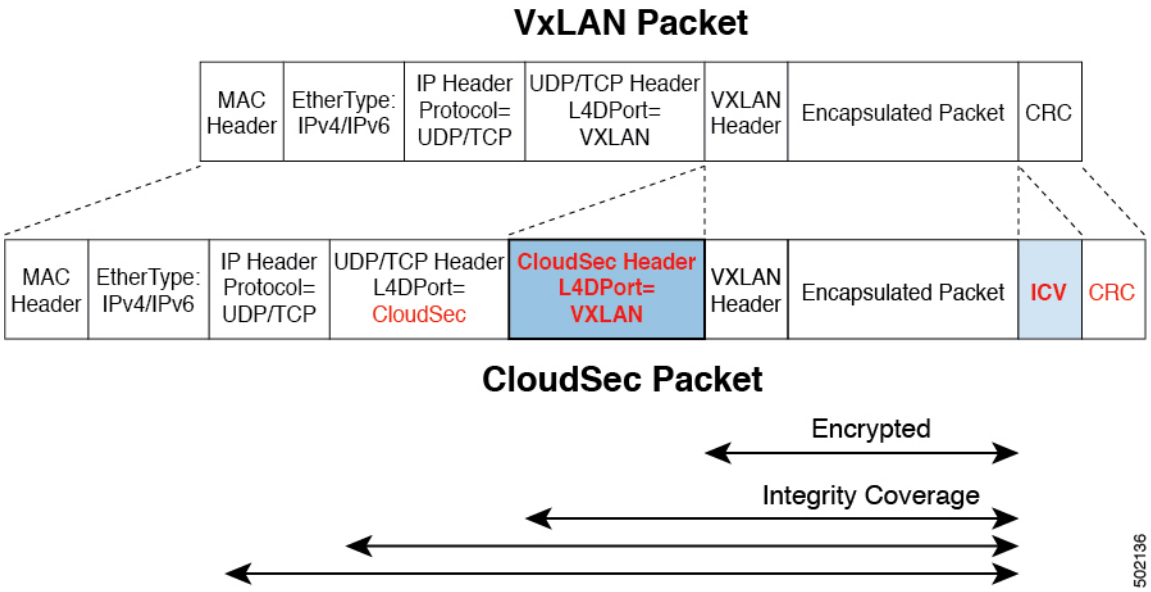
  In ACI only two association number values (`0` and `1`) are used for the two active RX keys and only one association number value (`0` or `1`) is used for the TX Key at any point in time.
- `Pre-shared key (PSK)` – One ore more keys must be configured in the Cisco APIC GUI to be used as a random seed 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 5: 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:

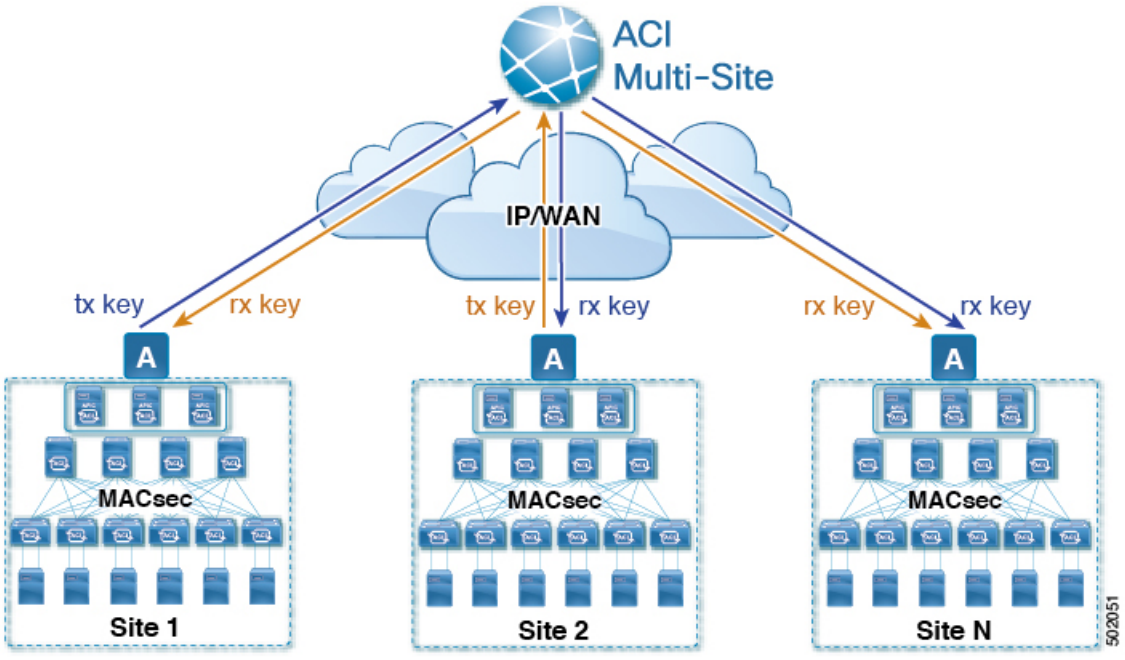
- If the received packet is a CloudSec packet, it is decrypted and the ICV is verified.

  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 6: 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 74 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 71
- Using the Cisco APIC NX-OS Style CLI, as described in Configuring Cisco APIC for CloudSec Encryption Using NX-OS Style CLI, on page 71
- Using the Cisco APIC REST API, as described in Configuring Cisco APIC for CloudSec Encryption Using REST API, on page 72

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.

Procedure

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.

Procedure

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.

Procedure

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*.

Procedure

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 74. 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 74. 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.

Procedure

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.

Procedure

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="">

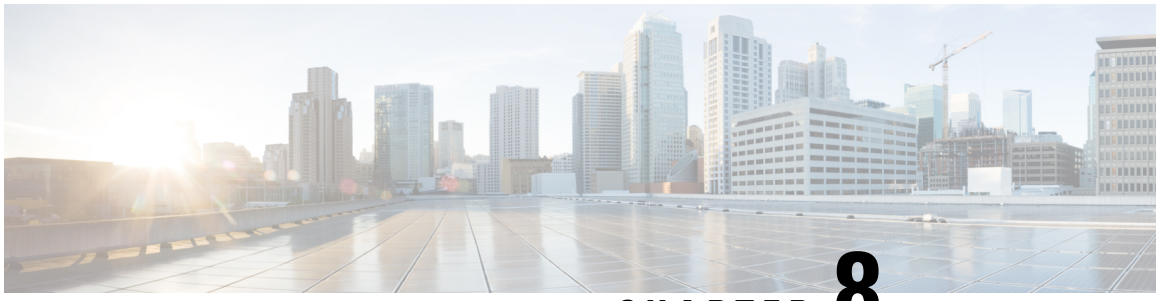
    <cloudsecIfPol descr="cloudsecifp" name="default" sakExpiryTime="10" stopRekey= "true" status=""
 />
</fvTenant>
```

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>
```



CHAPTER 8

Adding Tenants and Schemas

This chapter contains the following sections:

- Adding Tenants, on page 77
- Adding Schemas, on page 78

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.

Procedure

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

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

Step 3 In the top right of 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.

Note
If you select a site that is connected via an MPLS network, you will

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.

Adding Schemas

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

Procedure

Step 1 Log in to the Cisco ACI Multi-Site Orchestrator GUI, in the **Main menu**, click **Schemas**.

Step 2 In the **Schemas List** area, click **ADD SCHEMA**.

Step 3 In the **Untitled Schema** field, enter the new schema's name.

Step 4 Select a tenant.

In the main window pane, click **To build your schema please click here to select a tenant** then select a tenant from the **SELECT A TENANT** drop-down list.

Step 5 (Optional) Import fabric elements.

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. To import existing objects:

a) Click **IMPORT** button.
b) Select the site from which you want to import objects
c) 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.

Step 6 Add new fabric elements.

a) Click **+ Application profile**, in the **Master List**, enter the application profile name.
b) Click **+ Add EPG** field, in the **Master List**, perform the following actions:

1. In the **DISPLAY NAME** field, enter the EPG name.

2. Click **ADD SUBNET**, in the **Add Subnet** pane, perform the following actions:
   a. In the **GATEWAY IP** field, enter the gateway IP/netmask.
   b. In the **DESCRIPTION** field, enter a brief description.
   c. In the **SCOPE** section, choose **Private to VRF** or **Advertised Externally** radio button.
   d. In the **SHARED BETWEEN VRFS** section, place a check in the check box to share between VRFs.
   e. In the **NO DEFAULT SVI GATEWAY** section, place a check in the check box to not have a default SVI gateway.
   f. Click **SAVE**.
   g. Repeat 3d to create another EPG. You should have two EPGs.

c) In the **BRIDGE DOMAIN** field, from the drop-down list, choose a bridge domain or enter a bridge domain name to create one.
d) Click **+ CONTRACT** field, perform the following actions:
   1. In the **CONTRACT** field, from the drop-down list, choose a contract or enter a contract name to create one.
   2. In the **TYPE** field, from the drop-down list, choose **consumer**.
   3. Click **SAVE**.
e) Click **ADD CONTRACT** field to add a second contract, perform the following actions:
   1. In the **CONTRACT** field, from the drop-down list, choose a contract or enter a contract name to create one.
   2. In the **TYPE** field, from the drop-down list, choose **provider**.
   3. Click **SAVE**.
f) Click **+ VRF**, in the **Master List**, perform the following actions:
   1. In the **DISPLAY NAME** field, enter the VRF name.
g) Click **+ Add Bridge Domain**, in the **Master List**, perform the following actions:
   1. In the **DISPLAY NAME** field, enter the bridge domain name.
   2. In the **VIRTUAL ROUTING & FORWARDING** field, from the drop-down list, choose a VRF name or enter a VRF name to create one.
   3. In the **L2STRETCH** section, place a check in the check box to enable Layer 2 stretch.
   4. In the **INTERSITEBUMTRAFFICALLOW** section, place a check in the check box to allow intersite BUM traffic.
   5. In the **L2UNKNOWNUNICAST** field, from the drop-down list, choose **proxy** or **flood**.
   6. Click **[+] Add Subnet**, perform the following actions:
      a. In the **GATEWAY IP** field, enter the gateway IP address/netmask.
      b. In the **DESCRIPTION** field, enter a brief description of the subnet.
      c. In the **SCOPE** field, choose **Private to VRF** or **Advertised Externally**.

d. In the **SHARED BETWEEN VRFS** section, place a check in the check box to share between VRFs.

e. In the **NO DEFAULT SVI GATEWAY** section, place a check in the check box to not have a default SVI gateway.

f. In the **QUERIER** section, place a check in the check box to querier.

g. Click **OK**.

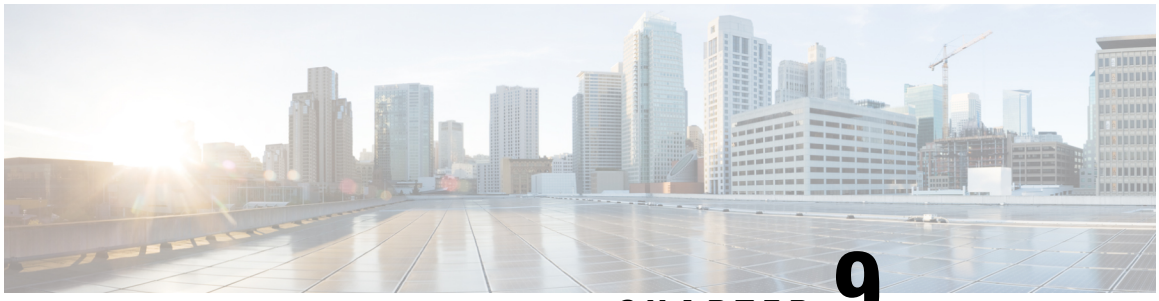
h) Click **Sites +**, place a check in the check box for each site.
i) Click **SAVE**.
j) Click **Click DEPLOY TO SITES**.



PART III

Cluster Upgrades and Downgrades

- Upgrading or Downgrading Orchestrator Deployments in Application Service Engine, on page 83
- Upgrading Orchestrator Deployments in VMware ESX , on page 91
- Downgrading Orchestrator Deployments in VMware ESX , on page 103



CHAPTER 9

Upgrading or Downgrading Orchestrator Deployments in Application Service Engine

This chapter contains the following sections:

- Prerequisites and Guidelines, on page 83
- Multi-Site Orchestrator and Cisco APIC Interoperability Support, on page 84
- Upgrading Multi-Site Orchestrator, Release 3.0(2) or Later, on page 85
- Upgrading Multi-Site Orchestrator, Release 3.0(1) or Earlier, on page 87

Prerequisites and Guidelines

This section describes how to upgrade or downgrade a 3-node Cisco ACI Multi-Site Orchestrator cluster that was deployed in Cisco Application Service Engine. If your Orchestrator cluster was deployed in VMware ESX VMs, see the Upgrading Orchestrator Deployments in VMware ESX , on page 91 chapter. If you deployed a single-node Orchestrator (for example, for testing purposes), the upgrade procedure differs slightly and is described in Installing Single Node Orchestrator, on page 111 chapter instead.

Before you upgrade or downgrade your Cisco ACI Multi-Site Orchestrator cluster:

- Ensure that your current Cisco ACI Multi-Site Orchestrator installation is running properly.
- If you are running Cisco ACI Multi-Site Orchestrator, Release 3.0(2) or later, follow the procedure described in Upgrading Multi-Site Orchestrator, Release 3.0(2) or Later, on page 85.

  That section provide steps specific to upgrading, however the same procedure can be used to switch to an earlier image to downgrade your installation. However, we do not recommend downgrading MSO deployments in Application Services Engine to releases prior to Release 3.0(2).

Note If you plan to downgrade from Release 3.1(1) or later and have Cloud APIC sites with Infra Tenant configuration, you must have a back up the configuration from the earlier release.

Due to Infra Tenant VRF changes, stateful downgrade from Release 3.1(1) is not supported and you will need to restore your configuration from the backup.

- If you are running Cisco ACI Multi-Site Orchestrator, Release 2.2(3)-3.0(1), follow the procedure described in Upgrading Multi-Site Orchestrator, Release 3.0(1) or Earlier, on page 87.
- If you are running a release prior to Release 2.2(3), your cluster was deployed in VMware ESX VMs and you must follow the Upgrading Orchestrator Deployments in VMware ESX , on page 91 chapter instead.

Multi-Site Orchestrator and Cisco APIC Interoperability Support

Multi-Site Orchestrator (MSO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the MSO itself can be upgraded independently of each other and run in mixed operation mode as long as each fabric is running APIC Release 3.2(6) or later. As such, we recommend that you always upgrade to the latest release of the Multi-Site Orchestrator.

However, keep in mind that if you upgrade the MSO before upgrading the APIC clusters in one or more sites, some of the new MSO 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 that 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)
vzAny	Release 3.2(6)
Host Based Routing	Release 4.0(1)
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)

Feature	Minimum APIC Version
Intersite L3Out	Release 4.2(1)

Upgrading Multi-Site Orchestrator, Release 3.0(2) or Later

This section describes how to upgrade Cisco ACI Multi-Site Orchestrator, Release 3.0(2) or later that is deployed on Cisco Application Service Engine.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 91

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator application.

You can download the required MSO image in one of the following two ways:

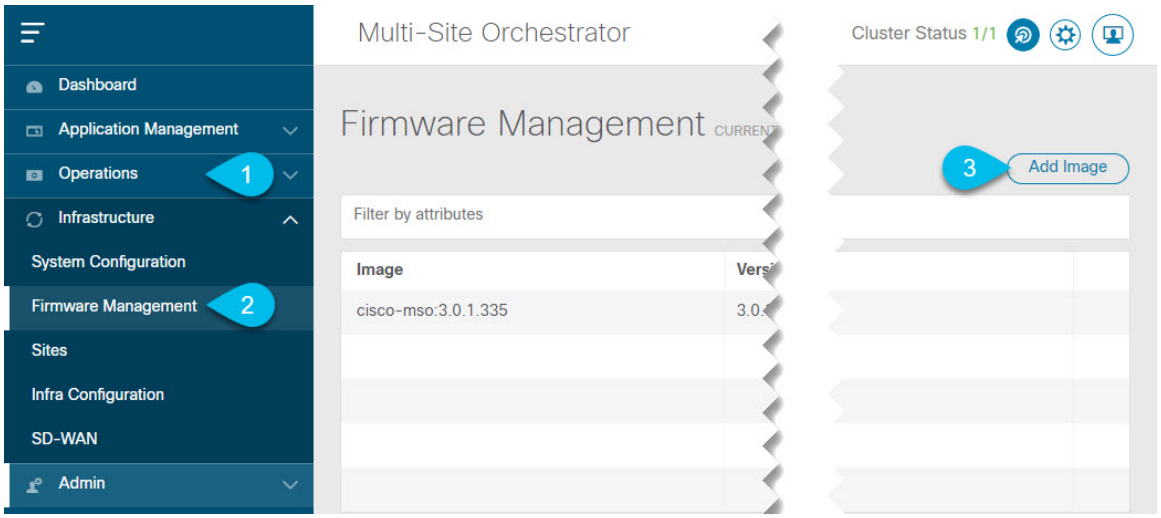
a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type
b) Click **ACI Multi-Site Software**.
c) From the left sidebar, choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site App Image* file (`Cisco-MSO-<version>.aci`) for the release.

Alternatively, you can download the image from the Cisco DC App Center:

a) Browse to the Multi-Site Orchestrator app page on DC App Center:
https://dcappcenter.cisco.com/multi-site-orchestrator.html
b) From the left sidebar **Version** dropdown, choose the Cisco ACI Multi-Site Orchestrator release version.
c) Click the **Download** button.
d) Click **Agree and download** to accept the license agreement and download the image.

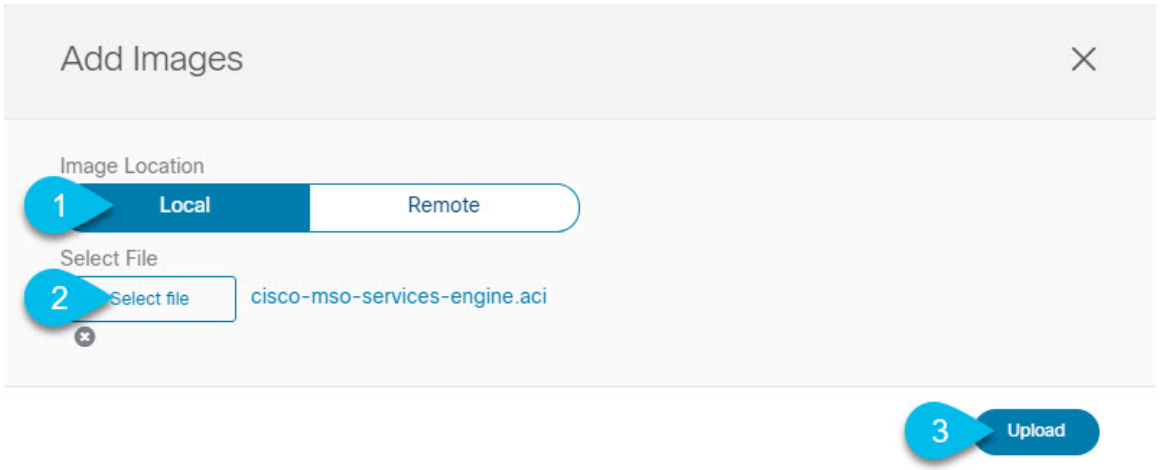
Step 2 Log in to your current Multi-Site Orchestrator GUI as a `power` user.

Step 3 Add a new image.



a) Navigate to **Operations** > **Firmware Management**.
b) Click **Add Image**.

Step 4 Select the new image.



a) In the **Add Images** window, select **Local**.

Alternatively, if you hosted the image on a web server, choose **Remote** instead.

b) Click **Select file** and select the ACI image you downloaded in the first step.

If you chose to upload a remote image, provide the file path for the image on the remote server.

c) Click **Upload** to add the image.

The image will be uploaded to the MSO Application Services Engine cluster, unpacked, processed, and made available for the upgrade or downgrade. The whole process may take several minutes and you will be able to see the status of the image.

Wait for the status to change to `Available` before proceeding to the next step.

Step 5 Activate the new image.

Ensure that the new image's status is `Available`.

a) In the main window, click the actions menu next to the image you added.
b) Then click **Activate**.
c) In the **Activation Confirmation** window, click **Continue**.

It may take up to 20 additional minutes for all the application services to start and the GUI to become available. The page will automatically reload when the process is completed.

Step 6 If you downgraded to an earlier release and had Cloud APIC sites with Infra Tenant configuration, restore the configuration from backup.

If you had Cloud APIC sites with Infra Tenant configuration, you will see schemas in an error state due to Infra Tenant VRF changes in Release 3.1(1). You will need to restore your configuration from the backup as you typically would to resolve these issues.

Upgrading Multi-Site Orchestrator, Release 3.0(1) or Earlier

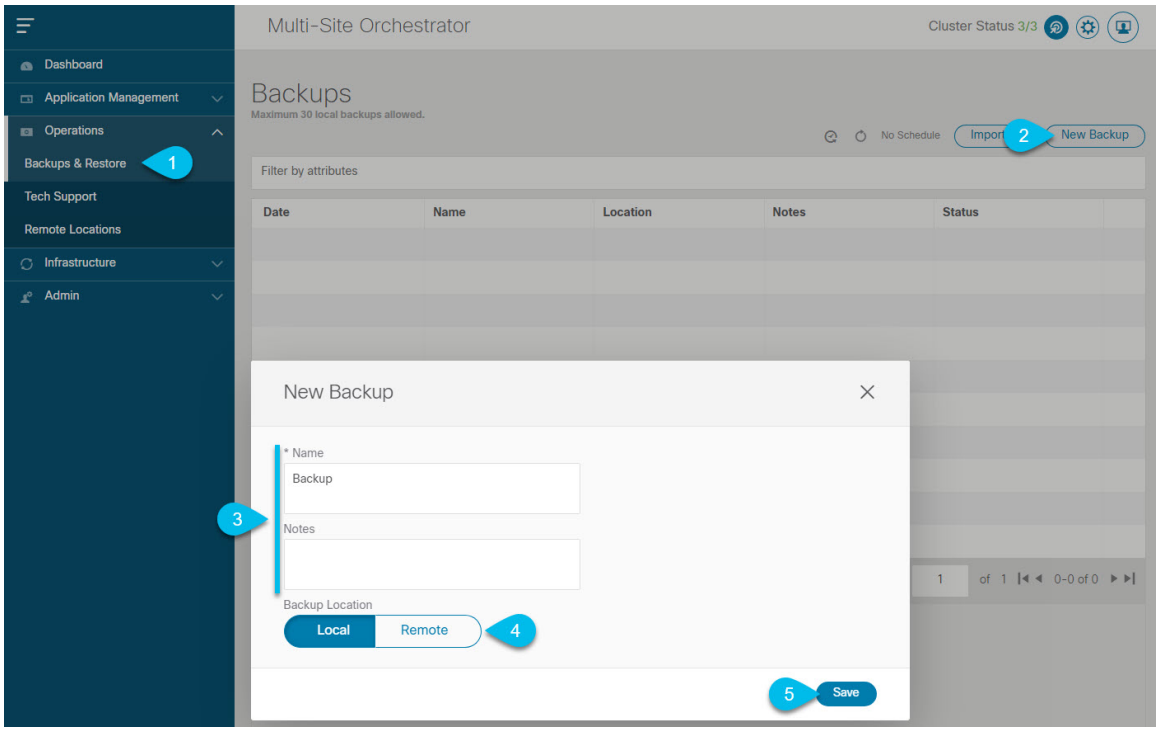
If you are running a release prior to Release 3.0(2) of Multi-Site Orchestrator in Application Services Engine, the upgrade process to Release 3.0(2) or later involves parallel deployment of a new cluster and manual transfer of the current configuration database from the existing MSO cluster.

Procedure

Step 1 Deploy a brand new Orchestrator cluster in Application Service Engine.

The procedure is described in the Deploying Multi-Site Orchestrator in Application Services Engine, on page 11 section.

Step 2 Backup existing deployment configuration.



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Operations** > **Backups & Restore**.
c) In the main window, click **New Backup**.

   A **New Backup** window opens.

d) 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 (_).

e) 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.

f) Click **Save** to create the backup.

Step 3 Copy the configuration backup file from the existing Orchestrator.

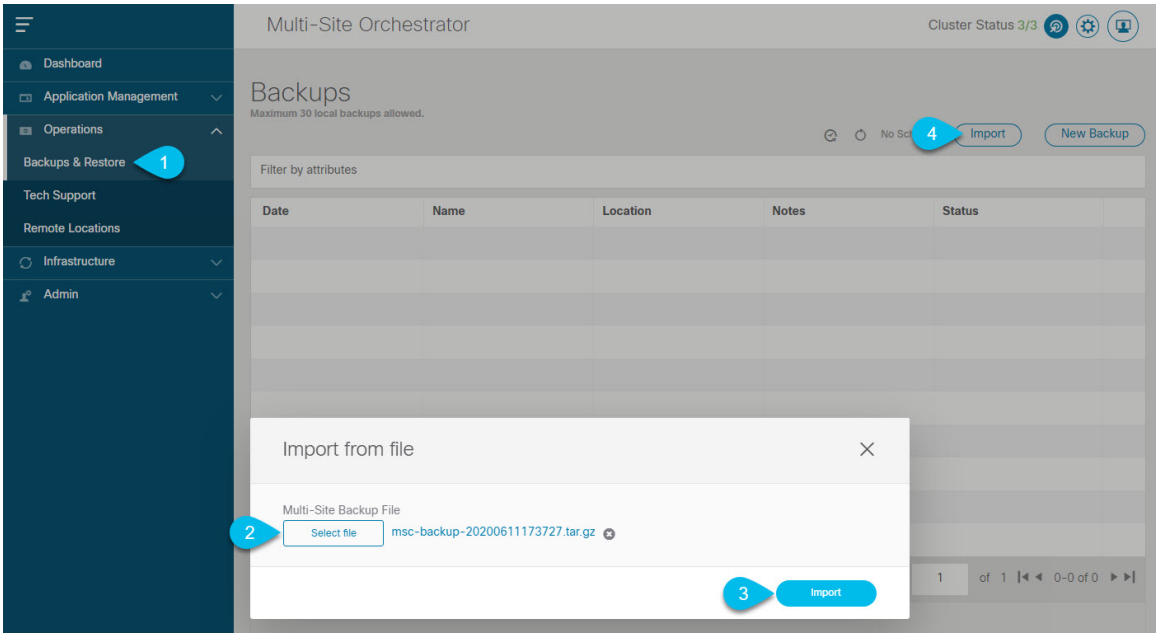
If you created the backup using a remote location, you can skip this step.

Otherwise, in the main window, click the actions ( ) icon next to the backup and select **Download**. This will download the backup file to your system.

Step 4 Bring down the existing Multi-Site Orchestrator cluster VMs.

Step 5 Import the backup file to your new Orchestrator cluster deployed on the Application Service Engine.

If you saved the backup locally, simply import the file:



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Operations** > **Backups & Restore**.
c) In the main window, click **Import**.
d) 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.

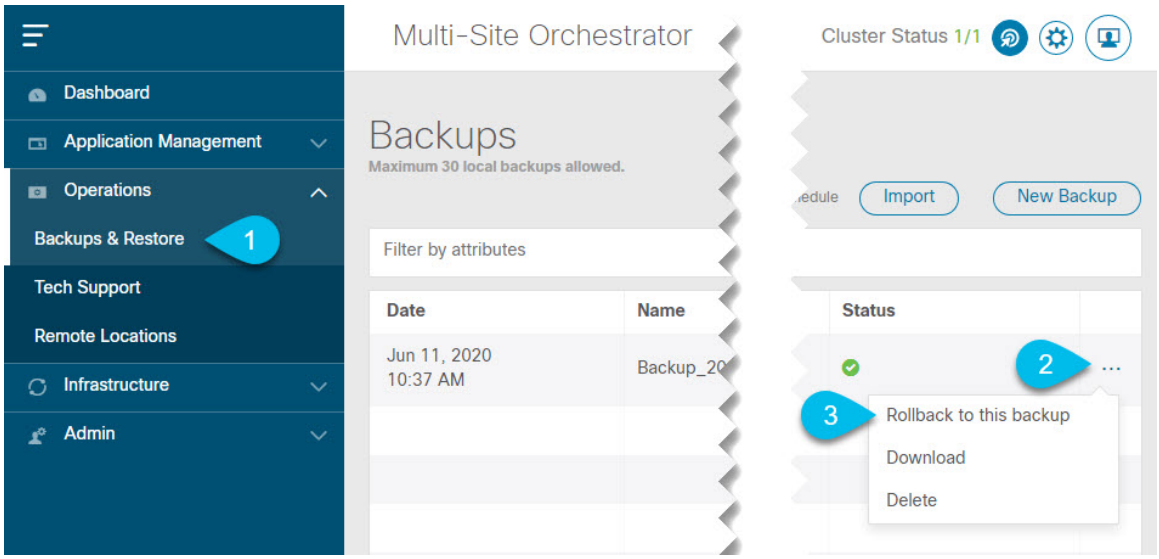
If you saved the backup to a remote location, add the remote location to the new Multi-Site Orchestrator:

a) Log in to your Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Admin** > **Remote Locations**.
c) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

d) Provide the same information for the remote location that you used in your old Orchestrator.
e) Click **Save** to add the remote server.

Step 6 Restore the configuration.



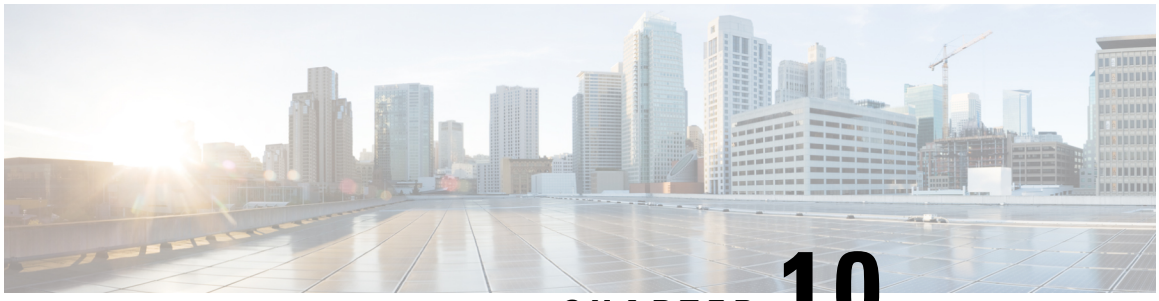
a) From the left navigation menu, select **Admin** > **Backups**.

b) 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.

c) 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.



CHAPTER 10

Upgrading Orchestrator Deployments in VMware ESX

This chapter contains the following sections:

- Prerequisites and Guidelines, on page 91
- Upgrading Cisco ACI Multi-Site Orchestrator Using Python, on page 93
- Upgrading Cisco ACI Multi-Site Orchestrator Using Backup/Restore, on page 98

Prerequisites and Guidelines

This section describes how to upgrade an existing 3-node Cisco ACI Multi-Site Orchestrator cluster that was deployed in VMware ESX VMs. If your Orchestrator cluster was deployed in Cisco Application Service Engine, see the Upgrading Multi-Site Orchestrator, Release 3.0(2) or Later, on page 85 chapter instead. If you deployed a single-node Orchestrator (for example, for testing purposes), the upgrade procedure differs slightly and is described in Upgrading Single Node Orchestrator, on page 117 chapter instead.

You can take one of the following two approaches to upgrade your cluster:

- Upgrade the existing cluster nodes in place by running a Python-based upgrade script from a remote machine (such as your laptop) which will connect to and upgrade all nodes.

  If you choose this approach, ensure that you complete the prerequisites listed below and then follow the steps described in the Upgrading Cisco ACI Multi-Site Orchestrator Using Python, on page 93 section.

- Bring up a brand new cluster and transfer a backup of the complete configuration from your existing cluster.

  If you choose this approach, ensure that you complete the prerequisites listed below and then follow the steps described in the Upgrading Cisco ACI Multi-Site Orchestrator Using Backup/Restore, on page 98 section.

Before you upgrade your Cisco ACI Multi-Site Orchestrator cluster, you must:

- Ensure that you are running at least Cisco ACI Multi-Site Orchestrator, Release 1.2(1). If you are running an earlier release, you must first upgrade it as described in *Upgrading Cisco ACI Multi-Site Orchestrator to Release 1.2(x)*.
- Ensure that your current Cisco ACI Multi-Site Orchestrator installation is running properly and each node in the cluster has at least 19 GB of free disk space.

- Ensure that all Cisco ACI Multi-Site Orchestrator node VMs have been upgraded to any new minimum CPU and RAM requirements that are listed in Prerequisites and Guidelines, on page 17.

  When upgrading the virtual machines:

  - It is recommended that all virtual machine CPU and RAM changes are done when the VM is powered down, as such we recommend updating the VMs one at a time to ensure that the cluster remains available.
  - Do not change the hard disk size of the Cisco ACI Multi-Site Orchestrator VMs.

- Ensure that you have set up the Python environment as described in Setting Up Python Environment, on page 19. (Note: This step is only necessary if you are upgrading Cisco ACI Multi-Site Orchestrator Using Python.)

  This is not required if you choose to upgrade using configuration backup and restore.

- If you are upgrading from a release prior to Release 2.1(1), configure at least 10GHz CPU cycle reservation for each Orchestrator VM.

  Specific steps are described as part of the upgrade procedure.

- If you have Cloud APIC sites with Infra Tenant configuration, you must back up your existing configuration in case you ever need to downgrade to a release prior to Release 3.1(1).

  Due to Infra Tenant VRF changes, stateful downgrade from Release 3.1(1) is not supported. If you ever downgrade from Release 3.1(1), you will need to restore your configuration from the backup.

Multi-Site Orchestrator and Cisco APIC Interoperability Support

Multi-Site Orchestrator (MSO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the MSO itself can be upgraded independently of each other and run in mixed operation mode as long as each fabric is running APIC Release 3.2(6) or later. As such, we recommend that you always upgrade to the latest release of the Multi-Site Orchestrator.

However, keep in mind that if you upgrade the MSO before upgrading the APIC clusters in one or more sites, some of the new MSO 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 that 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)

Feature	Minimum APIC Version
ACI Virtual Edge VMM Support	Release 3.2(6)
DHCP Support	Release 3.2(6)
Consistency Checker	Release 3.2(6)
vzAny	Release 3.2(6)
Host Based Routing	Release 4.0(1)
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)
Intersite L3Out	Release 4.2(1)

Upgrading Cisco ACI Multi-Site Orchestrator Using Python

The following sections describe how to prepare for and upgrade Cisco ACI Multi-Site Orchestrator using Python.

Setting Up Python Environment

This section describes how to set up the Python environment for deploying Cisco ACI Multi-Site Orchestrator using Python. You must set up the Python environment on the laptop or server from which you will run the installation scripts.



Note If you have already set up your python environment, for example for another Multi-Site deployment or upgrade, you can skip this section.

Before you begin

you will need:

- A laptop or server from which you will run the scripts.

  You must not use any of the Multi-Site Orchestrator nodes for this purpose.

- Python 3.4 or later already installed on the system from which you will run the scripts.

Procedure

Step 1 Download the **ACI Multi-Site Tools** image from Cisco ACI Multi-Site Software Download link.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.

d) Download the *ACI Multi-Site Tools Image* file (`tools-msc-<version>.tar.gz`).

Step 2 Extract the files.

```
# tar -xvzf tools-msc-<version>.tar.gz
```

Step 3 Change to the extracted directory.

```
# cd tools-msc-<version>
```

Step 4 Verify that you are running a correct version of Python.

Depending on your operating system and the versions of Python you have installed previously, the Python 3.x executable may be associated with `python` or `python3` command.

You can use one of the following examples to confirm which executable corresponds to Python 3.4 or later.

- If you have installed Python 2.x first, the `python` command will likely be associated with that version:

  ```
  # python -V
  Python 2.7.18
  ```

  In this case, you may need to use `python3` command instead:

  ```
  # python3 -V
  Python 3.4.5
  ```

- If you have installed only Python 3.x, you can use the `python` command as well:

  ```
  # python -V
  Python 3.7.7
  ```

Note
The following steps use `python` for all commands. If your Python 3.x executable is associated with `python3`, use that command instead.

Step 5 If you plan to use a proxy to access the Internet, make sure to configure the proxy as follows:

```
# export http_proxy=<proxy-ip-address>:<proxy-port>
# export https_proxy=<proxy-ip-address>:<proxy-port>
```

Step 6 Install or update the Python package manager.

```
# python -m ensurepip
```

If the package is already installed, update it to the latest version:

```
# python -m ensurepip --upgrade
```

Step 7 (Optional) Set up Python virtual environment.

We recommend using `virutalenv` to install the packages, so they do not impact the existing packages in the system. The following steps provide a brief overview of how to set up `virutalenv`. For additional information on how to use `virtualenv`, see Installing packages using pip and virtualenv.

a) Install `virtualenv`.

```
# python -m pip install --user virtualenv
```

b) Change into the directory where you want the virtual environment files to be created.

c) Create a virtual environment.

In the following command, provide a name for the virtual environment, for example *mso-deployments*.

```
# python -m venv <env-name>
```

d) Activate the virtual environment.

You need to activate the virtual environment you created before installing the packages required for Orchestrator deployment or upgrade in the next step.

For Windows:

```
# .\<env-name>\Scripts\activate.bat
```

For Linux:

```
# source ./<env-name>/bin/activate
```

Step 8 Install the required packages.

The required packages are listed in the `requirements.txt` file.

```
# python -m pip install -r requirements.txt
```

Note
The Python installation must complete successfully. If you encounter any errors, you must address them before proceeding to the next section or the Cisco ACI Multi-Site Orchestrator Python scripts will not work.

Step 9 If you used virtual Python environment, deactivate it now.

```
# deactivate
```

Sample Upgrade Configuration File

When you upgrade Multi-Site Orchestrator using Python, you can provide all the required information as command line arguments to the upgrade script or you can specify them all in a YAML configuration file and simply provide the configuration file.

This section provides sample `msc_cfg_upgrade.yml` files for two different upgrade scenarios:

- If your Multi-Site Orchestrator cluster was deployed using the Python installation scripts and within the same vCenter, you can use your vCenter login information to automatically find and upgrade all the nodes.

```
vcenter:
  name: 192.168.142.59
  user: administrator@vsphere.local

  # Update script will look for VMs with this prefix and Orchestrator label
```

```
  vm_name_prefix: msc


update:
  # Action can be 'upgrade' or 'downgrade'
  action: upgrade

  # Path to the Orchestrator upgrade image file
  msc_tgz_file: ~/tmp/msc-<version>.tar.gz
```

- If your Multi-Site Orchestrator cluster was deployed using the OVA, or in multiple different vCenters, or you simply do not wish to provide the vCenter login information, you can specify each node's IP address explicitly.

```
node1:
  ip: 192.64.136.204

node2:
  ip: 192.64.136.205

node3:
  ip: 192.64.136.206

update:
  # Action can be 'upgrade' or 'downgrade'
  action: upgrade

  # Path to the Orchestrator upgrade image file
  msc_tgz_file: ~/tmp/msc-<version>.tar.gz
```

Upgrading Multi-Site Orchestrator

This section describes how to upgrade Cisco ACI Multi-Site Orchestrator.

Note If you are upgrading from a release prior to Release 2.0(1), due to recent PSIRT updates the Cisco ACI Multi-Site Orchestrator nodes' kernels must be updated during the upgrade to Release 2.1(1) or later. This kernel update requires the nodes to be reloaded prior to performing the Orchestrator software upgrade. The python script performs the necessary update and reload automatically, followed by the Orchestrator software upgrade.

Before you begin

- Ensure that you have completed the prerequisites described in Prerequisites and Guidelines, on page 91
- Set up the Python environment as described in Setting Up Python Environment, on page 19

Procedure

Step 1 If you are upgrading from a release prior to Release 2.1(1), configure at least 10GHz CPU cycle reservation for each Orchestrator VM.

This release of Multi-Site Orchestrator requires at least 10GHz CPU cycle reservation for each VM. New deployments of Release 2.1(1) or later apply CPU cycle reservation automatically, however if you're upgrading from an earlier release, you must manually update each Orchestrator VM's settings.

a) Log in to the vSphere client.
b) Navigate to the ESX host where your Orchestrator VMs are located.
c) Shut down one of the VMs.
d) Right click the VM and choose **Edit Settings**
e) In the **Virtual Hardware** tab, expand the **CPU** category.
f) In the **Reservation** field, enter `10 GHz`.
g) Click **OK** to save the changes.
h) Power on the VM and wait for the Orchestrator cluster to stabilize with all nodes healthy.
i) Repeat the steps for the other Orchestrator VMs.

Step 2 Download the Cisco ACI Multi-Site Orchestrator upgrade image and tools.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type
b) Click **ACI Multi-Site Software**.
c) Choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site Upgrade Image* file (`msc-<version>.tar.gz`) for the release.
e) Download the *ACI Multi-Site Tools Image* file (`tools-msc-<version>.tar.gz`) for the release.

Step 3 Extract the `tools-msc-<version>.tar.gz` file to the directory from which you want to run the upgrade.

```
# tar -xvzf tools-msc-<version>.tar.gz
```

Then change into the extracted directory:

```
# cd tools-msc-<version>
```

Step 4 (Optional) Create a configuration file for the upgrade.

We recommend creating an upgrade configuration file with your deployment and upgrade details. This allows you to provide the required information once and then re-use the file for future upgrades.

If you would rather pass all the required information directly to the upgrade script, you can skip this step.

If you choose to use a configuration file, you can copy and rename the provided `msc_cfg_upgrade_example.yml` file, create one using one of the samples provided in Sample Upgrade Configuration File, on page 95, or update one from a previous upgrade with the new image path.

Step 5 If your last upgrade was from a release prior to Release 1.2(x), update the version database.

If your current Multi-Site Orchestrator installation was a fresh install of Release 1.2(1) or later, or you have upgraded to Release 2.0(x) or later in the past, skip this step.

Otherwise, run the following command replacing *1.2.3b* with the currently installed version:

```
# /opt/cisco/msc/builds/msc_1.2.3b/bin/save_msc_version.sh 1.2.3b
```

Step 6 Upgrade the Cisco ACI Multi-Site nodes.

If you created a configuration file for the upgrade as descibed in Step 4, simply run the following command:

```
# python msc_vm_util.py -c msc_cfg_upgrade.yml
```

If you would rather specify all the information on the command line, use the following command:

```
# python msc_vm_util.py -u -f msc-<version>.tar.gz -n1ip <node1-ip> -n2ip <node2-ip> -n3ip <node3-ip>
```

Step 7 Enter the passwords when prompted.

The script creates a backup of the MongoDB before the upgrade. It then copies the upgrade image to each node and executes the upgrade scripts.

It may take several minutes for the upgrade to complete. After the upgrade is complete, you can verify that the upgrade was successful and the Cisco ACI Multi-Site Orchestrator cluster is ready for use by accessing the Orchestrator GUI.

Step 8 If you upgraded from a release prior to Release 2.1(1), log in to your Orchestrator GUI and reset the password.

Due to password requirements change in Release 2.1(1), when you first log in to the Orchestrator GUI after upgrading to Release 2.1(1) or later, you will be prompted to update your password. The new password requirements are:

- At least 12 characters
- At least 1 letter
- At least 1 number
- At least 1 special character (`*` and `space` are not allowed

Upgrading Cisco ACI Multi-Site Orchestrator Using Backup/Restore

The following section will walk you through the process of upgrading the ACI Multi-Site Orchestrator cluster using a backup and restore method.



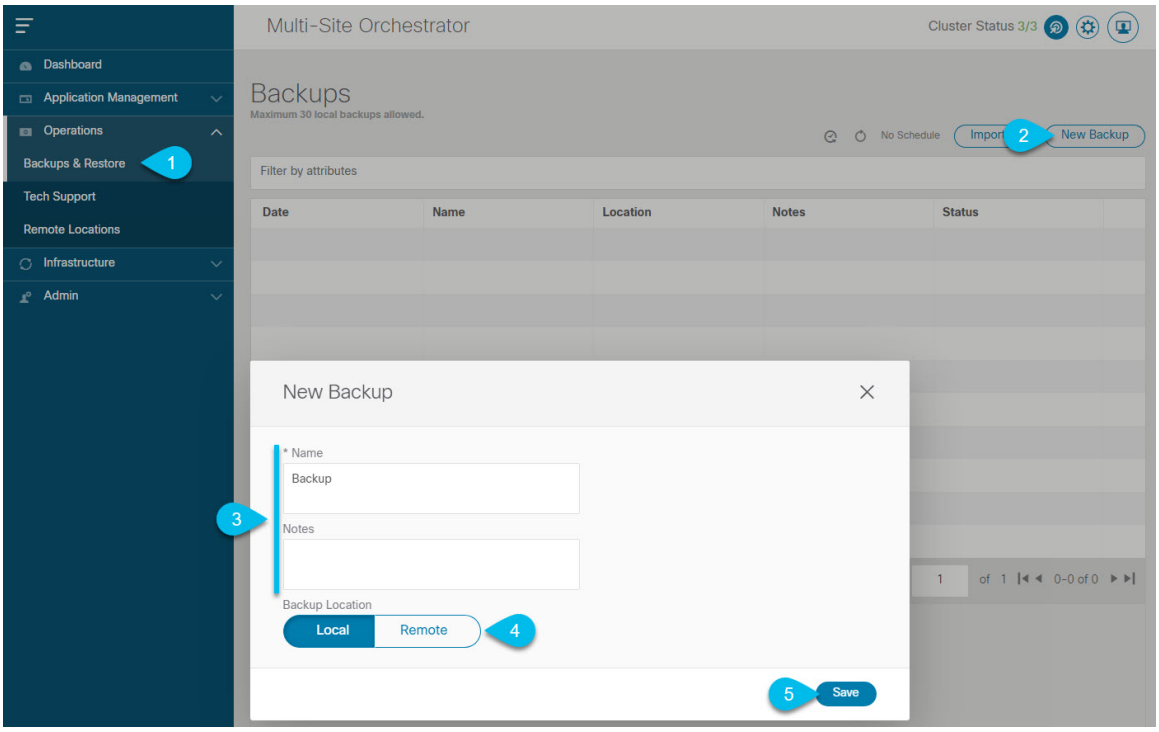
Note This involves bringing down your existing cluster and restoring the complete configuration in a brand new cluster. If you prefer to upgrade in place, follow the steps in Upgrading Cisco ACI Multi-Site Orchestrator Using Python, on page 93 instead.

Procedure

Step 1 Deploy a brand new Orchestrator cluster.

Since you are deploying a brand new cluster anyway, you can choose to continue to use the same form factor and deploy in VMware ESX as described in the Deploying in VMware ESX, on page 17 section, or you can deploy the cluster in Cisco Application Services Engine which is supported by this release. For more information about the Application Services Engine platform, see Deployment Overview, on page 5 and Deploying in Cisco Application Services Engine, on page 9.

Step 2 Backup existing deployment configuration.



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Operations** > **Backups & Restore**.
c) In the main window, click **New Backup**.

   A **New Backup** window opens.

d) 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 (_).

e) 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.

f) Click **Save** to create the backup.

Step 3 Copy the Backup file from the existing Orchestrator.

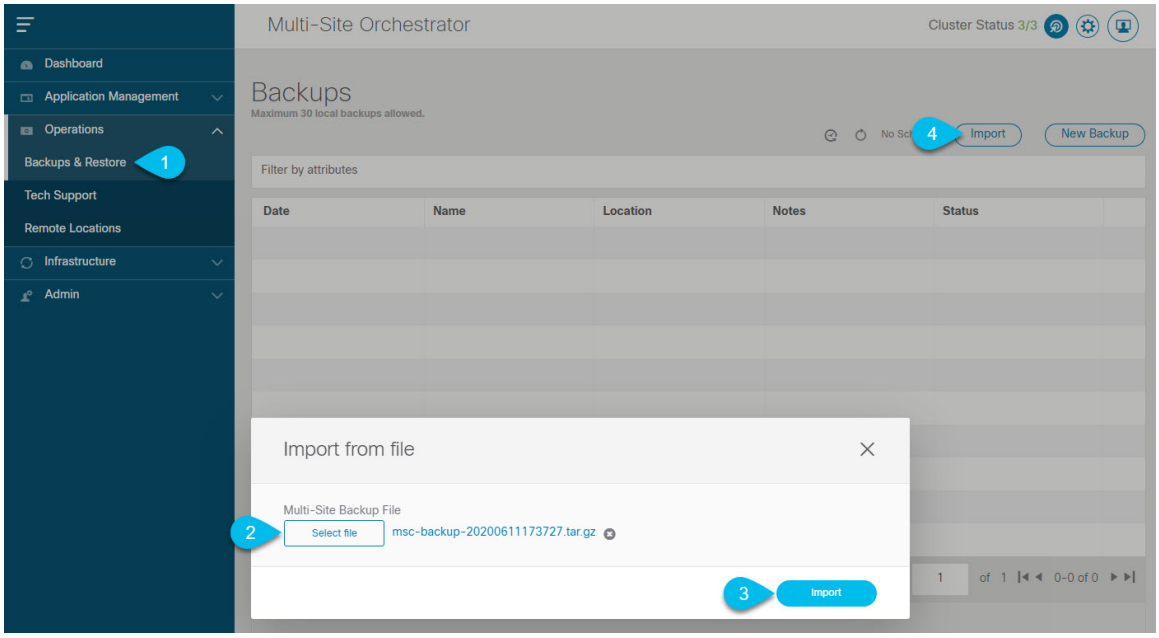
If you created the backup using a remote location, you can skip this step.

Otherwise, in the main window, click the actions ( ) icon next to the backup and select **Download**. This will download the backup file to your system.

Step 4 Bring down the existing Multi-Site Orchestrator cluster VMs.

Step 5 Import the backup file into your new Orchestrator cluster.

If you saved the backup locally, simply import the file:



a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Operations** > **Backups & Restore**.
c) In the main window, click **Import**.
d) 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.

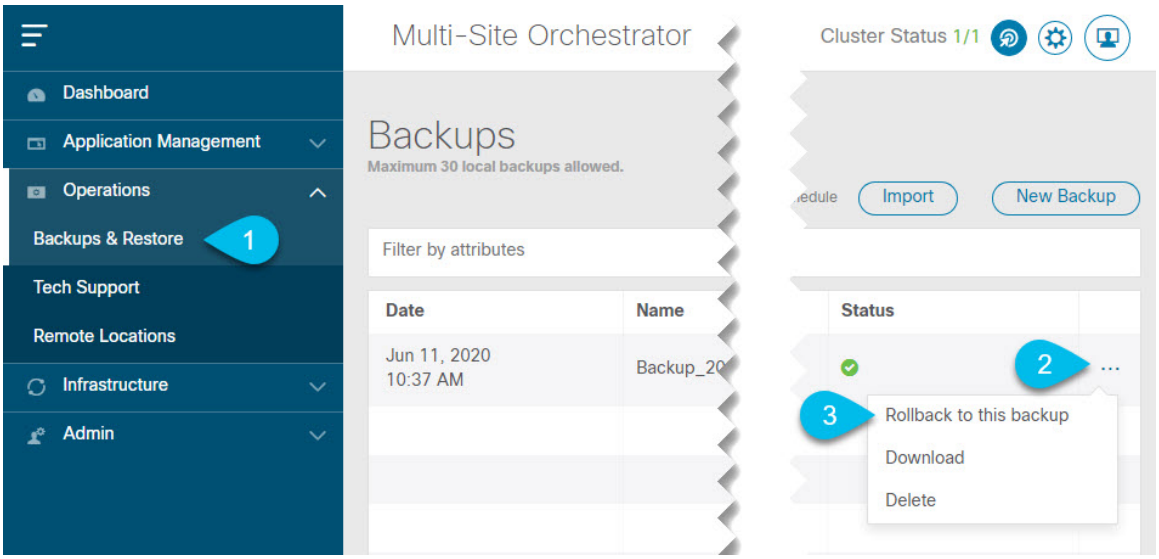
If you saved the backup to a remote location, add the remote location to the new Multi-Site Orchestrator:

a) Log in to your Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Admin** > **Remote Locations**.
c) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

d) Provide the same information for the remote location that you used in your old Orchestrator.
e) Click **Save** to add the remote server.

Step 6 Restore the configuration.



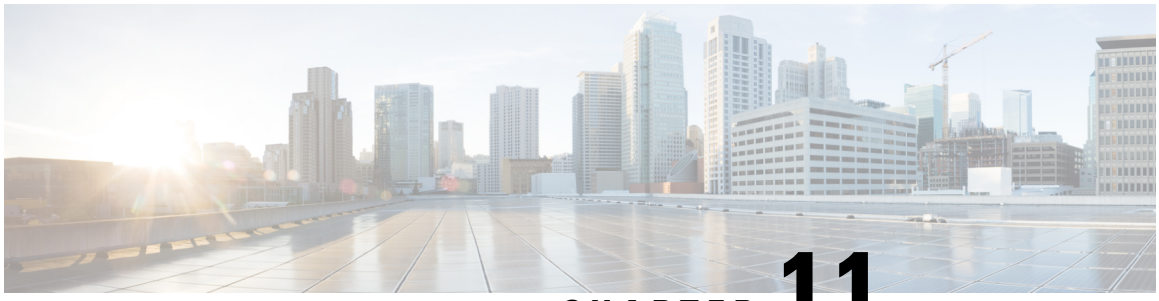
a) From the left navigation menu, select **Admin** > **Backups**.

b) 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.

c) 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.



CHAPTER 11

Downgrading Orchestrator Deployments in VMware ESX

This chapter contains the following sections:

- Downgrading Guidelines and Limitations, on page 103
- Downgrading Multi-Site Orchestrator, on page 104

Downgrading Guidelines and Limitations

Note This chapter describes how to downgrade Multi-Site Orchestrator that was deployed without using Cisco Application Service Engine. If you deployed the Orchestrator inside Application Service Engine, follow the downgrade instructions described in Upgrading or Downgrading Orchestrator Deployments in Application Service Engine, on page 83 instead.

The following list describes the guidelines and limitations for downgrading the Cisco ACI Multi-Site Orchestrator:

- If you have Cloud APIC sites with Infra Tenant configuration, you must restore from the older configuration backup after downgrading your Multi-Site Orchestrator.

  Due to Infra Tenant VRF changes, stateful downgrade from Release 3.1(1) is not supported. If you ever downgrade from Release 3.1(1), you will need to restore your configuration from the backup.

- If you plan to downgrade the Cisco APIC as well, you must downgrade Cisco ACI Multi-Site Orchestrator first.

- This release of Cisco ACI Multi-Site Orchestrator, can be downgraded to any Release 1.2(1) or later. If you plan to downgrade to an earlier release, you must first downgrade to a 1.2(x) release, then follow the instructions described in Downgrading Cisco ACI Multi-Site, Release 1.2(x) to downgrade further.

- When downgrading to a release prior to Release 2.1(1), you must remove any Cisco Cloud APIC sites you may have added to your Cisco ACI Multi-Site Orchestrator. Failing to remove the cloud sites will cause the downgrade to terminate.

- If you have configured any read-only user roles and are downgrading to a release prior to Release 2.1(2), the read-only roles will be removed from all users. This means that any user that has **only** read-only roles

will have no roles assigned to them and a Power User or User Manager will need to re-assign them new read-write roles.

In addition, if you used an external authentication server to configure the read-only user roles, you must reconfigure the authentication servers and remove those read-only user roles. The read-only user roles use a different format attribute-value (AV) string to specify read-write and read-only permissions and failing to update the configuration will cause those users to not authenticate correctly.

Additional details about external authentication servers configuration steps are described in the *Cisco ACI Multi-Site Configuration Guide*, but in short, you must update any user configuration strings from:

`cisco-av-pair=shell:msc-roles=writeRole1|writeRole2/readRole1|readRole2`

to:

`cisco-av-pair=shell:msc-roles=role1,role2`

- If you are downgrading to a release prior to Release 2.1(2), ensure that all users have the `Phone Number` field filled out. The field was required in earlier releases and any user created in Release 2.1(2) or later without a phone number provided will be unable to log into the GUI if the Orchestrator is downgraded to Release 2.1(1) or earlier. A Power User or User Manager can also update the field for any user after the downgrade.
- If you are downgrading to a release prior to Release 2.1(1), you will need to update all passwords stored by the Orchestrator, such as the passwords for all sites and authentication providers.
- Before you downgrade the Cisco ACI Multi-Site Orchestrator, remove the configuration of all features that are not supported in the release to which you are downgrading.

Downgrading Multi-Site Orchestrator

This section describes how to downgrade the Cisco ACI Multi-Site Orchestrator.

Before you begin

You must complete all the prerequisites detailed in Downgrading Guidelines and Limitations, on page 103.

Note If you have Cloud APIC sites with Infra Tenant configuration, you must have a configuration backup from the release to which you are downgrading. Due to Infra Tenant VRF changes, stateful downgrade from Release 3.1(1) is not supported. After you downgrade your Multi-Site Orchestrator, you will need to restore your configuration from the backup.

When downgrading to a release prior to Release 2.1(1), you must remove any Cisco Cloud APIC sites you may have added to your Cisco ACI Multi-Site Orchestrator. Failing to remove the cloud sites will cause the downgrade to terminate.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator downgrade (target) image.

a) Go to the Software Download link:

https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.
c) Choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site Upgrade Image* file (`msc-<version>.tar.gz`) for the release.

Step 2 Copy the downgrade image to each node.

Copy the `msc-<version>.tar.gz` file you downloaded to the `/opt/cisco/msc/builds/` directory on each node. You can use SCP or SFTP to transfer the file.

Example:

SFTP:

```
# sftp root@<node-ip>sftp> cd /opt/cisco/msc/builds/sftp> put msc-<version>.tar.gzsftp> quit
```

Example:

SCP:

```
# scp ./msc-<version>.tar.gz  root@<node-ip>:/opt/cisco/msc/builds/
```

Step 3 On each node, extract the file.

Example:

```
# cd /opt/cisco/msc/builds/# tar -xvzf msc-<version>.tar.gz
```

Step 4 On `node2` and `node3`, load the downgrade image.

On `node2` and `node3` only, run the following commands, replacing:

- <*current-version*> with the currently installed Cisco ACI Multi-Site Orchestrator release, for example *msc_2.2.1c*
- <*downgrade-version*> with the target downgrade version you downloaded and extracted in previous steps, for example *msc_1.2.1h*

Example:

```
# cd /opt/cisco/msc/builds/<current-version>/downgrade/# ./downgrade.sh <downgrade-version>
```

Step 5 From `node1`, downgrade Cisco ACI Multi-Site Orchestrator cluster.

On `node1` only, run the following commands, replacing:

- <*current-version*> with the currently installed Cisco ACI Multi-Site Orchestrator release
- <*node2-ip*> with the IP address of node2
- <*node2-password*> with the root user password for node2
- <*node3-ip*> with the IP address of node3
- <*node3-password*> with the root user password for node3
- <*downgrade-version*> with the version you are downgrading to

Note
If you leave the IP and password arguments out, the script will prompt you to enter them.

Example:

```
# cd /opt/cisco/msc/builds/<current-version>/downgrade/# ./downgrade.sh -1 <node2-ip> -2
<node2-password> -3 <node3-ip> -4 <node3-password> <downgrade-version>
```

It may take several minutes for the downgrade to complete. After the downgrade is complete, you can verify that it was successful and the Cisco ACI Multi-Site Orchestrator cluster is ready for use by accessing the Orchestrator GUI.

Step 6 If necessary, update stored passwords.

Starting with Release 2.1(1), Multi-Site Orchestrator encrypts all stored passwords, such as each site's APIC passwords and the external authentication provider passwords. As a result, when downgrading to a release prior to Release 2.1(1), you must re-enter all the password after the Orchestrator downgrade is completed.

To update APIC passwords:

a) Log in to the Orchestrator after the downgrade.
b) From the main navigation menu, select **Sites**.
c) For each site, edit its properties and re-enter its APIC password.

To update external authentication passwords

a) Login into the Orchestrator after the downgrade.
b) From the navigation menu, select **Admin** > **Providers**.
c) For each authentication provider, edit its properties and re-enter its password.

Step 7 If you had Cloud APIC sites with Infra Tenant configuration, restore the configuration from backup.

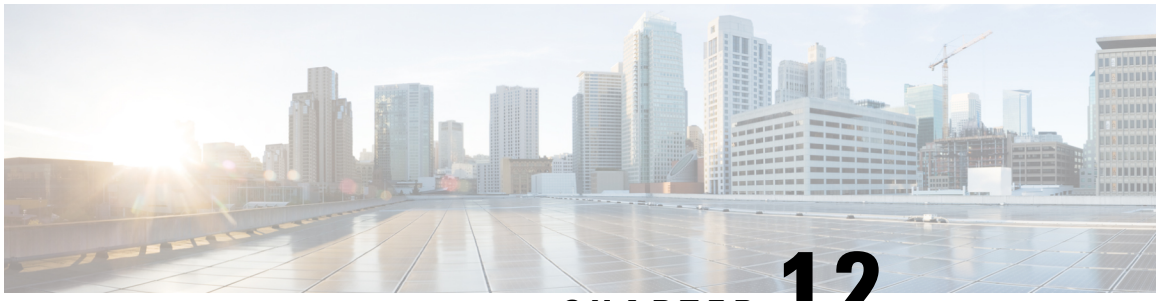
If you had Cloud APIC sites with Infra Tenant configuration, you will see schemas in an error state due to Infra Tenant VRF changes in Release 3.1(1). You will need to restore your configuration from the backup as you typically would to resolve these issues.



PART IV

Single Node Deployments

- Single Node Overview, on page 109
- Installing Single Node Orchestrator, on page 111
- Upgrading Single Node Orchestrator, on page 117
- Converting to Production Cluster, on page 123



CHAPTER 12

Single Node Overview

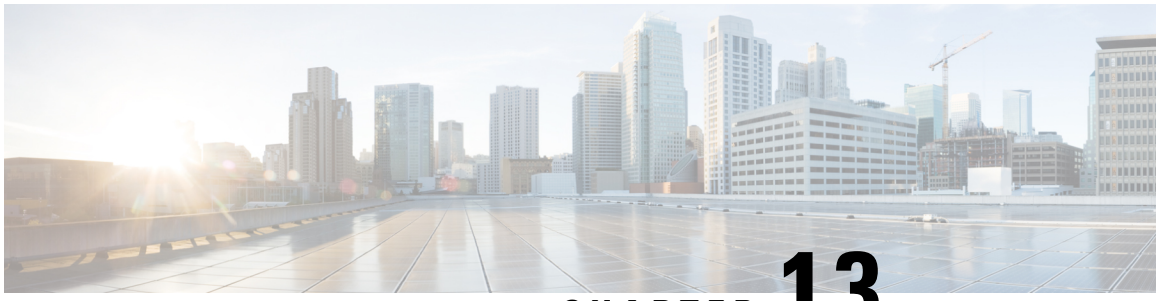
This chapter contains the following sections:

- Overview, on page 109

Overview

The following sections describe installation and upgrade procedures for single node installations. Single node is supported for lab and testing purposes only. Production deployments require 3-node high availability (HA) Orchestrator clusters and are described in other chapters in this book.

The basic requirements and workflows for single node installations and upgrades are similar to production deployments, so we encourage you to read through the Deployment Overview, on page 5 and the basic requirements of the *Cluster Deployments* section. However, certain steps may be skipped when installing a single node Orchestrator, so when you are ready, proceed with the installation or upgrade described in one of the following sections.



CHAPTER 13

Installing Single Node Orchestrator

This chapter contains the following sections:

- Installing Single Node Orchestrator in VMware ESX, on page 111
- Installing Single Node Orchestrator in Service Engine, on page 113

Installing Single Node Orchestrator in VMware ESX

This section describes how to deploy a single node Cisco ACI Multi-Site Orchestrator in an ESX VM. Single node installations are supported for testing purposes only. Production Multi-Site deployments require a 3-node Orchestrator cluster, which is described in Deployment Overview, on page 5.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator Image.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type

b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.

d) Download the *ACI Multi-Site Image* file (`msc-<version>.ova`) for the release.

Step 2 Deploy the OVA using the vCenter either the WebGUI or the vSphere Client.

Note
The OVA cannot be deployed directly in ESX, it must be deployed using vCenter.

Step 3 Configure the OVA properties.

In the **Properties** dialog box, enter the appropriate information for each VM:

- In the **Enter password** field, enter the root password for the VM.
- In the **Confirm password** field, enter the password again.
- In the **Hostname** field, enter the hostnames for each Cisco ACI Multi-Site Orchestrator node. You can use any valid Linux hostname.

- In the **Management Address** (network address) field, enter the network address or leave the field blank to obtain it via DHCP.

  Note
  The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- In the **Management Netmask** (network netmask) field, enter the netmask netmask or leave the field blank to obtain it via DHCP.
- In the **Management Gateway** (network gateway) field, enter the network gateway or leave the field blank to obtain it via DHCP.
- In the **Domain Name System Server** (DNS server) field, enter the DNS server or leave the field blank to obtain it via DHCP.
- In the **Time-zone string (Time-zone)** field, enter a valid time zone string.

  You can find the time zone string for your region in the IANA time zone database or using the `timedatectl list-timezones` Linux command. For example, *`America/Los_Angeles`* .
- In the **NTP-servers** field, enter Network Time Protocol servers separated by commas.
- In the **Application overlay** field, enter the default address pool to be used for Docker internal bridge networks.

  Application overlay must be a `/16` network. Docker then splits this network into two `/24` subnets used for the internal `bridge` and `docker_gwbridge` networks.

  For example, if you set the application overlay pool to `192.168.0.0/16`, Docker will use `192.168.0.0/24` for the `bridge` network and `192.168.1.0/24` for the `docker_gwbridge` network.

  You must ensure that the application overlay network is unique and does not overlap with any existing networks in the environment.

  Note
  The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- In the **Service overlay** field, enter the default Docker overlay network IP.

  Service overlay must be a `/24` network and is used for the `msc_msc` Orchestrator Docker service network.

  You must ensure that the service overlay network is unique and does not overlap with any existing networks in the environment.

  Note
  The field is not validated prior to installation, providing an invalid value for this field will cause the deployment to fail.

- Click **Next**.
- In the **Deployment settings** pane, check all the information you provided is correct.
- Click **Power on after deployment**.
- Click **Finish**.

In addition to the above parameters, a 10GHz CPU cycle reservation is automatically applied to each Orchestrator VM when deploying the OVA.

Step 4 Log in to the VM using SSH.

Step 5 Change into the deployment scripts directory.

```
# cd /opt/cisco/msc/builds/<build_number>/prod-standalone
```

Step 6 Run the initialization script.

```
# ./msc_cfg_init.py
```

Step 7 Run the deployment script.

```
# ./msc_deploy.py
```

Step 8 Log in to the Cisco ACI Multi-Site Orchestrator GUI.

You can access the GUI using any of the 3 nodes' IP addresses.

The default log in is **admin** and the default password is **We1come2msc!**.

When you first log in, you will be prompted to change the password.

What to do next

For more information about Day-0 Operations, see Adding Tenants and Schemas, on page 77.

Installing Single Node Orchestrator in Service Engine

This section describes how to deploy a single node Cisco ACI Multi-Site Orchestrator in Cisco Application Service Engine. Single node installations are supported for testing purposes only. Production Multi-Site deployments require a 3-node Orchestrator cluster, which is described in Deployment Overview, on page 5.

Before you begin

- You must have Cisco Application Services Engine installed and the cluster configured as described in *Cisco Application Services Engine User Guide*.

  Note that if you are deploying Services Engine in AWS, by default only PEM-based login is enabled for each node. If you'd like to be able to SSH into the nodes using a password, you will need to explicitly enable password-based logins. You can do that by logging into each node separately using the PEM file the first time, then executing the following command:

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator Image.

a) Browse to the ACI Multi-Site Orchestrator download page on Cisco DC App Center.

b) Click **Download** to download the image.

Step 2 Copy the Orchestrator image to the Application Services Engine.

If your Cisco Application Services Engine is deployed in VMware ESX (`.ova`), Linux KVM (`.qcow`), or as a physical appliance (`.iso`), or you have enabled password-based logins for your AWS (`.ami`) deployment, use the following command to copy the Orchestrator image into the `tmp` directory on the Services Engine:

```
# scp <app-local-path> rescue-user@<service-engine-ip>:/tmp/
```

However, if your Service Engine is deployed in AWS and you have not enabled password-based login, you must use the certificate (`.pem`) file that you created during the Application Services Engine deployment:

```
# scp -i <pem-file-name>.pem <app-local-path>.aci rescue-user@<service-engine-ip>:/tmp/
```

For example, assuming you're running the `scp` command from the same directory where you saved the Orchestrator image:

- For password-based authentication:

  ```
  # scp ./Cisco-MSO-2.2.3.aci rescue-user@10.30.11.147:/tmp/
  ```

- For PEM-based authentication:

  ```
  # scp -i <pem-file-name>.pem ./Cisco-MSO-2.2.3.aci rescue-user@10.30.11.147:/tmp/
  ```

Step 3 Install the Orchestrator app in your Application Services Engine.

a) Log in to your Services Engine as `rescue-user`.

If your Cisco Application Services Engine is deployed in VMware ESX (`.ova`), Linux KVM (`.qcow`), or as a physical appliance (`.iso`), simply SSH in using the following command:

```
# ssh rescue-user@<service-engine-ip>
```

However, if your Application Services Engine is deployed in AWS (`.ami`), you must login using the certificate (`.pem` file) that you created during the Application Services Engine deployment:

```
# ssh -i <pem-file-name>.pem rescue-user@<service-engine-ip>
```

b) Verify Services Engine health.

```
# acidiag health
All components are healthy
```

c) Install the Orchestrator.

In the following command, replace <*application-path*> with the full path to the application image you copied in the previous step.

```
# acidiag app install <application-path>
```

For example:

```
# acidiag app install /tmp/Cisco-MSO-2.2.3.aci
Image uploaded succesfully
check image status using: acidiag image show cisco-mso-2.2.3.aci
```

d) Verify that the application was loaded.

Use the following command to check the `operState` of the application.

While the application is loading and installing it will go through a number of operational states, which will be reflected in the `operState` field, for example `'operState': 'Initialize'`. This process can take up to 20 minutes and you must ensure that the state changes to `Disabled` before proceeding to the next step.

After the application's state changes to `Disabled`, make a note of the application's `id`, you will use it in the next step to enable the application.

```
# acidiag app show
[   {   'adminState': 'Disabled',
        'apiEntrypoint': '/query',
        'appID': 'MSO',
        'creationTimestamp': '2020-02-10T20:30:36.195960295Z',
        'description': 'Multi-Site Orchestrator application',
        'displayName': 'ACI Multi-Site Orchestrator',
        'id': 'cisco-mso:2.2.3',
        'name': 'cisco-mso',
        'operStage': 'PostInstall',
        'operState': 'Disabled',
        'schemaversion': '',
        'uiEntrypoint': '/ui/app-start.html',
        'vendorID': 'Cisco',
        'version': '2.2.3'}]
```

Step 4 Enable the Orchestrator app.

After installation is complete, the application will remain in the `Disabled` state by default and you must enable it.

In the following command, replace <*app-id*> with the application ID from the previous step:

```
# acidiag app enable <app-id>
```

For example:

```
# acidiag app enable cisco-mso:2.2.3
Application enabled succesfully
```

Step 5 Verify that the cluster was deployed successfully.

a) Verify that the application was enabled successfully.

While the application is being enabled, it will go through multiple operational states. You can use `acidiag app show` command to check the current state.

In the following output, ensure that the highlighted fields are `Enabled`, `Enable`, and `Running` respectively.

```
## acidiag app show
[   {   'adminState': 'Enabled',
        'apiEntrypoint': '/query',
        'appID': 'MSO',
        'creationTimestamp': '2020-02-10T20:30:36.195960295Z',
        'description': 'Multi-Site Orchestrator application',
        'displayName': 'ACI Multi-Site Orchestrator',
        'id': 'cisco-mso:2.2.3',
        'name': 'cisco-mso',
        'operStage': 'Enable',
        'operState': 'Running',
        'schemaversion': '',
        'uiEntrypoint': '/ui/app-start.html',
        'vendorID': 'Cisco',
        'version': '2.2.3'}]
```

b) Log in to the Cisco ACI Multi-Site Orchestrator GUI.

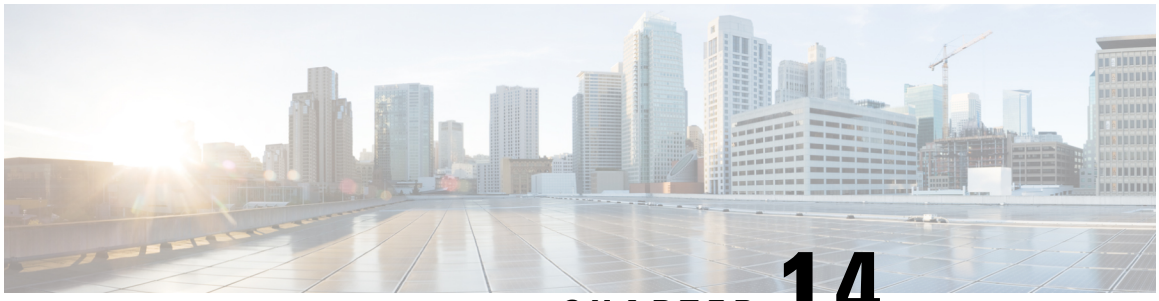
Note
After the application is enabled as described in the previous step, it may take up to 20 additional minutes for all the Orchestrator services to start and the GUI to become available.

After the GUI becomes available, you can access it by browsing to any one of your Application Services Engine nodes' IP addresses. The default log in is **admin** and the default password is **We1come2msc!**.

When you first log in, you will be prompted to change the password.

What to do next

For more information about Day-0 Operations, see Adding Tenants and Schemas, on page 77.



CHAPTER 14

Upgrading Single Node Orchestrator

This chapter contains the following sections:

- Upgrading Single Node ESX VM, on page 117
- Upgrading Single Node Service Engine VM, on page 119

Upgrading Single Node ESX VM

This section describes how to upgrade single node Cisco ACI Multi-Site Orchestrator deployed in an ESX VM. If you are running a 3-node Orchestrator cluster, follow the upgrade procedure described in Upgrading Orchestrator Deployments in VMware ESX , on page 91 instead.

Procedure

Step 1 If you are upgrading from a release prior to Release 2.1(1), configure at least 10GHz CPU cycle reservation for each Orchestrator VM.

This release of Multi-Site Orchestrator requires at least 10GHz CPU cycle reservation for each VM. New deployments of Release 2.1(1) or later apply CPU cycle reservation automatically, however if you're upgrading from an earlier release, you must manually update each Orchestrator VM's settings.

a) Log in to the vSphere client.
b) Navigate to the ESX host where your Orchestrator VMs are located.
c) Shut down one of the VMs.
d) Right click the VM and choose **`Edit Settings`**
e) In the **Virtual Hardware** tab, expand the **CPU** category.
f) In the **Reservation** field, enter `10 GHz`.
g) Click **OK** to save the changes.
h) Power on the VM and wait for the Orchestrator cluster to stabilize with all nodes healthy.
i) Repeat the steps for the other Orchestrator VMs.

Step 2 Download the Cisco ACI Multi-Site Orchestrator upgrade image.

a) Browse to the Software Download link:
https://software.cisco.com/download/home/285968390/type
b) Click **ACI Multi-Site Software**.

c) Choose the Cisco ACI Multi-Site Orchestrator release version.
d) Download the *ACI Multi-Site Upgrade Image* file (`msc-<version>.tar.gz`) for the release.

Step 3 Copy the image to your Orchestrator node.

The following command copies the image into the `/tmp` directory on the node.

```
# scp msc-<version>.tar.gz <mso-ip-address>:/tmp/
```

Step 4 Log in to your Orchestrator node and extract the `tools` directory from the Orchestrator image you copied.

Note
In the following command, the image filename contains an `msc-` prefix with a dash, while the second argument has a `msc_` prefix with an underscore, because you are extracting a specific directory within the image file.

```
# cd /tmp
# tar xfz msc-<version>.tar.gz msc_<version>/tools
```

Step 5 Copy the upgrade script from the extracted directory into the Orchestrator `scripts` directory.

```
# cp msc_<version>/tools/msc_setup.py /opt/cisco/msc/scripts
```

Step 6 Change into the Orchestrator scripts directory.

```
# cd /opt/cisco/msc/scripts
```

Step 7 Run the upgrade script.

In the following command, provide the image filename you uploaded to the Orchestrator node.

```
# python3 msc_setup.py -u -f /tmp/msc-<version>.tar.gz -st
```

Note
If the kernel needs to be upgraded, the scripts will perform the upgrade and reboot.

Step 8 If the system restarts due to a kernel upgrade, re-run the upgrade script.

When you re-run the script after a reboot, add the `--pass2` argument:

```
# python3 msc_setup.py -u -f /tmp/msc-<version>.tar.gz -st --pass2
```

Step 9 Verify upgrade was successful.

The installation can take up to 15 minutes. After it is completed, use the `docker service ls` command to verify that all docker `REPLICAS` are up:

```
# docker service ls
ID             NAME                      MODE            REPLICAS
yljvxfmjt3kb   msc_auditservice          replicated      1/1
kyvaqpehau15   msc_authyldapservice      replicated      1/1
y2fh16599hi5   msc_authytacacsservice    replicated      1/1
6pajp3kjktli   msc_backupservice         replicated      1/1
9a6tnu7wwb6j   msc_cloudsecservice       replicated      1/1
dmwkjl7het8i   msc_consistencyservice    replicated      1/1
l01mbez8j4sy   msc_endpointservice       replicated      1/1
qerrp08i6hsq   msc_executionengine       replicated      1/1
vsitso4b9xu6   msc_jobschedulerservice   replicated      1/1
l1sldx735iut   msc_kong                  replicated      1/1
zk6s5f9h6l93   msc_kongdb                replicated      1/1
t4wbsstsps6r   msc_mongodb               replicated      1/1
qi5aj3zygc2w   msc_pctagvnidservice      replicated      1/1
olxke4nk7me9   msc_platformservice       replicated      1/1
tlsjms2kwl64   msc_policyservice         replicated      1/1
9owa824s83a7   msc_schemaservice         replicated      1/1
```

```
zb3dy4d7j775    msc_siteservice          replicated      1/1
miubr6yw135n    msc_syncengine           replicated      1/1
wxle65d6ag1g    msc_ui                   replicated      1/1
jr01hfpmrcbw    msc_userservice          replicated      1/1
```

Step 10 (Optional) Free up disk space by deleting old images.

Use the following command to display all docker images in the system:

```
# docker images
```

You can then delete the old Orchestrator images using the following command:

```
# docker system prune -a
```

Step 11 If you upgraded from a release prior to Release 2.1(1), log in to your Orchestrator GUI and reset the password.

Due to password requirements change in Release 2.1(1), when you first log in to the Orchestrator GUI after upgrading to Release 2.1(1) or later, you will be prompted to update your password. The new password requirements are:

- At least 12 characters
- At least 1 letter
- At least 1 number
- At least 1 special character (`*` and `space` are not allowed

Upgrading Single Node Service Engine VM

This section describes how to upgrade single node Cisco ACI Multi-Site Orchestrator deployed in Cisco Application Service Engine. If you are running a 3-node Orchestrator cluster, follow the upgrade procedure described in Upgrading or Downgrading Orchestrator Deployments in Application Service Engine, on page 83 instead.

Procedure

Step 1 Download the Cisco ACI Multi-Site Orchestrator Image.

a) Browse to the ACI Multi-Site Orchestrator download page on Cisco DC App Center.
b) Click **Download** to download the image.

Step 2 Make the image accessible by the Orchestrator.

Note
This release supports GUI image upload from an HTTP or HTTPS server only, so you must either make the image available on a web server accessible by the Orchestrator or manually upload the image to the Application Server Engine where the Orchestrator is hosted.

If you have a web server running in your environment, simply host the `.aci` image you downloaded on that server and proceed to the next step.

Otherwise, to manually upload the image:

a) Copy the application to the Application Service Engine.

If your Cisco Application Services Engine is deployed in VMware ESX (`.ova`), Linux KVM (`.qcow`), or as a physical appliance (`.iso`), or you have enabled password-based logins for your AWS (`.ami`) deployment, use the following command to copy the Orchestrator image into the `tmp` directory on the Services Engine:

```
# scp <app-local-path> rescue-user@<service-engine-ip>:/tmp/
```

However, if your Service Engine is deployed in AWS and you have not enabled password-based login, you must use the certificate (`.pem`) file that you created during the Application Services Engine deployment:

```
# scp <app-local-path>.aci -i <pem-file-name>.pem rescue-user@<service-engine-ip>:/tmp/
```

For example, assuming you're running the `scp` command from the same directory where you saved the Orchestrator image:

- For password-based authentication:

  ```
  # scp ./cisco-mso-2.2.3c.aci rescue-user@10.30.11.147:/tmp/
  ```

- For PEM-based authentication:

  ```
  # scp ./cisco-mso-2.2.3c.aci -i <pem-file-name>.pem rescue-user@10.30.11.147:/tmp/
  ```

b) Log in to your Service Engine as `rescue-user`.

If your Cisco Application Service Engine is deployed in VMware ESX (`.ova`), Linux KVM (`.qcow`), or as a physical appliance (`.iso`), simply SSH in using the following command:

```
# ssh rescue-user@<service-engine-ip>
```

However, if your Application Service Engine is deployed in AWS (`.ami`), you must login using the certificate (`.pem` file) that you created during the Application Service Engine deployment:

```
# ssh -i <pem-file-name>.pem rescue-user@<service-engine-ip>
```

c) Add the new image.

In the following command, replace <*application-path*> with the full path to the application image you copied in the previous step.

```
# acidiag app install <application-path>
```

For example:

```
# acidiag app install /tmp/cisco-mso-2.2.3c.aci
```

d) Verify that the application was loaded.

Use the following command to check the `operState` of the application.

While the application is loading and installing it will go through a number of operational states, which will be reflected in the `operState` field, for example `'operState': 'Initialize'`. This process can take up to 20 minutes and you must ensure that the state changes to `Disabled` before proceeding to the next step.

```
# acidiag app show
[   {   'adminState': 'Disabled',
        'apiEntrypoint': '/query',
        'appID': 'MSO',
        'creationTimestamp': '2020-02-10T20:30:36.195960295Z',
        'description': 'Multi-Site Orchestrator application',
        'displayName': 'ACI Multi-Site Orchestrator',
        'id': 'cisco-mso:2.2.3',
        'name': 'cisco-mso',
        'operStage': 'PostInstall',
```

```
'operState': 'Disabled',
'schemaversion': '',
'uiEntrypoint': '/ui/app-start.html',
'vendorID': 'Cisco',
'version': '2.2.3'}]
```

Step 3 Log in to your Orchestrator.

Step 4 From the left navigation pane, select **Admin** > **Firmware Management**.

Step 5 Add the new image to the Application Service Engine cluster.

Note
If you manually uploaded the image to the Service Node cluster, the image will be already available and you can skip this step.

a) In the main window, click **Add Image**.

An **Add Image** window opens.

b) In the **File Path** field, provide the URL to the new Orchestrator image.

For example, `https://www.my-web-server.com/mso/cisco-mso-2.2.3c.aci`.

c) Click **OK** to add the image.

The image will be uploaded to the Orchestrator's Service Engine nodes, unpacked, processed, and made available for the upgrade. The whole process may take several minutes and you will be able to see the status of the image.

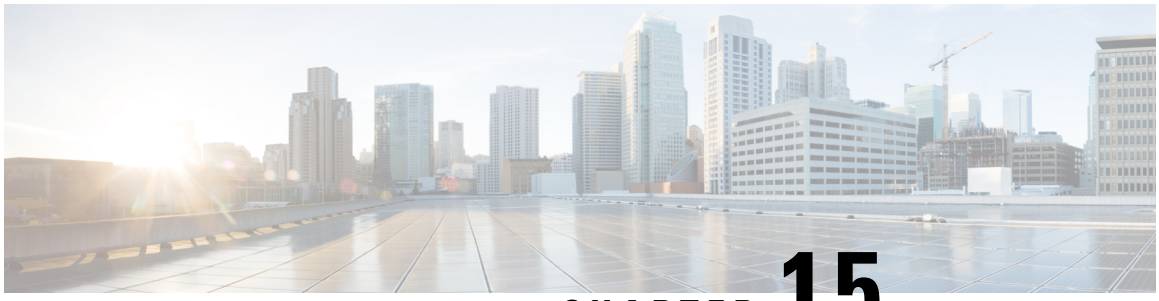
Wait for the status to change to `Available` before proceeding to the next step.

Step 6 Activate the new image.

Ensure that the new image's status is `Available`.

a) In the main window, click the actions menu next to the image you added.
b) Then click **Activate**.
c) In the **Activation Confirmation** window, click **Continue**.

Wait for the new image to be activated. The page automatically reloads when the process is completed.



CHAPTER 15

Converting to Production Cluster

This chapter contains the following sections:

- Converting Single Node to Production Cluster, on page 123

Converting Single Node to Production Cluster

The easiest way to convert your single node Orchestrator installation into a full-fledged production cluster is to simply deploy a brand new cluster and then restore existing configuration database on it.

Procedure

Step 1 Deploy a brand new Orchestrator cluster as described in *Cluster Deployments* section.

You can deploy the Multi-Site Orchestrator cluster in multiple different form factors. We recommend deploying a Service Engine Orchestrator cluster as described in Deploying in Cisco Application Services Engine, on page 9. Other available options are described in the Deployment Overview, on page 5.

Step 2 Backup existing deployment configuration.

a) Log in to your existing Cisco ACI Multi-Site Orchestrator.

b) From the left navigation pane, select **Admin** > **Backups**.

c) In the main window, click **New Backup**.

A **New Backup** window opens.

d) 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 (_).

e) 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.

f) Click **Save** to create the backup.

Step 3 Copy the Backup file from the existing Orchestrator.

If you created the backup using a remote location, you can skip this step.

Otherwise, in the main window, click the actions ( ) icon next to the backup and select **Download**. This will download the backup file to your system.

Step 4 Bring down you single node Orchestrator instance.

Step 5 Import the backup file to your new Orchestrator cluster.

If you saved the backup locally, simply import the file:

a) Log in to your existing Cisco ACI Multi-Site Orchestrator.
b) From the left navigation menu, select **Admin** > **Backups**.
c) In the main window, click **Import**.
d) 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.

If you saved the backup to a remote location, add the remote location to the new Multi-Site Orchestrator:

a) Log in to your Cisco ACI Multi-Site Orchestrator.
b) From the left navigation pane, select **Admin** > **Remote Locations**.
c) In the top right of the main window, click **Add Remote Location**.

An **Add New Remote Location** screen appears.

d) Provide the same information for the remote location that you used in your old Orchestrator.
e) Click **Save** to add the remote server.

Step 6 Restore the configuration.

a) From the left navigation menu, select **Admin** > **Backups**.
b) 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.

c) 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.