

# Dell EMC PowerEdge MX SmartFabric and Cisco ACI Integration Guide

## Abstract

This document provides the steps for integrating Dell EMC PowerEdge MX Networking switches in SmartFabric mode with the Cisco Application Centric Infrastructure (ACI) environment. It also includes steps to configure the Cisco Application Policy Infrastructure Controller (APIC).

October 2019

## Revisions

Date	Description
October 2019	Initial Release

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

© 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

# Table of contents

Revisions.....	2
<b>1 Introduction.....</b>	<b>6</b>
1.1 Dell EMC SmartFabric OS10.....	8
1.2 Cisco Application Centric Infrastructure (ACI).....	8
1.3 Typographical conventions.....	9
<b>2 Process flow and checklist.....</b>	<b>10</b>
<b>3 SmartFabric mode requirements.....</b>	<b>12</b>
3.1 PowerEdge MX requirements.....	12
3.1.1 Physically cable MX7000 chassis and upstream switches.....	12
3.1.2 Create multi-chassis management group.....	12
3.2 Application Centric Infrastructure.....	12
<b>4 SmartFabric connections to Cisco ACI leaf switches.....</b>	<b>14</b>
4.1 Validated environment.....	15
4.2 Cisco APIC configuration.....	17
4.2.1 Create a VLAN Pool.....	17
4.2.2 Create a Physical Domain.....	18
4.2.3 Create an Attachable Access Entity Profile.....	19
4.2.4 Create a Port Channel Policy.....	20
4.2.5 Create a vPC Interface Policy Group.....	20
4.2.6 Create a Leaf Access Port Policy Group.....	21
4.2.7 Create a Leaf Interface Profile.....	22
4.2.8 Create a VPC Domain Policy.....	24
4.2.9 Create a VPC Explicit Protection Group.....	24
4.2.10 Create a Leaf Profile.....	25
4.2.11 Create a Tenant.....	27
4.2.12 Create a VRF.....	28
4.2.13 Create Bridge Domains.....	29
4.2.14 Create an Application Profile.....	31
4.2.15 Create Application EPGs.....	32
4.2.16 Configure the Access Entity Profile with EPGs and VLANs.....	33
4.2.17 Create vCenter domain for Cisco ACI and Virtual Machine Manager (VMM) Domain Integration.....	35
4.2.18 Create a Contract Filter.....	39
4.2.19 Create a Contract.....	40
4.2.20 Apply the contract to the VRF.....	41

4.3	Deploy the SmartFabric .....	43
4.3.1	Define VLANs .....	43
4.3.2	LLDP setting for SmartFabric .....	44
4.3.3	Create the SmartFabric .....	44
4.3.4	Create the Uplink .....	45
4.4	Deploy servers .....	46
4.4.1	Create Server Templates .....	46
4.4.2	Add VLANs to the server templates .....	46
4.4.3	Deploy the Server Templates .....	47
4.5	vCenter configuration overview .....	48
4.6	SmartFabric connected with MX5108n Ethernet switch and Cisco ACI Leaf switches .....	50
5	Validate the configuration .....	52
5.1	MX Validation using OME-M console .....	52
5.1.1	Show the MCM group topology .....	52
5.1.2	Show the SmartFabric status .....	53
5.1.3	Show port status .....	55
5.2	Validation using the MX9116n FSE CLI .....	56
5.2.1	show switch-operating-mode .....	56
5.2.2	show discovered-expanders .....	56
5.2.3	show unit-provision .....	57
5.2.4	show vlt <i>domain-id</i> .....	57
5.2.5	show vlt <i>domain-id</i> vlt-port-detail .....	57
5.2.6	show interface port channel summary .....	58
5.2.7	show lldp neighbors .....	58
5.2.8	show qos system .....	58
5.2.9	show policy-map .....	59
5.2.10	show class-map .....	59
5.3	SmartFabric Services – Troubleshooting commands .....	60
5.3.1	show smartfabric cluster .....	60
5.3.2	show smartfabric cluster member .....	60
5.3.3	show smartfabric details .....	60
5.3.4	show smartfabric uplinks .....	61
5.3.5	show smartfabric nodes .....	61
5.3.6	show smartfabric networks .....	61
5.4	Cisco ACI validation .....	62
5.4.1	Verify vPC configuration .....	62

5.4.2	Verify physical interface configuration .....	63
5.4.3	Verify ACI learning endpoints .....	65
5.4.4	Verify ACI VMM domain integration .....	65
5.5	Verify connectivity between VMs .....	68
<b>A</b>	<b>Hardware supported in this document .....</b>	<b>69</b>
A.1	Dell EMC PowerSwitch S3048-ON management switch .....	69
A.2	Dell EMC Networking MX9116n Fabric Switching Engine (FSE).....	69
A.3	Dell EMC Networking MX5108n Ethernet switch .....	69
A.4	Cisco Nexus C93180YC-EX.....	70
A.5	Cisco Nexus C9336-PQ .....	70
<b>B</b>	<b>Validated components .....</b>	<b>71</b>
B.1	Dell EMC PowerSwitch.....	71
B.2	Dell EMC PowerEdge MX7000 chassis and components .....	71
B.3	Cisco ACI components .....	72
<b>C</b>	<b>Technical resources .....</b>	<b>73</b>
<b>D</b>	<b>Support and feedback .....</b>	<b>74</b>

# 1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company from the edge, to the core, and to the cloud. Dell EMC Networking ensures modernization for today's applications and for the emerging cloud-native world. Dell EMC is committed to disrupting the fundamental economics of the market with an open strategy that gives you the freedom of choice for networking operating systems and top-tier merchant silicon. The Dell EMC strategy enables business transformations that maximize the benefits of collaborative software and standards-based hardware, including lowered costs, flexibility, freedom, and security. Dell EMC provides further customer enablement through validated deployment guides which demonstrate these benefits while maintaining a high standard of quality, consistency, and support.

The Dell EMC PowerEdge MX is a unified, high-performance data center infrastructure. PowerEdge MX provides the agility, resiliency, and efficiency to optimize a wide variety of traditional and new, emerging data center workloads and applications. With its kinetic architecture and agile management, PowerEdge MX dynamically configures compute, storage, and fabric, increases team effectiveness, and accelerates operations. The responsive design delivers the innovation and longevity that customers need for their IT and digital business transformations.

As part of the PowerEdge MX platform, the Dell EMC SmartFabric OS10 network operating system includes SmartFabric Services. SmartFabric Services is a network automation and orchestration solution that is fully integrated with the MX Platform.

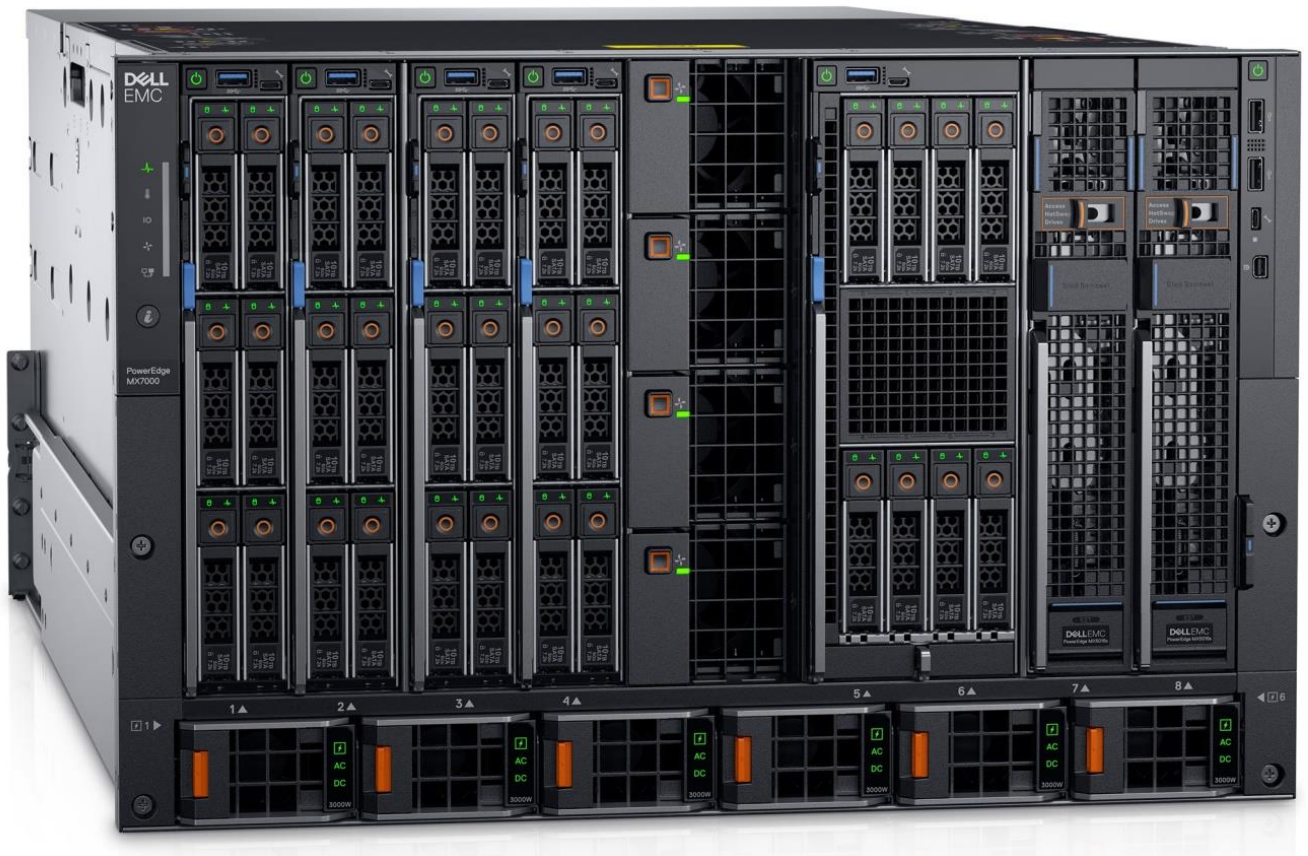


Figure 1 Dell EMC PowerEdge MX7000 chassis

This document provides examples for integrating Dell EMC PowerEdge MX platform running SmartFabric Services with Cisco Application Centric Infrastructure (ACI).

The examples in this document assume that the MX7000 chassis are configured in a multi-chassis management group and the reader has a basic understanding of the PowerEdge MX platform.

SmartFabric mode, SmartFabric Services (SFS), Full Switch mode, and Scalable Fabric are each defined in the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

---

**Note:** For an overview of hardware components supported in this document, see Appendix A.

**Note:** For a general overview of PowerEdge MX networking concepts, see the [Dell EMC PowerEdge MX Network Architecture Guide](#).

---

## 1.1 Dell EMC SmartFabric OS10

The networking market is transitioning from a closed, proprietary stack to open hardware supporting various operating systems. Dell EMC SmartFabric OS10 is designed to allow multi-layered disaggregation of the network functionality. While OS10 contributions to Open Source provide users freedom and flexibility to pick their own third-party networking, monitoring, management and orchestration applications, SmartFabric OS10 bundles industry hardened networking stack featuring standard L2 and L3 protocols over a standard and well accepted CLI interface. The MX9116n Fabric Switching Engine (FSE) and MX5108n switches in this guide use the Dell EMC SmartFabric OS10 network operating system.

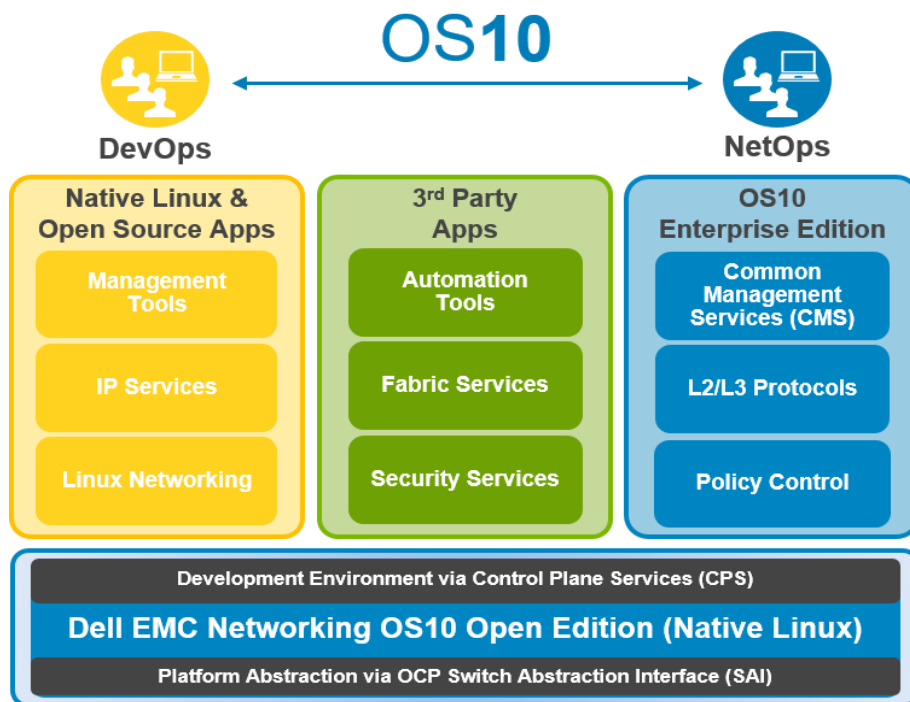


Figure 2 Dell EMC SmartFabric OS10 High-Level Architecture

**Note:** For detailed information about Dell EMC SmartFabric OS10, see [Dell EMC SmartFabric OS10 User Guide](#).

## 1.2 Cisco Application Centric Infrastructure (ACI)

Cisco ACI is an application focused, software-defined networking solution that utilizes both software and traditional switching hardware. The solution is an overlay on Cisco's high-performance switches, operating in an ACI mode managed by a controller. The Cisco Application Policy Infrastructure Controller (APIC) is a central management appliance that handles policy, visibility, security, and overall network control for the ACI environment.

Cisco ACI provides the following features within the ACI domain:

- Multi-tenant security
- Microsegmentation
- Application-specific policy management
- Network availability and QoS
- Network automation



## 1.3 Typographical conventions

The CLI and GUI examples in this document use the following conventions:

Monospace Text	CLI examples
<u>Underlined Monospace Text</u>	CLI examples that wrap the page
<i>Italic Monospace Text</i>	Variables in CLI examples
<b>Bold Monospace Text</b>	Commands entered at the CLI prompt, or to highlight information in CLI output
<b>Bold text</b>	UI elements and information entered in the GUI

## 2 Process flow and checklist

This guide is used with other documentation to configure the validated MX networking SmartFabric and Cisco ACI environment that is shown in Figure 4 on page 15.

Table 1 shows the ordered steps and locations that are referenced in the duration of this guide. Each step is covered in detail either in this guide or a link that is referenced in Table 1. The table may also be used as a checklist to ensure full coverage of all instructions in the guide.

**Note:** While some steps can be performed in a different order than shown in the table, this guide was validated using the order mentioned in Table 1.

Table 1 Step reference table and checklist

<input checked="" type="checkbox"/>	Step	Description	Reference	Where to implement
<input type="checkbox"/>	1	Physically cable the MX Chassis and upstream switches	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i>, section 3.2</a>	Hardware
<input type="checkbox"/>	2	Create multi-chassis management group	This document, section 3.1.2	OME-M
<input type="checkbox"/>	3	Deploy APIC and register Nexus leaf and spine switches	<a href="#"><i>Fabric Initialization and switch discovery</i></a> Cisco document	APIC
<input type="checkbox"/>	4	Create VLAN Pool	This document, section 4.2, step 4.2.1	APIC
<input type="checkbox"/>	5	Create a Physical Domain	This document, section 4.2, step 4.2.2	APIC
<input type="checkbox"/>	6	Create an Attachable Access Entity Profile	This document, section 4.2, step 4.2.3	APIC
<input type="checkbox"/>	7	Create a Port Channel Policy	This document, section 4.2, step 4.2.4	APIC
<input type="checkbox"/>	8	Create a vPC Interface Policy Group	This document, section 4.2, step 4.2.5	APIC
<input type="checkbox"/>	9	Create a Leaf Access Port Policy Group	This document, section 4.2, step 4.2.6	APIC
<input type="checkbox"/>	10	Create a Leaf Interface Profile	This document, section 4.2, step 4.2.7	APIC
<input type="checkbox"/>	11	Create a vPC Domain Policy	This document, section 4.2, step 4.2.8	APIC
<input type="checkbox"/>	12	Create a vPC Explicit Protection Group	This document, section 4.2, step 4.2.9	APIC
<input type="checkbox"/>	13	Create a Leaf Profile	This document, section 4.2, step 4.2.10	APIC
<input type="checkbox"/>	14	Create a Tenant	This document, section 4.2, step 4.2.11	APIC
<input type="checkbox"/>	15	Create a VRF	This document, section 4.2, step 4.2.12	APIC
<input type="checkbox"/>	16	Create Bridge Domains	This document, section 4.2, step 4.2.13	APIC
<input type="checkbox"/>	17	Create an Application Profile	This document, section 4.2, step 4.2.14	APIC
<input type="checkbox"/>	18	Create Application EPGs	This document, section 4.2, step 4.2.15	APIC

<input type="checkbox"/>	19	Configure Access Entity profile with EPGs and VLANs	This document, section 4.2, step 4.2.16	APIC
<input type="checkbox"/>	20	Create vCenter Domain for Cisco ACI and Virtual Machine Manager (VMM) domain integration:	This document, section 4.2, step 4.2.17	APIC
<input type="checkbox"/>	21	Create a contract filter	This document, section 4.2, step 4.2.18	APIC
<input type="checkbox"/>	22	Create a Contract	This document, section 4.2, step 4.2.19	APIC
<input type="checkbox"/>	23	Apply the Contract to the VRF	This document, section 4.2, step 4.2.20	APIC
<input type="checkbox"/>	24	Define VLANs	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 4.2	OME-M
<input type="checkbox"/>	25	Create a SmartFabric	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 4.3	OME-M
<input type="checkbox"/>	26	Create an Uplink	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 4.5	OME-M
<input type="checkbox"/>	27	Create Server Template	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 5.2	OME-M
<input type="checkbox"/>	28	Add VLANs to the server templates	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 5.4	OME-M
<input type="checkbox"/>	29	Deploy Server Templates	<a href="#"><i>Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide</i></a> , section 5.6	OME-M
<input type="checkbox"/>	30	Create a data center	VMware: <a href="#"><i>Create a data center</i></a>	vCenter
<input type="checkbox"/>	31	Create a cluster	VMware: <a href="#"><i>Create a cluster</i></a>	vCenter
<input type="checkbox"/>	32	Configure a cluster	VMware: <a href="#"><i>Configure a cluster</i></a>	vCenter
<input type="checkbox"/>	33	Add a host	VMware: <a href="#"><i>Add a host</i></a>	vCenter
<input type="checkbox"/>	34	Create a virtual machine	VMware: <a href="#"><i>Create a virtual machine</i></a>	vCenter
<input type="checkbox"/>	35	Create VDS and set up networking	VMware: <a href="#"><i>Setting up Networking with vSphere Distributed Switches</i></a>	vCenter

For more information about configuring VMware, see [\*Organizing your Inventory with VMware\*](#).

## 3 SmartFabric mode requirements

Before beginning SmartFabric deployment, ensure that the requirements and guidelines in this section are followed.

Configuration of SmartFabric on MX Chassis with Cisco Application Centric Infrastructure (ACI) makes the following assumptions:

- All MX7000 chassis and management modules are cabled correctly (see Section 3.1.1) and in a multi-chassis management group (see Section 3.1.2)
- The VLTi cables between switches have been connected (see Section 3.1.1)
- OME-Modular is at version 1.10.00 or later, and SmartFabric OS10 is at version 10.5.0.1 or later

---

**Note:** This document assumes that all of the server, network, and chassis hardware for the MX platform has been updated to the latest firmware, and ESXi is installed on the MX7000 compute sleds. Cisco APIC is also updated to version 4.0(3d). See Appendix B for the minimum recommended firmware versions.

---

### 3.1 PowerEdge MX requirements

#### 3.1.1 Physically cable MX7000 chassis and upstream switches

Use the following guidelines to cable the MX7000 chassis and upstream switches:

- For Management Module cabling, see the [PowerEdge MX7000 Chassis Management Networking Cabling](#) guide
- For VLTi cabling of different IOM placements, see *Switch slot placement for SmartFabric mode* in Section 3.5 of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).
- For information and requirements on cabling MX chassis to the upstream switches, see Section 3.2 of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

For more information about cabling the PowerEdge MX, see the [Dell EMC PowerEdge MX Networking Architecture Guide](#) and the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

#### 3.1.2 Create multi-chassis management group

For a scalable fabric that uses more than one MX chassis, the chassis must be in a multi-chassis management (MCM) group. See [Dell EMC OpenManage Enterprise-Modular Edition for PowerEdge MX7000 Chassis](#) for information about how to create the MCM group.

---

**Note:** SmartFabric mode can be enabled on a single chassis having two MX9116n FSEs or two MX5108n switches. For a SmartFabric implemented using a single chassis, creating an MCM group is not mandatory but recommended. The chassis must be in an MCM group for a SmartFabric containing more than one MX chassis.

---

### 3.2 Application Centric Infrastructure

Before using this guide, one or more Cisco APICs should already be deployed with the Nexus leaf and spine switches already discovered and registered with the APIC. The node ID numbers and names used in the examples in this guide are listed in Table 2.

Table 2 APIC leaf and spine node IDs and names

Node ID	Node name
101	Leaf1
102	Leaf2
201	Spine1

The networks used are shown in Table 3 along with the corresponding bridge domain and application EPG names used in APIC configuration in this guide.

Table 3 Network information

VLAN ID	VLAN name	Gateway IP address/mask	Bridge domain name	Application EPG name
1611	ESXi_Mgmt	172.16.11.254/24	ESXiMgmtBD1	ESXiMgmtEPG1
1612	vMotion	172.16.12.254/24	vMotionBD1	vMotionEPG1
1613	vSAN	172.16.13.254/24	vSANBD1	vSANEPG1
1614	web	172.16.14.254/24	webBD1	webEPG1
1615	app	172.16.15.254/24	appBD1	appEPG1
1616	db	172.16.16.254/24	dbBD1	dbEPG1

## 4 SmartFabric connections to Cisco ACI leaf switches

This chapter covers deploying a PowerEdge MX SmartFabric connected to a Cisco ACI environment. By integrating PowerEdge MX into an ACI environment, compute resources in the MX environment can use ACI gateways and access ACI resources.

The Cisco ACI environment that is validated includes a pair of Nexus C93180YC-EX switches as leaf switches as shown in Figure 3. Both C93180YC-EX leaves are connected to a single Nexus C9336-PQ spine using 40GbE uplinks shown in Figure 4. Connecting the MX9116n FSE directly to the Cisco Nexus spine in an ACI environment is not supported.

Connections from MX9116n FSE switches to C93180YC-EX leaves are 100GbE. These connections are shown in blue in Figure 3.

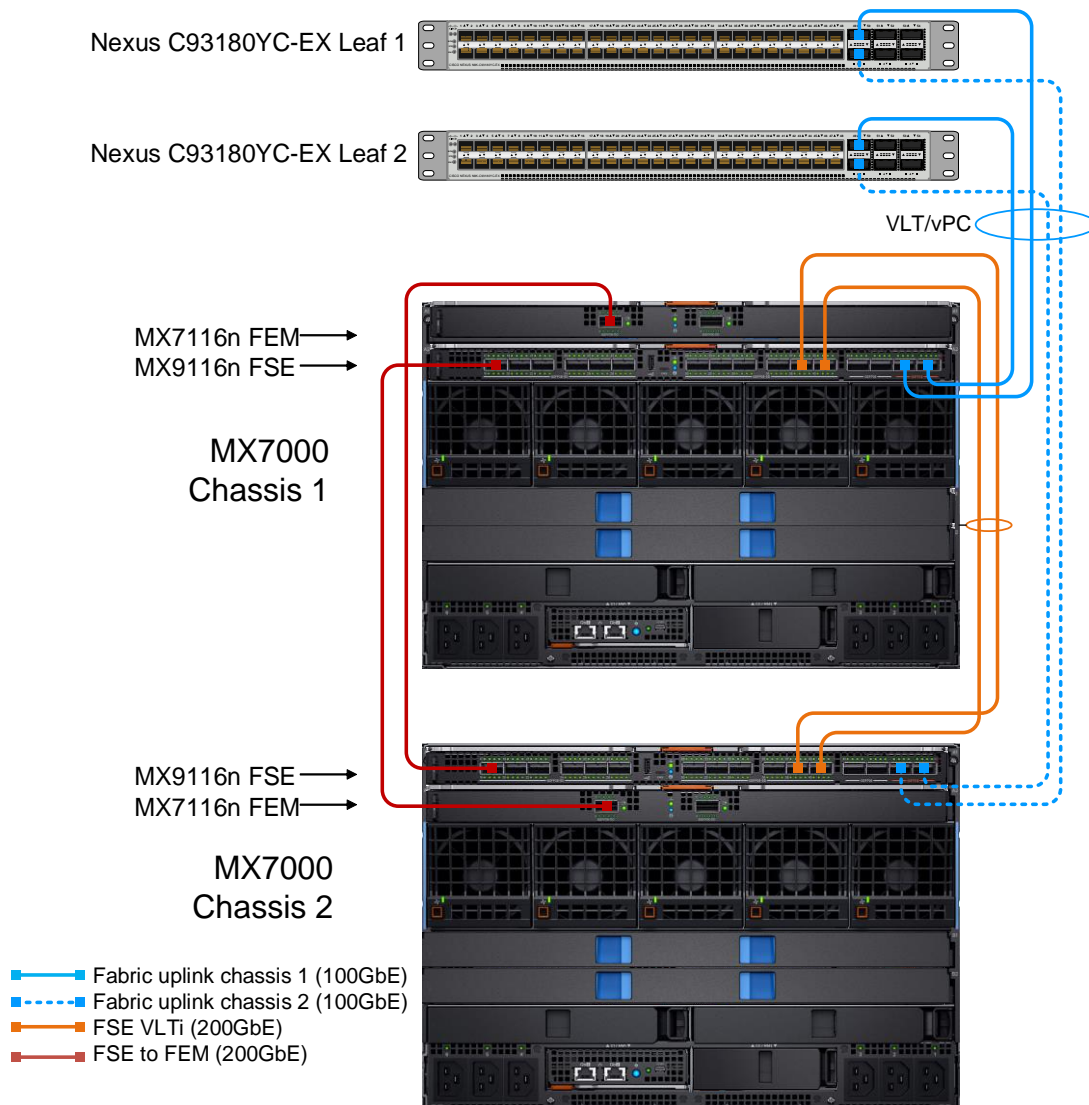


Figure 3 PowerEdge MX connected to Cisco ACI leaf switches

**Note:** For information about supported cable types for this example, such as QSFP+ and QSFP28DD, see the [PowerEdge MX I/O Guide](#) and [Dell EMC PowerEdge MX Network Architecture Guide](#).

## 4.1 Validated environment

In this scenario, two MX7000 chassis are joined to an existing Cisco ACI environment. The MX chassis environment consists of two MX9116n FSEs, two MX7116n Fabric Expander Modules (FEMs), and four MX compute sleds.

The connections between the ACI environment and the MX chassis are made using a double-sided multi-chassis link aggregation group (MLAG). The MLAG is called a vPC on the Cisco ACI side and a VLT on the PowerEdge MX side.

All devices in the validated environment that is covered in this chapter are connected as shown in Figure 4.

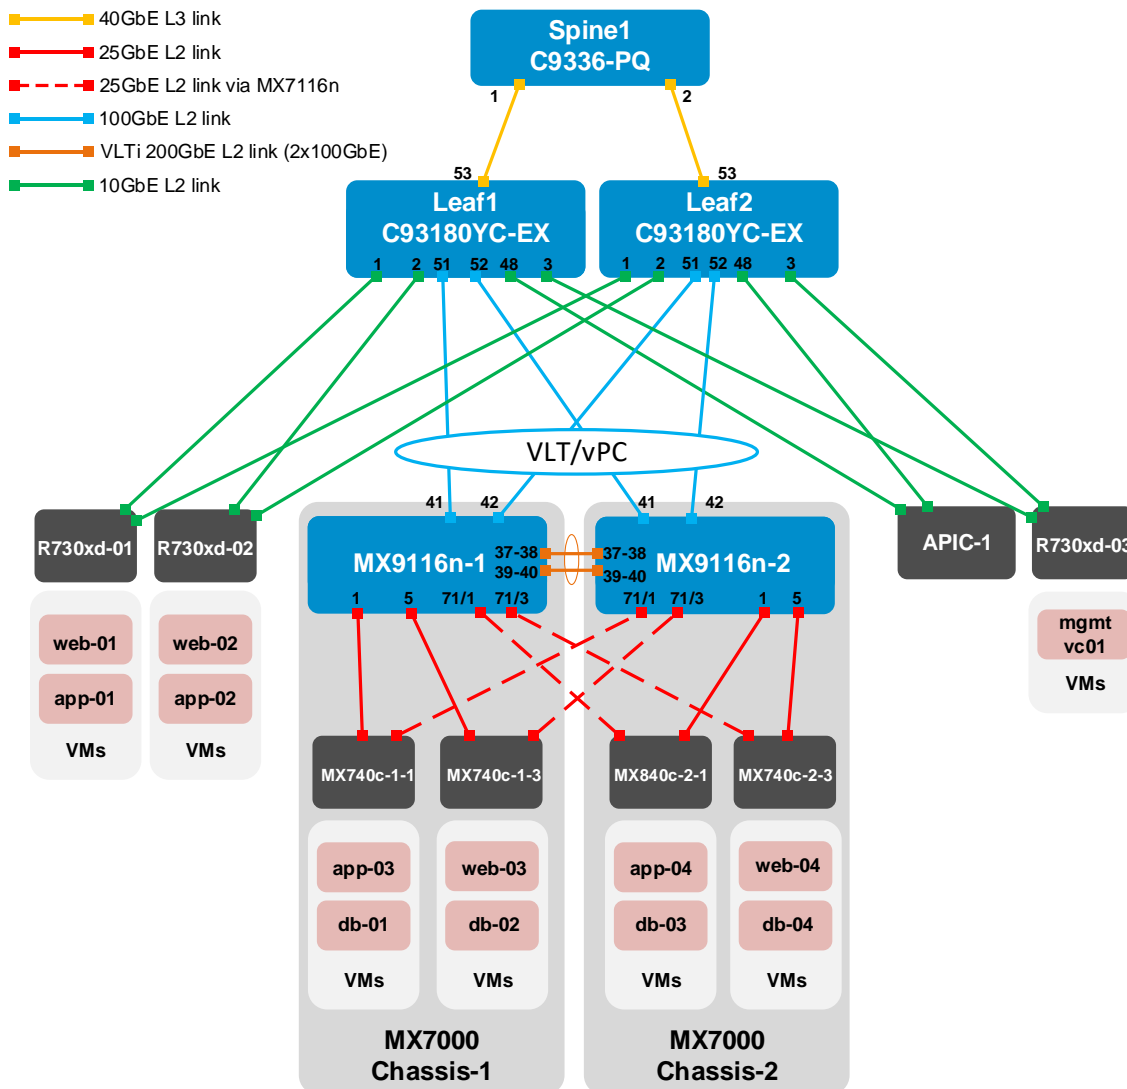


Figure 4 Validated SmartFabric and ACI environment

**Note:** The MX7116n FEMs are not shown in Figure 4 as they are transparent to the topology.

There is no peer link that is used between the Cisco ACI leaf switches. While a typical production environment has multiple Application Policy Infrastructure Controllers (APICs), for this example, a single APIC (APIC-1) is used.

All Dell EMC PowerEdge R730xd rack servers and MX compute sleds in this example are running VMware ESXi 6.7.0. To install ESXi on Dell EMC PowerEdge servers, follow the instructions on [Installation of VMware ESXi on Dell EMC PowerEdge servers](#).

VMs named “web,” “app,” and “db” on the ESXi hosts are running Ubuntu Linux guest operating systems. A third R730xd server is added to assist with vCenter configuration and is accessible over the OOB management network.

The Cisco ACI environment has three PowerEdge R730xd rack servers that are directly connected to the ACI leafs. These rack servers are in a VMware vSphere cluster, with a vCenter VM named mgmtvc01 on the R730xd-03 as shown in Figure 4.

Integrating PowerEdge MX into the Cisco ACI environment enables the MX compute sleds to join the existing VMware vSphere cluster. This enables the hosts and VMs to communicate using the relevant networks.

The environment uses the six networks that are shown in Table 4.

Table 4 Networks used

VLAN ID	VLAN name	Description	Network address	Gateway address
1611	ESXi_Mgmt	ESXi host in-band management	172.16.11.0/24	172.16.11.254
1612	vMotion	VM migration	172.16.12.0/24	172.16.12.254
1613	vSAN	Storage	172.16.13.0/24	172.16.13.254
1614	web	VM data network	172.16.14.0/24	172.16.14.254
1615	app	VM data network	172.16.15.0/24	172.16.15.254
1616	db	VM data network	172.16.16.0/24	172.16.16.254

---

**Note:** While the VMware vMotion and vSAN networks are configured in this example, their use is out of scope for this guide.

---

VMs in the validated environment use the IP addresses shown in Table 5.

Table 5 VM IP addresses

VM Name	VLAN name	IP address
mgmtvc01	ESXi_Mgmt	172.16.11.171
web01-web04	web	172.16.14.1-4
app01-app04	app	172.16.15.1-4
db01-db04	db	172.16.16.1-4



## 4.2 Cisco APIC configuration

The Cisco APIC configuration includes the ports connected to the R730xd rack servers and the vPC that connects to the MX9116n FSE VLT port channel. This includes configuration of the ACI fabric interfaces, switches, creating VLAN Pool, policies, policy group and profiles, as well as configuring application-level elements such as ACI endpoint groups (EPGs) and bridge domains (BDs). This configuration should be done before creating the SmartFabric.

The networks used in the validated environment are shown in Table 3 on page 13, along with the corresponding bridge domain, and application EPG names used in APIC configuration.

Before creating a SmartFabric, steps need to be performed to configure ACI.

The following steps were performed in the Cisco APIC GUI to configure ACI for the environment shown in Section 4.1.

### 4.2.1 Create a VLAN Pool

1. Go to **Fabric > Access Policies > Pools > VLAN**.
2. From the VLAN screen, right-click on **VLAN** and select **Create VLAN Pool**.
3. In the **Name** field, enter **VLANPool1**.
4. From the **Allocation mode** option, select **Static**.

In this example, the static allocation mode is used because it is important that the VLAN ID is the same as used in ACI and on the MX platform. Dynamic allocation mode enables APIC to choose VLANs from the pool dynamically.

---

**Note:** Always use static mode when the VLAN pool is referenced from a static source, such as a static path binding for an EPG for use with servers.

---

**Create VLAN Pool**

Specify the Pool identity

Name: VLANPool1

Description: optional

Allocation Mode: Dynamic Allocation **Static Allocation**

Encap Blocks:

VLAN Range	Allocation Mode	Role
[1611-2000]	Inherit allocMode from par...	External or On the wire en...

Cancel Submit

Figure 5 Create VLAN Pool

5. From the **Encap Blocks** field, click the **Add(+)** icon.
6. In the **VLAN Range** fields, enter **1611** and **2000** as shown in Figure 6.
7. From the **Allocation Mode** field, click to select **Static Allocation**.
8. For the Role, select **External or On the wire encapsulations**.

**Create Ranges**

Specify the Encap Block Range

Type: VLAN

Range: VLAN [1611] - VLAN [2000]  
Integer Value Integer Value

Allocation Mode: Dynamic Allocation Inherit allocMode from parent **Static Allocation**

Role: **External or On the wire encapsulations** Internal

Cancel Submit

Figure 6 VLAN Range

9. Click **OK** and then **Submit**.

## 4.2.2 Create a Physical Domain

A physical domain acts as a link between the VLAN pool and the Access Entity Profile (AEP).

1. Go to **Fabric > Access Policies > Physical and External Domains > Physical Domains**.
2. Right-click on **Physical Domain** and select **Create Physical Domain**.
3. In the **Name** field, enter **physDomain1**.
4. From the **VLAN Pool** drop-down, select the **VLANPool1** option (created above in section 4.2.1).
5. Click **Submit**.

Figure 7 Create Physical Domain

### 4.2.3 Create an Attachable Access Entity Profile

To create an Attachable Access Entity Profile, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Global > Attachable Access Entity Profiles**.
2. Right-click on **Attachable Access Entity Profiles** and select **Create Attachable Access Entity Profile**.
3. In the **Name** field, enter **AEP1**.
4. In the **Domains** field, click **Add(+)** icon.
5. Select **physDomain1** (created above in step 4.2.2) and then click **Update**.
6. Click **Next** and then **Finish**.

Figure 8 Create Attachable Access Entity Profile

## 4.2.4 Create a Port Channel Policy

To create Port Channel Policy:

1. Go to **Fabric > Access Policies > Policies > Interface > Port Channel**.
2. Right-click on **Port Channel** and select **Create Port Channel Policy**.
3. In the **Name** field, enter **LACPPol1**.
4. From the **Mode** drop-down, select **LACP Active**.

---

**Note:** When LACP is enabled on the leaf switch, it must also be enabled on the connected devices

---

5. Keep default settings that are shown in the **Control** field.
6. Click **Submit**.

The screenshot shows a web-based configuration interface for creating a Port Channel Policy. The title is "Create Port Channel Policy" and the subtitle is "Specify the Port Channel Policy". The form includes the following fields:

- Name:** LACPPol1
- Description:** optional
- Alias:** (empty)
- Mode:** LACP Active (dropdown menu, with a note "Not Applicable for FC PC" below it)
- Control:** Suspend Individual Port, Graceful Convergence, Fast Select Hot Standby Ports (multi-select menu)

Figure 9 Create Port Channel Policy

## 4.2.5 Create a vPC Interface Policy Group

When interfaces are configured in vPC, interface policy group needs to be created. vPC policy group contains the port channel behavior definition and the identifier.

1. Go to **Fabric > Access Policies > Interfaces > Leaf Interfaces > Policy Groups > VPC Interface**.
2. Right-click on **VPC Interface** and select **Create VPC Interface Policy Group**.
3. In the **Name** field, enter **vPCPolGrp1**.
4. From the **Attached Entity Profile** drop-down, select **AEP1** (created above in step 4.2.3).
5. From the **Port Channel Policy** drop-down, select **LACPPol1** (created above in step 4.2.4).
6. Click **Submit**.

## Create VPC Interface Policy Group

Specify the Policy Group identity



Name:	<input type="text" value="vPCPolGrp1"/>
Description:	<input type="text" value="optional"/>
Link Level Policy:	<input type="text" value="select a value"/> ▾
CDP Policy:	<input type="text" value="select a value"/> ▾
MCP Policy:	<input type="text" value="select a value"/> ▾
CoPP Policy:	<input type="text" value="select a value"/> ▾
LLDP Policy:	<input type="text" value="select a value"/> ▾
STP Interface Policy:	<input type="text" value="select a value"/> ▾
L2 Interface Policy:	<input type="text" value="select a value"/> ▾
Port Security Policy:	<input type="text" value="select a value"/> ▾
Egress Data Plane Policing Policy:	<input type="text" value="select a value"/> ▾
Ingress Data Plane Policing Policy:	<input type="text" value="select a value"/> ▾
Priority Flow Control Policy:	<input type="text" value="select a value"/> ▾
Fibre Channel Interface Policy:	<input type="text" value="select a value"/> ▾
Slow Drain Policy:	<input type="text" value="select a value"/> ▾
MACsec Policy:	<input type="text" value="select a value"/> ▾
Attached Entity Profile:	<input type="text" value="AEP1"/> ▾ 
Port Channel Policy:	<input type="text" value="LACPPol1"/> ▾ 

Figure 10 Create VPC Interface Policy Group


### 4.2.6 Create a Leaf Access Port Policy Group


1. Go to **Fabric > Access Policies > Interfaces > Leaf Interfaces > Policy Groups > Leaf Access Port**.
2. Right-click on **Leaf Access Port** and select **Create Leaf Access Port Policy Group**.
3. In the **Name** field, enter **LeafHostPortGrp1**.
4. From the **Attached Entity Profile** drop-down, select **AEP1** (created above in step 4.2.3).
5. Click **Submit**.


## Create Leaf Access Port Policy Group


Name:


Description:

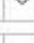
Link Level Policy:  

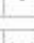
CDP Policy:  

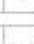
MCP Policy:  

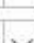
CoPP Policy:  


LLDP Policy:  


STP Interface Policy:  

Storm Control Interface Policy:  


L2 Interface Policy:  


Port Security Policy:  


Egress Data Plane Policing Policy:  


Ingress Data Plane Policing Policy:  


Monitoring Policy:  


Priority Flow Control Policy:  


Fibre Channel Interface Policy:  

PoE Interface Policy:  

Slow Drain Policy:  

MACsec Policy:  

802.1x Port Authentication Policy:  

DWDM Policy:  



Attached Entity Profile:   

Figure 11 Create Leaf Access Port Policy Group

### 4.2.7 Create a Leaf Interface Profile

Once the vPC Interface Policy Group and Leaf Access Port Policy Group is created to bundle the interfaces, the interfaces need to be added to the policy groups. To achieve that, leaf interface profile is created, and access port selectors connect the interfaces to the policy groups.

1. Go to **Fabric > Access Policies > Interfaces > Leaf Interfaces > Profiles.**

2. Right-click on **Profiles** and select **Create Leaf Interface Profile**.
3. In the **Name** field, enter **LeafIntProf1**.
4. From the **Interface Selectors** field, click the **Add(+)** icon.

**Create Leaf Interface Profile**

Specify the profile identity

Name: LeafIntProf1

Description: optional

Figure 12 Create Leaf Interface Profile

5. Create Access Port Selectors:
  - a. In the **Name** field, enter **LeafHostSel1**.
  - b. From the **Interface IDs**, enter **1/1-3**. These ports are connected directly to the R730xd servers.
  - c. From the **Interface Policy Group** drop-down, select **LeafHostPortGrp1** (created above in step 4.2.6).
  - d. Click **OK**.
  - e. From the **Interface Selectors listing**, click the **Add(+)** icon.

**Create Access Port Selector**

Specify the selector identity

Name: LeafHostSel1

Description: optional

Interface IDs: 1/1-3  
valid values: All or Ranges. For Example:  
 1/13, 1/15 or 2/22-2/24, 2/16-3/16, or  
 1/21-23/1-4, 1/24/1-2

Connected To Fex:

Interface Policy Group: LeafHostPortGrp1

Figure 13 Access Port Selector for host interfaces

- f. **LeafvPCSel1** contains vPC interfaces **1/51-52**. The ports on the Nexus leaf switches are vPC ports, connected to MX9116n FSEs. Associate it to **vPCPolGrp1** (created above in step 4.2.5) and click **OK**.

## Create Access Port Selector

Specify the selector identity

Name:

Description:

Interface IDs:   
valid values: All or Ranges. For Example:  
 1/13, 1/15 or 2/22-2/24, 2/16-3/16, or  
 1/21-23/1-4, 1/24/1-2

Connected To Fex:

Interface Policy Group:  

Figure 14 Access Port Selector for vPC interfaces

g. Click **Submit**.

### 4.2.8 Create a VPC Domain Policy

To create VPC Domain Policy, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Switch > VPC Domain**.
2. Right-click on **VPC Domain** and select **Create VPC Domain Policy**.
3. In the **Name** field, enter **vPCDom1**.
4. Click **Submit**.

## Create VPC Domain Policy

Specify the Domain Policy Identity

Name:

Description:

Peer Dead Interval:  

Figure 15 Create vPC Domain Policy

### 4.2.9 Create a VPC Explicit Protection Group

1. Click **Fabric > Access Policies > Policies > Switch** and select **Virtual Port Channel default**.
2. Leave Pairing Type set to **Explicit (default)**.
3. Next to Explicit **VPC Protection Groups**, click the **Add(+)** icon.
4. In the **Name** field, enter **vPCExpProGrp1**.
5. In the **ID** field, enter **101**.
6. From the **VPC Domain Policy** drop-down, select **vPCDom1** (created above in step 4.2.8).
7. For **Switch 1**, select the first leaf switch, **101/Leaf1**.



8. For **Switch 2**, select the second leaf switch, **102/Leaf2**.
9. Click **Submit**.

Create VPC Explicit Protection Group

Specify the Explicit Group settings

Name: vPCExpProGrp1

ID: 101

VPC Domain Policy: vPCDom1

Switch 1: 101

Switch 2: 102

Figure 16 Create vPC Explicit Protection Group

#### 4.2.10 Create a Leaf Profile

1. Go to **Fabric > Access Policies > Switches > Leaf Switches > Profiles**.
2. Right-click on **Profiles** and select **Create Leaf Profile**.
3. In the **Name** field, enter **LeafProf1**.
4. Next to **Leaf Selectors**, click the **Add(+)** to create a Leaf Selector:
  - a. In the **Name** field, enter **LeafSel1**.
  - b. **Blocks** - select switches **101** and **102** and click **Update**.

## Create Leaf Profile

STEP 1 > Profile

1. Profile 2. Associations

Specify the profile Identity

Name: LeafProf1

Description: optional

Leaf Selectors:

Name	Blocks	Policy Group
LeafSel1	101,102	

Previous Cancel Next

Figure 17 Create Leaf Profile

- c. Click **Next**.
- d. From the **Interface Selector Profiles**, select **LeafIntProf1** (created above in step 4.2.7), then click **Finish**.
- e. Leaf 101 and 102 display in the **Leaf Profile** shown in Figure 17.

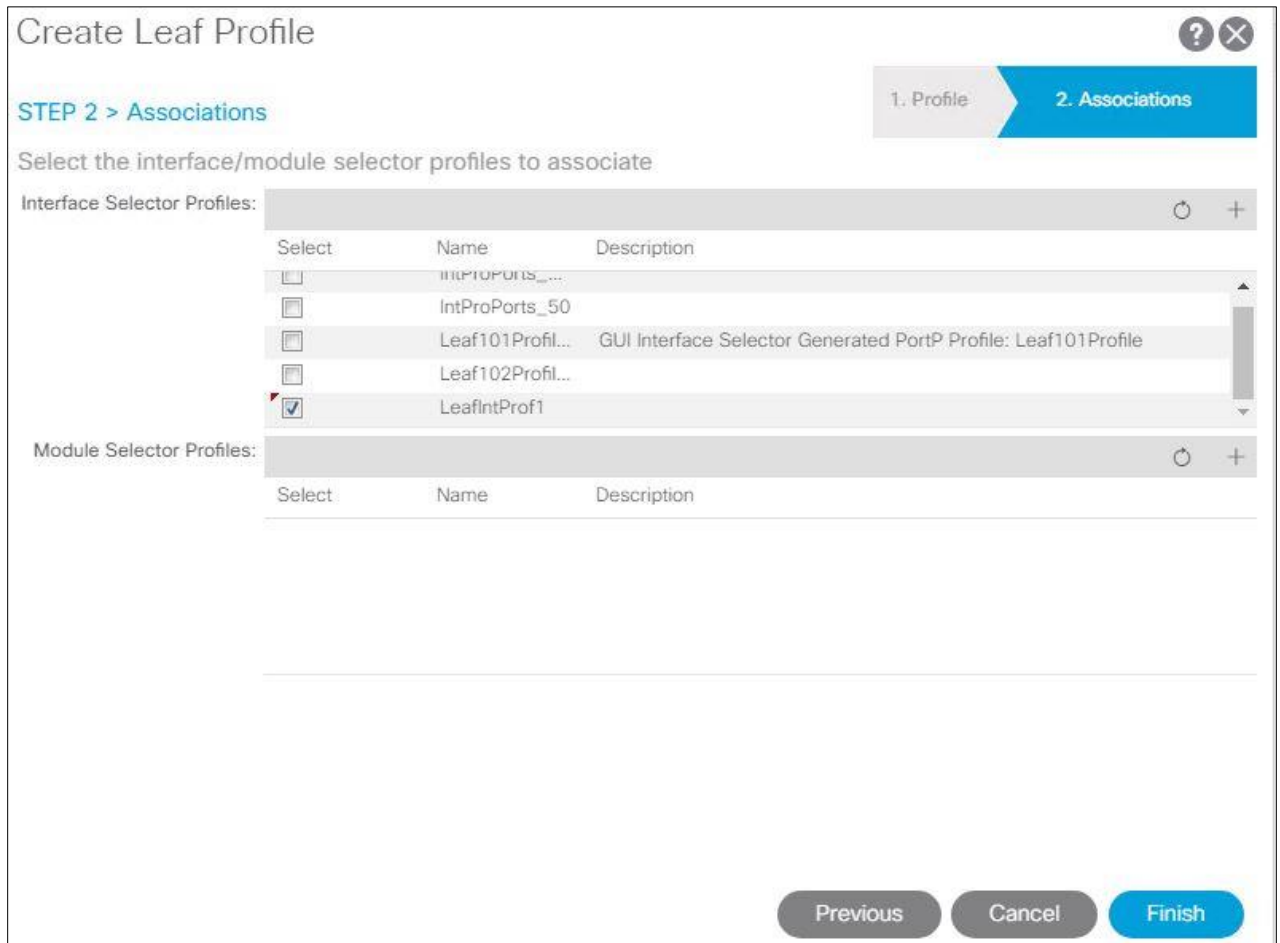


Figure 18 Choose Interface selector profile

## 4.2.11 Create a Tenant

To create a Tenant:

1. Go to **Tenants > Add Tenant**.
2. In the **Name** field, enter **Customer-TN1**.
3. Click **Submit**.

## Create Tenant

Specify tenant details

Name:

Alias:

Description:

Tags:    
enter tags separated by comma

GUID:

Provider	GUID

Monitoring Policy:

Security Domains:

Name	Description

Figure 19 Create a Tenant

### 4.2.12 Create a VRF

Virtual Routing and Forwarding (VRF) also called private networks are a unique layer 3 forwarding and application policy domain. Private networks contain Bridge domains.

1. Go to **Tenants > Customer-TN1 > Networking > VRFs**.
2. Right-click on **VRFs** and select **Create VRF**.
3. In the **Name** field, enter **VRF1**.
4. Click to deselect the **Create a Bridge Domain** option and then click **Finish**.

## Create VRF

**STEP 1 > VRF**

Specify Tenant VRF

Name:

Alias:

Description:

Tags:    
enter tags separated by comma

Policy Control Enforcement Preference:  Enforced  Unenforced

Policy Control Enforcement Direction:  Egress  Ingress

BD Enforcement Status:

Endpoint Retention Policy:    
This policy only applies to remote L3 entries

Monitoring Policy:

DNS Labels:   
enter names separated by comma

Route Tag Policy:

IP Data-plane Learning:  Disabled  Enabled

Create A Bridge Domain:

Figure 20 Create VRF

### 4.2.13 Create Bridge Domains

Layer 2 forwarding domain within the fabric is a Bridge Domain. Bridge domain is linked to a private network and it can have multiple subnets.

---

**Note:** Refer to Table 3 as needed to complete the following steps.

---

Bridge domains are created for each VLAN as follows:

1. Click **Tenants > Customer-TN1 > Networking > Bridge Domains**.
2. Right-click on **Bridge Domains** and then select **Create Bridge Domain**.
3. In the field provided, enter the name of the first bridge domain, **webBD1**.
4. From the VRF drop-down, select **VRF1** (created above in step 4.2.12), and click **Next**.

# Create Bridge Domain

STEP 1 > Main 1. Main

Specify Bridge Domain for the VRF

Name:

Alias:

Description:

Tags:    
enter tags separated by comma

Type:

Advertise Host Routes:

VRF:

Forwarding:

Endpoint Retention Policy:    
This policy only applies to local L2 L3 and remote L3 entries

IGMP Snoop Policy:

Figure 21 Create Bridge Domain

5. Next to the **Subnets** listing, click the **Add(+)** icon.
6. In the **Gateway IP** field, enter 172.16.14.254/24 for the address and mask for the bridge domain. Leave the remaining values at their defaults settings.

**Create Subnet** ? ✕

Specify the Subnet Identity

Gateway IP:   
address/mask

Treat as virtual IP address:

Make this IP address primary:

Scope:  Private to VRF  
 Advertised Externally  
 Shared between VRFs

Description:

Subnet Control:  No Default SVI Gateway  
 Querier IP

L3 Out for Route Profile:

Route Profile:

ND RA Prefix policy:

Figure 22 Create Subnet

7. Click **OK, Next** and then click **Finish**.
8. Repeat the steps in this section as needed for each VLAN. Note that the additional bridge domains created in this example are appBD1, dbBD1, ESXiMgmtBD1, vMotionBD1, and vSANBD1.

#### 4.2.14 Create an Application Profile

1. Go to **Tenants > Customer-TN1 > Application Profiles**.
2. Right-click on **Application Profiles** and select **Create Application Profile**.
3. In the **Name** field, enter **ap1**.
4. Click **Submit**.

## Create Application Profile

Specify Tenant Application Profile

Name:

Alias:

Description:

Tags:    
enter tags separated by comma

Monitoring Policy:

### EPGs

Name	Alias	BD	Domain	Switching Mode	Static Path

Figure 23 Create Application Profile

### 4.2.15 Create Application EPGs

End point groups (EPGs) are logically grouped hosts or servers that share similar policies and perform similar functions within the fabric.

---

**Note:** Refer to Table 3 for the required network information.

---

1. Click **Tenants > Customer-TN1 > Application Profiles > ap1 > Application EPGs**.
2. Right-click on **Application EPGs** and then select **Create Application EPG**.
3. In the **Name** field, enter **webEPG1** as the name of the first EPG.
4. From the **Bridge Domain** drop-down, select **webBD1**.
5. Click **Finish**.



## Create Application EPG

**STEP 1 > Identity**

Specify the EPG Identity

Name:

Alias:

Description:

Tags:    
enter tags separated by comma

Contract Exception Tag:

QoS class:

Custom QoS:

Data-Plane Policer:

Intra EPG Isolation:

Preferred Group Member:

Flood on Encapsulation:

Bridge Domain:

Figure 24 Create Application EPG

6. Create a separate EPG for each of the remaining bridge domains using the EPG names provided in Table 3: appEPG1, dbEPG1, ESXiMgmtEPG1, vMotionEPG1, and vSANEPG1.

### 4.2.16 Configure the Access Entity Profile with EPGs and VLANs

**Note:** Refer to Table 3 for the necessary information.

1. Go to **Fabric > Access policies > Policies > Global > Attachable Access Entity Profiles**.
2. From the profiles listed, select **AEP1** (created above in step 4.2.3).

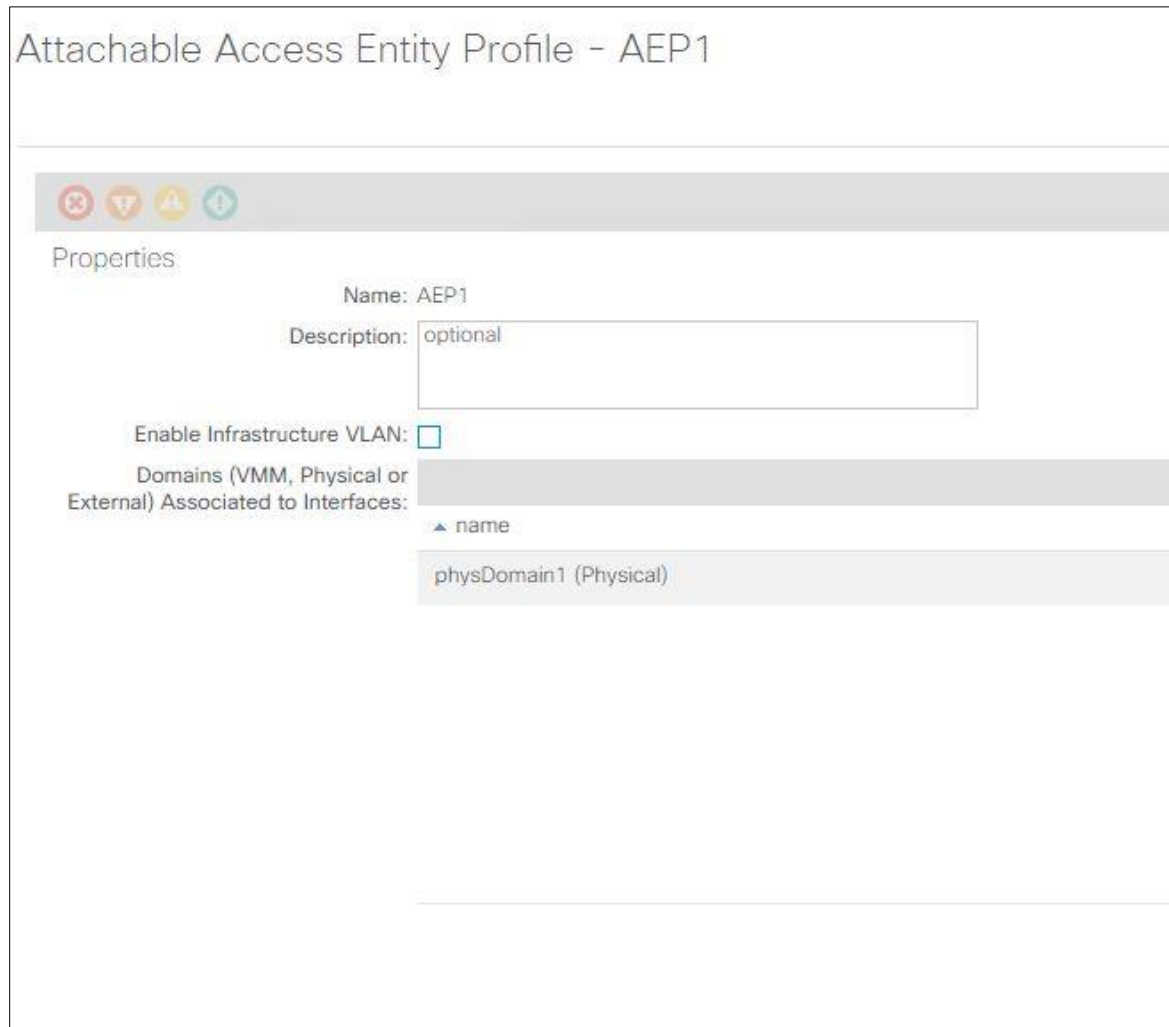


Figure 25 Create Attachable Access Entity Profile

3. At bottom of page next to **Application EPGs**, click the **Add(+)** icon.
4. For the first EPG, webEPG1, select the following options:
  - a. From the **Tenant** drop-down, select **Customer-TN1**.
  - b. From the **Application Profile** menu, select **ap1**.
  - c. From the **EPG** menu, select **webEPG1**.
  - d. In the **Encap** field, enter **vlan-1614**.
  - e. Leave the **Primary Encap** field blank.
  - f. From the **Mode** menu, select **Trunk**.
  - g. Click **Update**.

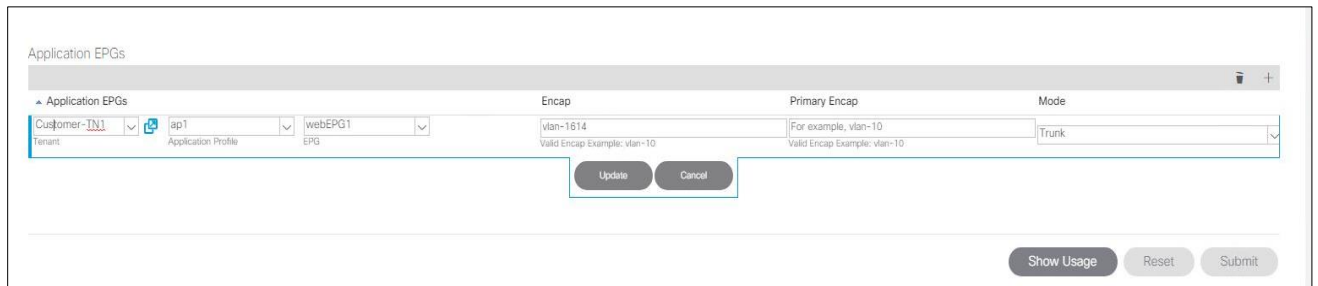


Figure 26 Attach AEP to EPGs and Bridge Domains

- Repeat the steps in this section for all remaining EPGs using their associated VLAN IDs.

## 4.2.17 Create vCenter domain for Cisco ACI and Virtual Machine Manager (VMM) Domain Integration

By creating vCenter domain, user can connect the VMs by creating and configuring policies and EPGs in the Cisco APIC. These EPGs as well as policies in turn are pushed to vCenter as port groups.

---

**Note:** The name of the Datacenter created in APIC under vCenter domain must be same as the Datacenter name on vCenter mentioned in section 4.5.

---

To create VMware vCenter domain:

- Click **Virtual Networking > VMM Domains**.
- Right-click on **VMware** and choose **Create vCenter Domain**.
- In **Virtual Switch Name** field, enter **VDS-ACI**.
- From the **Virtual Switch**, select **VMware vSphere Distributed Switch**.
- From **Associated Attachable Entity Profile** menu, select **AEP-1**.
- Select **VLAN Pool**. In this example **VLANPool1** is selected. A new VLAN pool can also be created and attached.

## Create vCenter Domain

Specify vCenter domain users and controllers

Virtual Switch Name:

Virtual Switch: **VMware vSphere Distributed Switch** | Cisco AVS | Cisco AVE

Associated Attachable Entity Profile:

Delimiter:

Enable Tag Collection:

Access Mode:

Endpoint Retention Time (seconds):

VLAN Pool:

Security Domains:

Name	Description

vCenter Credentials:

Profile Name	Username	Description
vCenter-Credentials	administrator@dell.l...	

Figure 27 Create vCenter Domain

7. From the **vCenter Credentials** listing, click the **Add(+)** icon.
  - a. In the **Name** field, enter **vCenter-Credentials**.
  - b. In the **Username** field, enter `administrator@dell.local`.
  - c. In the fields provided, enter and confirm **Password**, then click **OK**.

Create vCenter Credential

Specify account profile

Name: vCenter-Credentials

Description: optional

Username: administrator@dell.local

Password: .....

Confirm Password: .....

Cancel OK

Figure 28 vCenter Credential

8. Next to the **vCenter** listing, click the **Add(+)** to add the **vCenter Controller**.
  - a. In the **Name** field, enter **vCenter**.
  - b. Enter **Host Name or IP Address** as per the configuration.
  - c. In the **Datacenter** field, enter **MgmtDatacenter**.
  - d. Associate **vCenter-Credentials** created above and click **Submit**.

---

**Note:** The **Management EPG** field is optional. New Management EPG can also be created and associated by choosing **Create EPG under Tenant mgmt** from this menu.

---

## Create vCenter Controller

Specify controller profile

Name:

Host Name (or IP Address):

DVS Version:  ▾

Datacenter:

Stats Collection:

Management EPG:  ▾


Associated Credential:  ▾ 

Figure 29 Create vCenter Controller

9. Select the **Port Channel Mode**, **vSwitch Policy** and **NetFlow Exporter Policy** as per configuration. For this example, these options are not required.
10. Click **Submit**.

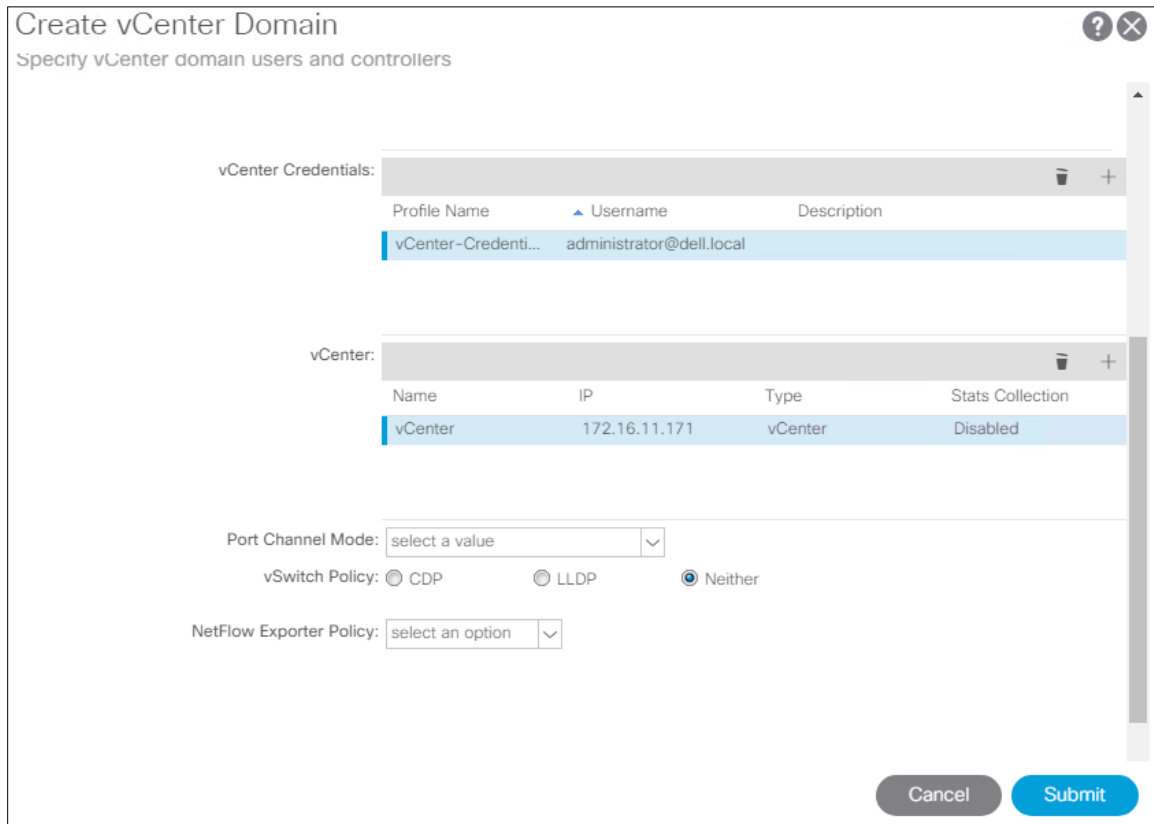


Figure 30 vCenter domain after adding vCenter

## 4.2.18 Create a Contract Filter

Contracts are necessary in order to communicate between EPGs.

1. Go to **Tenants > Customer-TN1 > Contracts > Filters**.
2. Right-click on **Filters** and select **Create Filter**.
3. In the **Name** field, enter **AllowAllFilter1**.
4. In the **Entries** section, click the **Add(+)** icon:
  - a. In the **Name** field, enter **Allow**.
  - b. Select the **IP** as **EtherType**.
  - c. Leave remaining items at their defaults and click **Update** and then **Submit**.

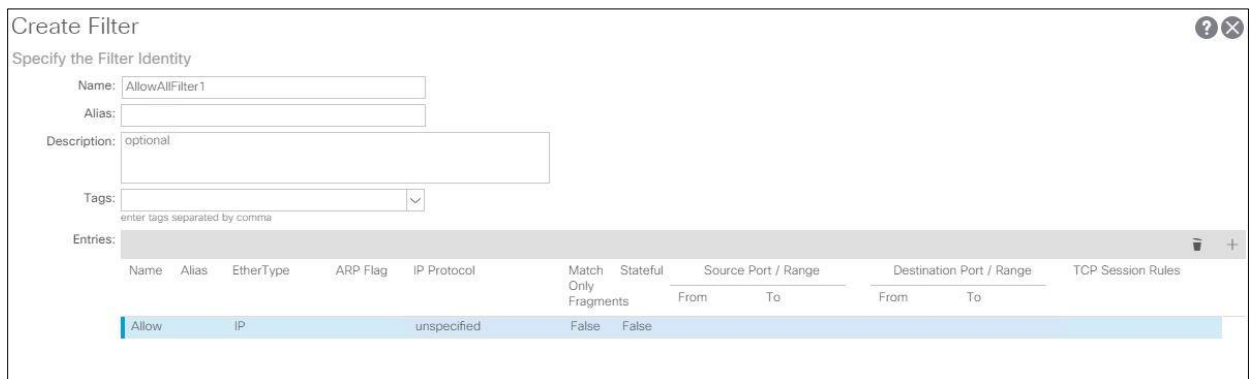


Figure 31 Create contract Filter

## 4.2.19 Create a Contract

Contract provides a way to control traffic flow within the ACI fabric between EPGs. To create Contract, perform the following steps:

1. Go to **Tenants > Customer-TN1 > Contracts > Standard**.
2. Right-click **Standard** and select **Create Contract**.
3. In the **Name** field, enter **AllowAllContract1**.

**Create Contract** [?] [X]

Specify Identity Of Contract

Name: AllowAllContract1

Alias:

Scope: VRF

QoS Class: Unspecified

Target DSCP: Unspecified

Description: optional

Tags:

enter tags separated by comma

Subjects:

Name	Description
------	-------------

Cancel Submit

Figure 32 Create Contract

4. In the **Subjects** field, click the **Add(+)** icon.
5. In the **Name** field, enter **AllowAllSub1**.
6. In the **Filters** field, click the **Add(+)** icon.
7. Under filter Name, select **AllowAllFilter1** (created above in step 4.2.18).



**Create Contract Subject**

Name: AllowAllSub1

Alias:

Description: optional

Target DSCP: Unspecified

Apply Both Directions:

Reverse Filter Ports:

**Filter Chain**

L4-L7 Service Graph: select an option

QoS Priority:

Filters			
Name	Directives	Action	Priority
Customer-TN1/AllowAllFilter1	none	permit	default

Cancel OK

Figure 33 Create Subject

8. Click **Update** > **OK** > **Submit**.

#### 4.2.20 Apply the contract to the VRF

1. Go to **Tenant** > **Customer-TN1** > **Networking** > **VRFs** > **VRF1**.
2. Expand the **VRF1** section and select EPG collection for VRF.
3. Next to **Provided Contracts** listing, click the **Add(+)** icon:
  - a. In the **Name** field, select **AllowAllContract1** (created above in step 4.2.19).
  - b. Click **Update**.
4. Next to **Consumed Contracts** listing, click the **Add(+)** icon:
  - a. In the **Name** field, select **AllowAllContract1** (created above in step 4.2.19).
  - b. Click **Update**.

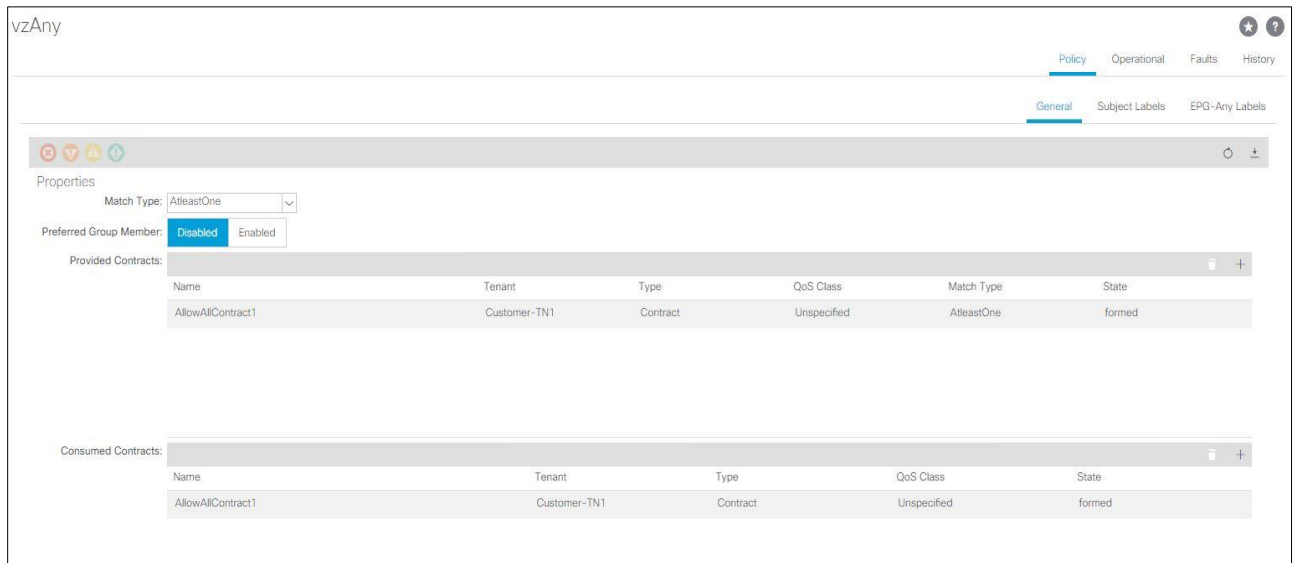


Figure 34 Apply the Contract to VRF

In this deployment, EPGs are extended outside of the ACI fabric by mapping EPGs to external VLANs. This is so when a frame tagged with, VLAN 1611 for example, enters the ACI fabric, ACI knows that it belongs to the ESXi Management EPG and treats it accordingly.

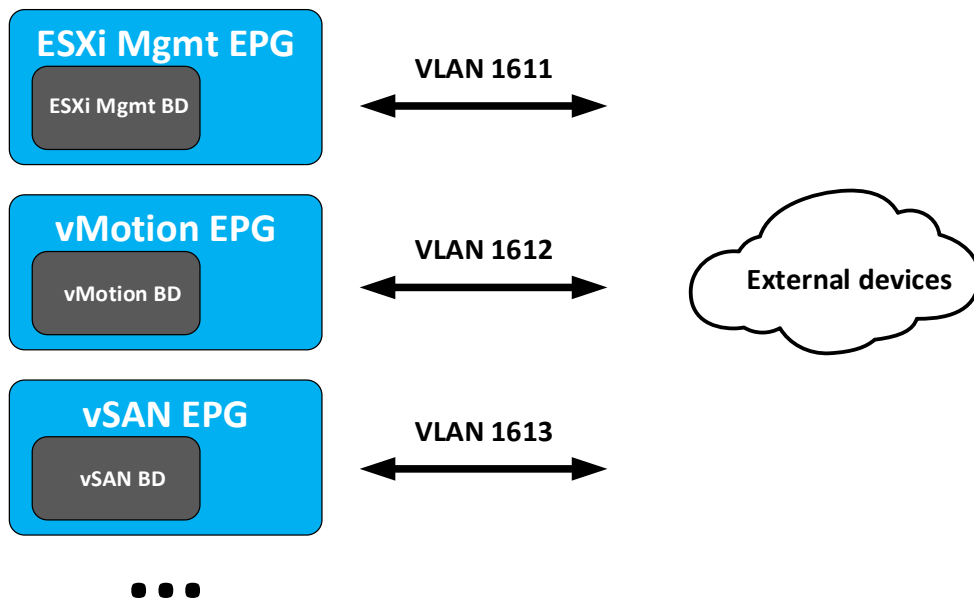


Figure 35 Bridge domains are associated with EPGs, which are mapped to external VLANs.

## 4.3 Deploy the SmartFabric

This section provides the details used to deploy the SmartFabric that is used in the example provided in this guide. Download the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#), which is referenced in this section.

### 4.3.1 Define VLANs

The VLAN settings used during the SmartFabric deployment for this environment, are shown in Table 6.

Table 6 SmartFabric VLAN settings

VLAN ID	VLAN name	Description	Network type (QoS)	Tagged/Untagged
1611	ESXi_Mgmt	ESXi host in-band management	Hypervisor Management	Tagged
1612	vMotion	VM migration	VM migration	Tagged
1613	vSAN	Storage	Storage – Data Replication	Tagged
1614	web	VM data network	General Purpose (Silver)	Tagged
1615	app	VM data network	General Purpose (Silver)	Tagged
1616	db	VM data network	General Purpose (Silver)	Tagged

**Note:** For instructions on Defining VLANs for the SmartFabric on OME-M console, see *Section 4.2 - Define VLANs* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

**Note:** For information about network type and QoS group settings, see *Section 2.7 - Network and Automated QoS* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

The configured VLANs for this example are shown in Figure 36.

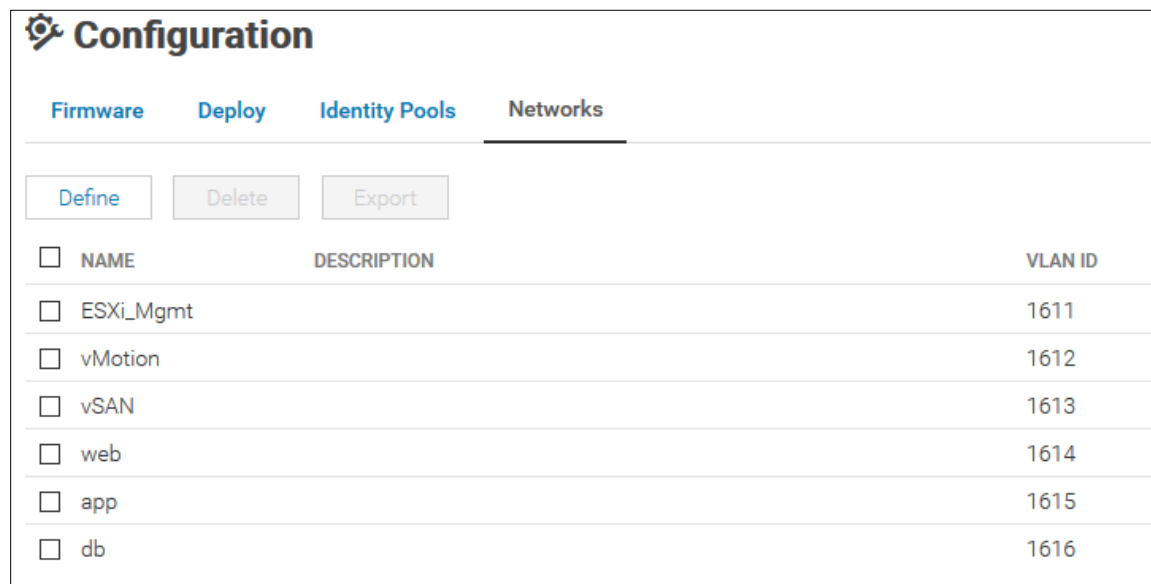


Figure 36 Defined VLANs

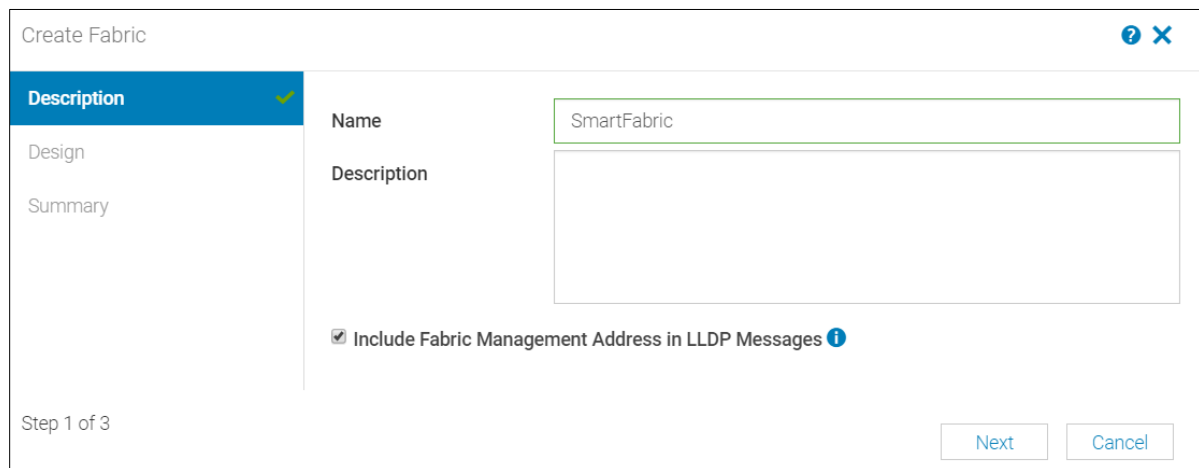
## 4.3.2 LLDP setting for SmartFabric

Cisco ACI uses Link Layer Discovery Protocol (LLDP) to discover and build the network topology that includes the Distributed Virtual Switch (DVS) hosted in the hypervisor. To enable this functionality, click the checkbox next to **Include Fabric Management Address in LLDP Messages** on the **Create Fabric** screen, as shown in Figure 37, during deployment.

---

**Note:** Without the **Include Fabric Management Address in LLDP Messages** feature enabled, the ACI fabric will not be able to discover the complete network topology.

---



The screenshot shows a 'Create Fabric' dialog box. On the left, there is a sidebar with tabs: 'Description' (selected and highlighted in blue with a checkmark), 'Design', and 'Summary'. The main area contains a 'Name' text box with 'SmartFabric' entered, a larger 'Description' text box, and a checkbox labeled 'Include Fabric Management Address in LLDP Messages' which is checked. At the bottom left, it says 'Step 1 of 3'. At the bottom right, there are 'Next' and 'Cancel' buttons. A help icon (?) and a close icon (X) are in the top right corner.

Figure 37 Enabling LLDP in SmartFabric

After creating the SmartFabric in section 4.3.3 and creating the uplink in section 4.3.4, the VMs display in the APIC under the **Tenants** tab after configuring vCenter. Select the **Tenant** and click **Networking** to view the network topology.

---

**Note:** If VMs are not present in APIC after creating the SmartFabric with this feature enabled, bring down the MX9116n downlink ports going to the VMs and then bring them back up.

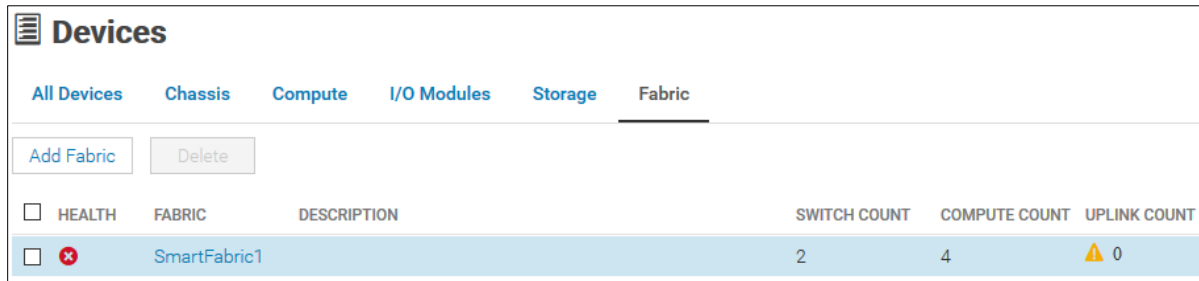
---

## 4.3.3 Create the SmartFabric

To create a SmartFabric using the OME-M console, perform the following steps in *Section 4.3 - Create the SmartFabric* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

The SmartFabric deployment takes several minutes to complete. During this time, the related IOMs reload, the operating mode of the IOMs change to SmartFabric, and the SmartFabric is created.

Figure 38 shows the new SmartFabric object.



The screenshot shows the 'Devices' page with the 'Fabric' tab selected. Below the navigation tabs are 'Add Fabric' and 'Delete' buttons. A table lists the fabric objects with columns for Health, Fabric, Description, Switch Count, Compute Count, and Uplink Count. The 'SmartFabric1' entry is highlighted and shows a red 'x' in the Health column, a value of 2 for Switch Count, 4 for Compute Count, and 0 for Uplink Count with a yellow warning triangle icon.

HEALTH	FABRIC	DESCRIPTION	SWITCH COUNT	COMPUTE COUNT	UPLINK COUNT
<input type="checkbox"/>	SmartFabric1		2	4	0

Figure 38 SmartFabric after deployment before uplinks are created

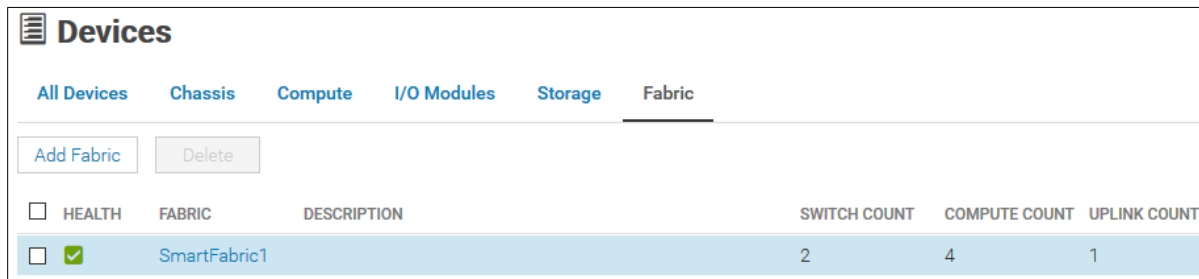
After creation, the SmartFabric shows the **Uplink Count** as zero with the icon displayed. The **Health** column displays the icon until uplinks are defined.

### 4.3.4 Create the Uplink

**Note:** To change the port speed or breakout configuration, see *Section 4.4 - Configure uplink port speed or breakout* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#) and make those changes before creating the uplinks. No port breakout was used in this example.

To create an ethernet uplink from the MX9116n FSEs to the Cisco ACI leafs, see *Section 4.5 - Create Ethernet uplink* in the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

After creating uplinks, the SmartFabric creates the uplink object. If the connected Cisco ACI vPC is configured correctly, the uplink comes up and the status for the fabric changes to **Ok** on the **Devices > Fabric** page as shown in Figure 39.



The screenshot shows the 'Devices' page with the 'Fabric' tab selected. Below the navigation tabs are 'Add Fabric' and 'Delete' buttons. A table lists the fabric objects with columns for Health, Fabric, Description, Switch Count, Compute Count, and Uplink Count. The 'SmartFabric1' entry is highlighted and shows a green checkmark in the Health column, a value of 2 for Switch Count, 4 for Compute Count, and 1 for Uplink Count.

HEALTH	FABRIC	DESCRIPTION	SWITCH COUNT	COMPUTE COUNT	UPLINK COUNT
<input type="checkbox"/>	SmartFabric1		2	4	1

Figure 39 SmartFabric status after uplink is created

## 4.4 Deploy servers

### 4.4.1 Create Server Templates

Create a server template for each unique server and NIC combination used in the chassis group. For identical servers, only create one template.

---

**Note:** For the hardware used in this example, three templates were created:

---

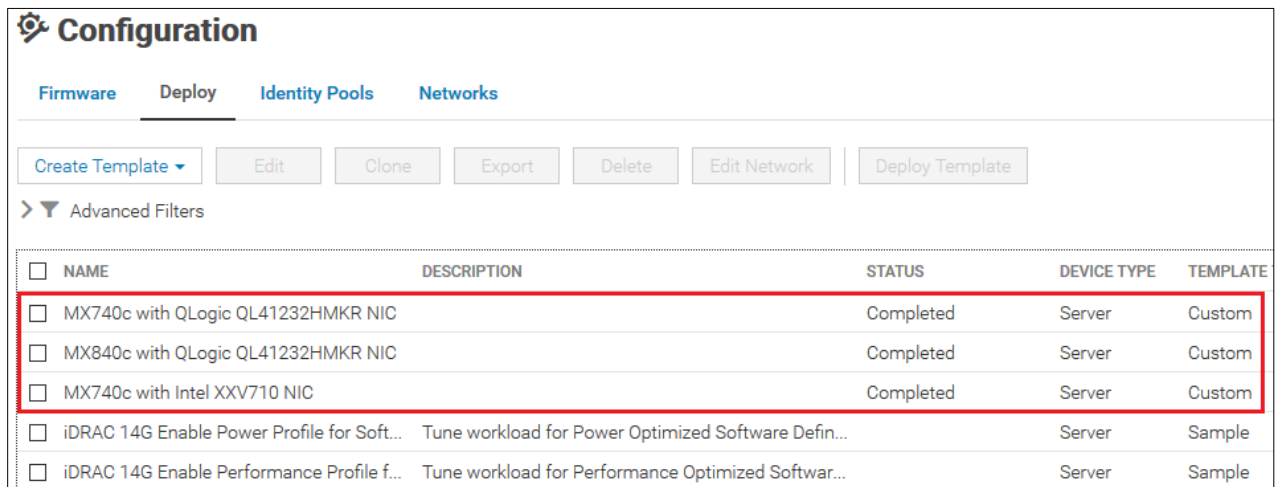
- MX740c with QLogic QL41232HMKR NIC
- MX740c with Intel XXV710 NIC
- MX840c with QLogic QL41232HMKR NIC

---

**Note:** To create a server template, follow the steps in *Section 5.2 - Create a server template* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

---

The templates created for this example are shown in Figure 40.



<input type="checkbox"/>	NAME	DESCRIPTION	STATUS	DEVICE TYPE	TEMPLATE
<input type="checkbox"/>	MX740c with QLogic QL41232HMKR NIC		Completed	Server	Custom
<input type="checkbox"/>	MX840c with QLogic QL41232HMKR NIC		Completed	Server	Custom
<input type="checkbox"/>	MX740c with Intel XXV710 NIC		Completed	Server	Custom
<input type="checkbox"/>	iDRAC 14G Enable Power Profile for Soft...	Tune workload for Power Optimized Software Defini...		Server	Sample
<input type="checkbox"/>	iDRAC 14G Enable Performance Profile f...	Tune workload for Performance Optimized Softwar...		Server	Sample

Figure 40 Server templates created

### 4.4.2 Add VLANs to the server templates

After successfully creating server templates, associate each template with appropriate VLANs. See *Section 5.4 - Associate server template with networks* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#) for the steps necessary.

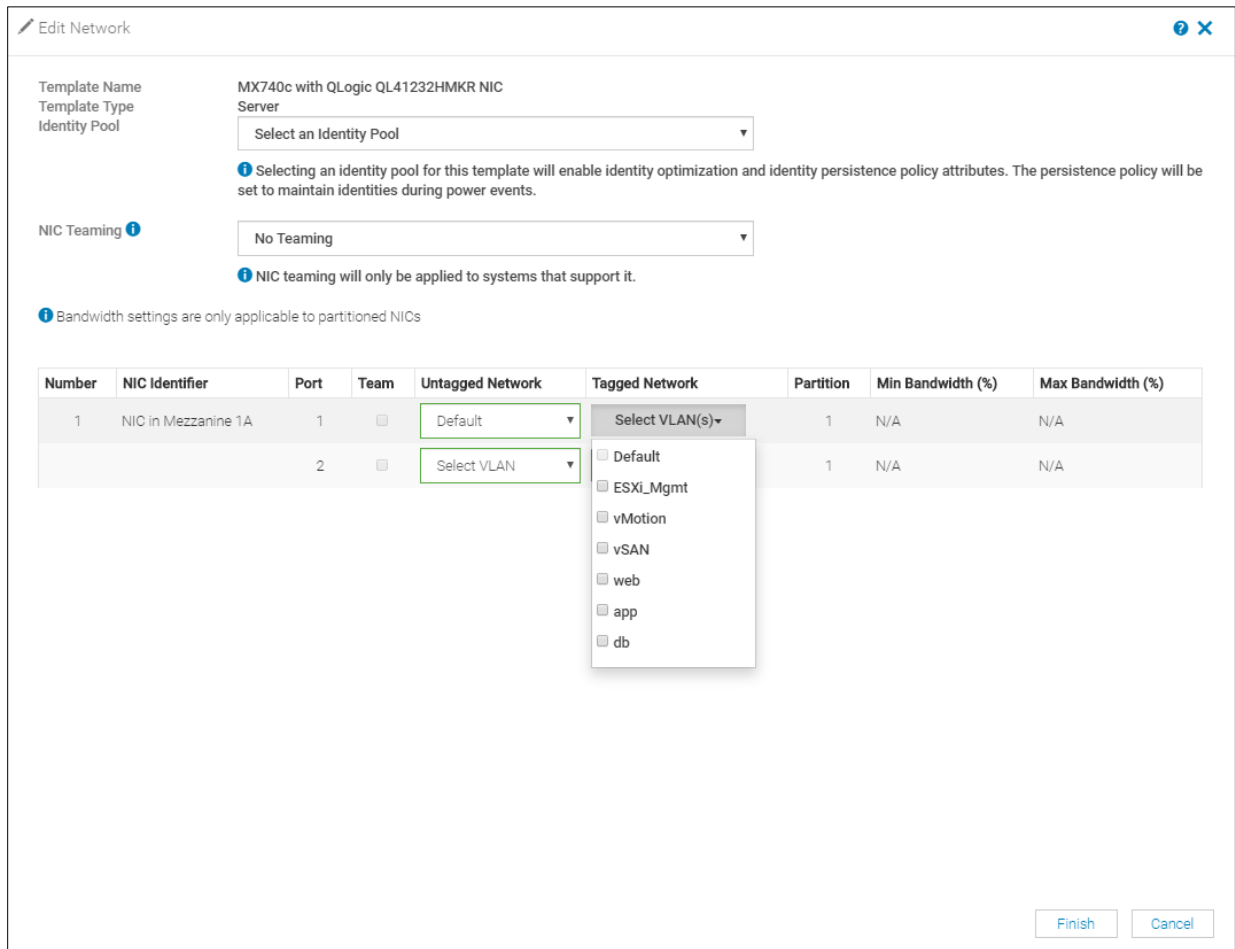


Figure 41 VLANs added to server template

### 4.4.3 Deploy the Server Templates

To deploy the server templates, complete the steps in *Section 5.6 - Deploy a server template* of the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#).

## 4.5 vCenter configuration overview

The existing ACI environment has two PowerEdge R730xd rack servers connected to the ACI leaves. The rack servers are in a vSphere cluster named **Management**.

After the SmartFabric is deployed and uplink is created, the rack servers can be added to vCenter. To create a data center, create a cluster, add a host, create virtual machine, configure a cluster and create VDS, see [Documentation related to configure vCenter](#).

The MX compute sleds can now communicate with the rack servers and the vCenter, mgmtvc01. The MX compute sleds are joined to the vSphere cluster by an administrator as shown in Figure 42.

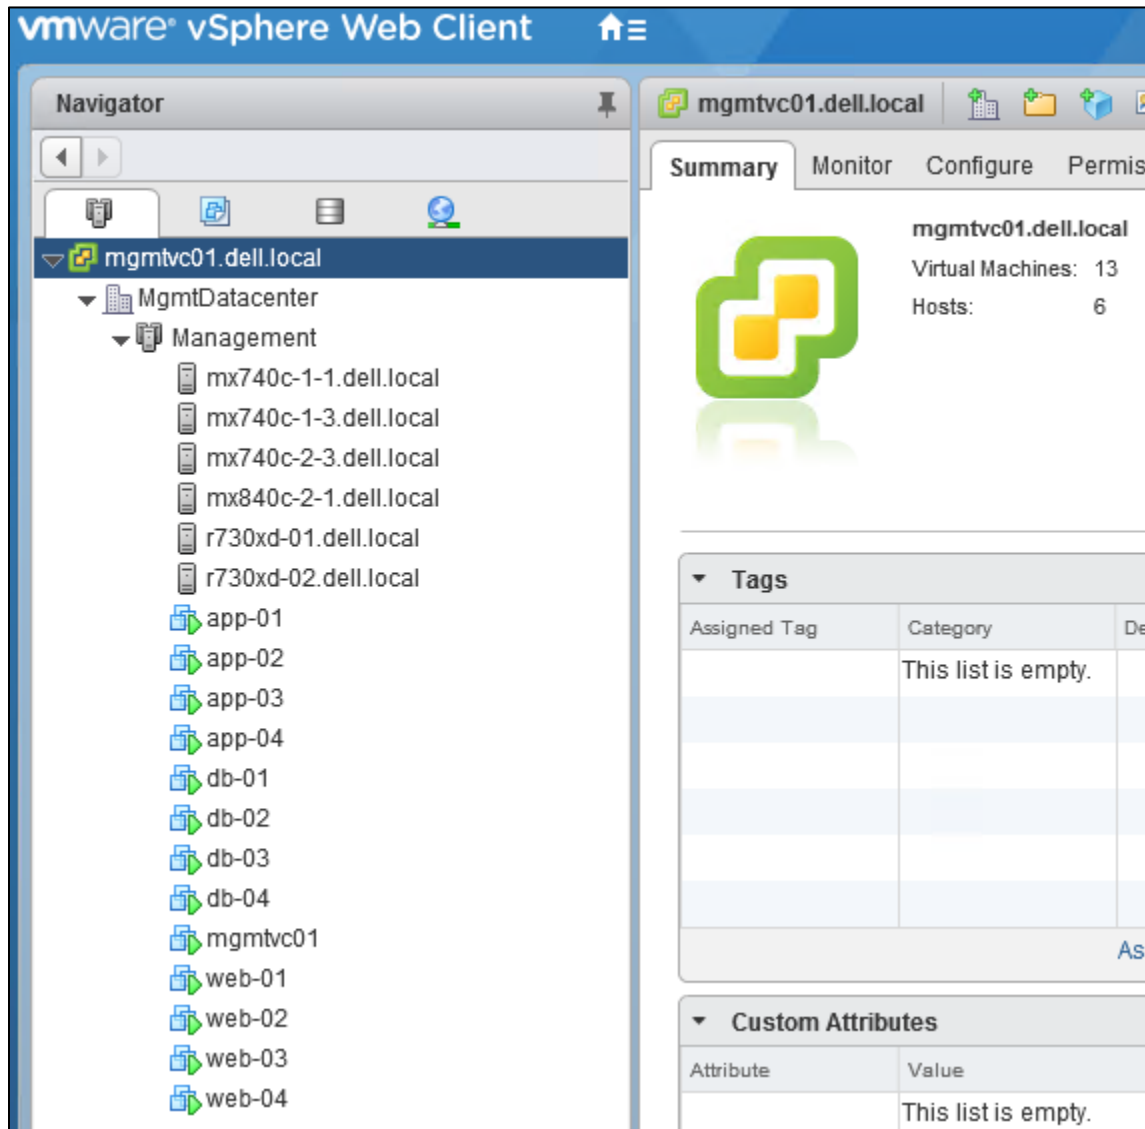


Figure 42 Hosts and VMs used in the validated environment in a single vSphere cluster

For information on creating VDS and configuring networking for VDS, see [Setting up Networking with vSphere Distributed Switches](#).



A VDS named **VDS-Mgmt**, along with six distributed port groups, one for each VLAN, are used as shown in Figure 43.

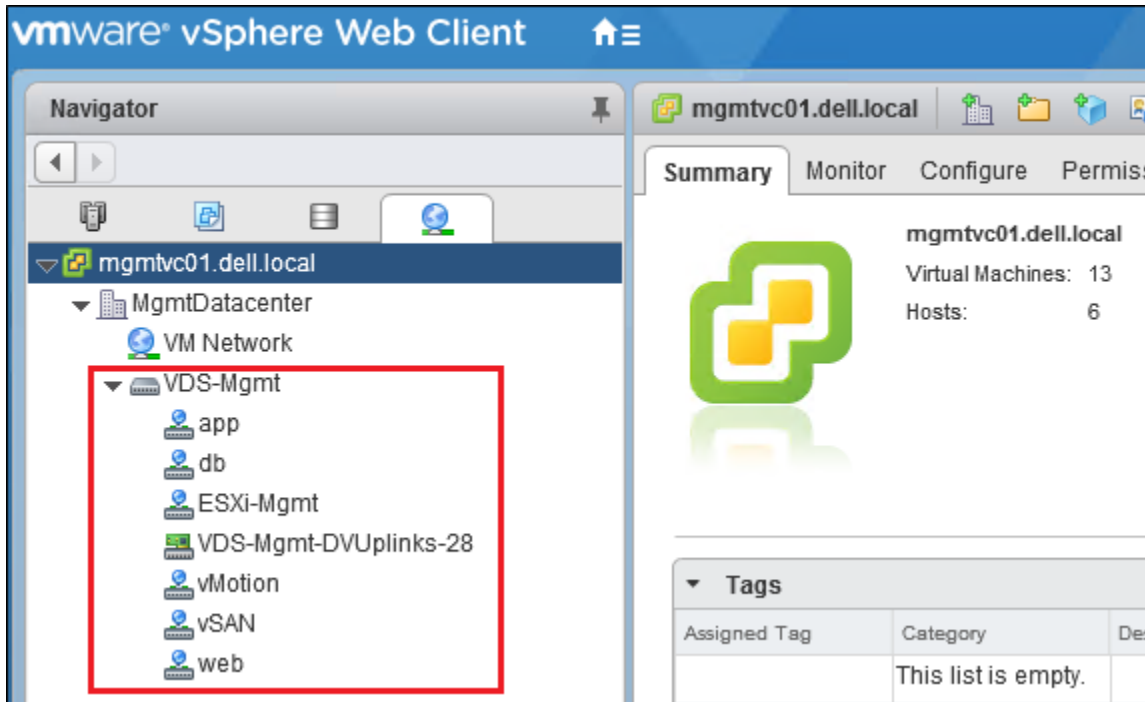


Figure 43 VDS and port groups used in the validated environment

**Note:** For each port group in the VDS in this example, both uplinks are active and the load balancing method used is Route based on physical NIC load as recommended in [VMware Validated Design Documentation](#). Detailed vCenter configuration is beyond the scope of this document.

For more information on vCenter configuration, see the [VMware vSphere Documentation](#).

## 4.6 SmartFabric connected with MX5108n Ethernet switch and Cisco ACI Leaf switches

A single MX7000 chassis may also join an existing Cisco ACI environment by using the MX5108n ethernet switch. The MX chassis in this example has two MX5108n ethernet switches and two MX compute sleds.

The connections between the ACI environment and the MX chassis are made using a double-sided multi-chassis link aggregation group (MLAG). The MLAG is called a vPC on the Cisco ACI side and a VLT on the PowerEdge MX side. The environment is depicted in Figure 44.

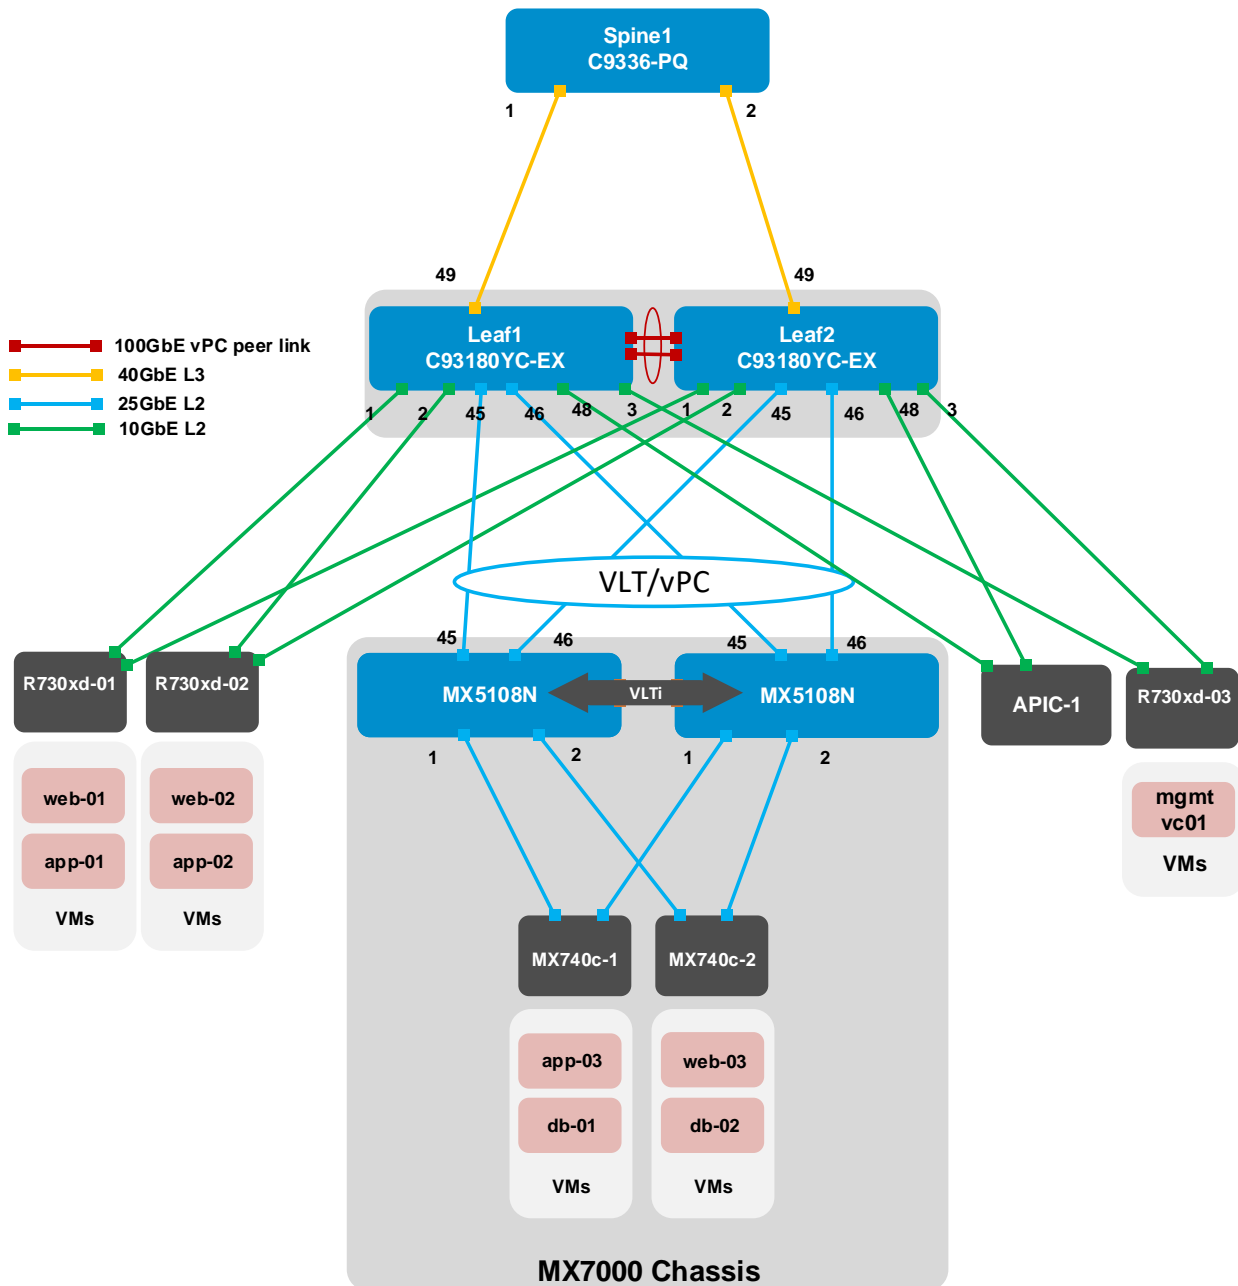


Figure 44 SmartFabric and ACI environment using MX5108n Ethernet switches

The SmartFabric creation and APIC configuration steps are the same as mentioned in Sections 4.2 through 4.5. Refer to these sections to deploy the ACI infrastructure on the MX7000 Chassis in SmartFabric mode using MX5108n switches.

## 5 Validate the configuration

This section covers methods to verify the SmartFabric and ACI environment is configured properly. The screens shown in this chapter depict the MX9116n FSE configuration. Steps for validating the MX5108n will be similar.

### 5.1 MX Validation using OME-M console

This section covers the methods used to verify the SmartFabric and ACI environment is configured properly.

#### 5.1.1 Show the MCM group topology

OME-M console can be used to show the physical cabling of the SmartFabric, perform the following steps:

1. Open the OME-M console and click **Home**.
2. In the **Chassis group** pane, click **View Topology**.
3. Click the lead chassis image and then click **Show Cabling**.
4. Click the  icons to view cable connections as shown in Figure 45.

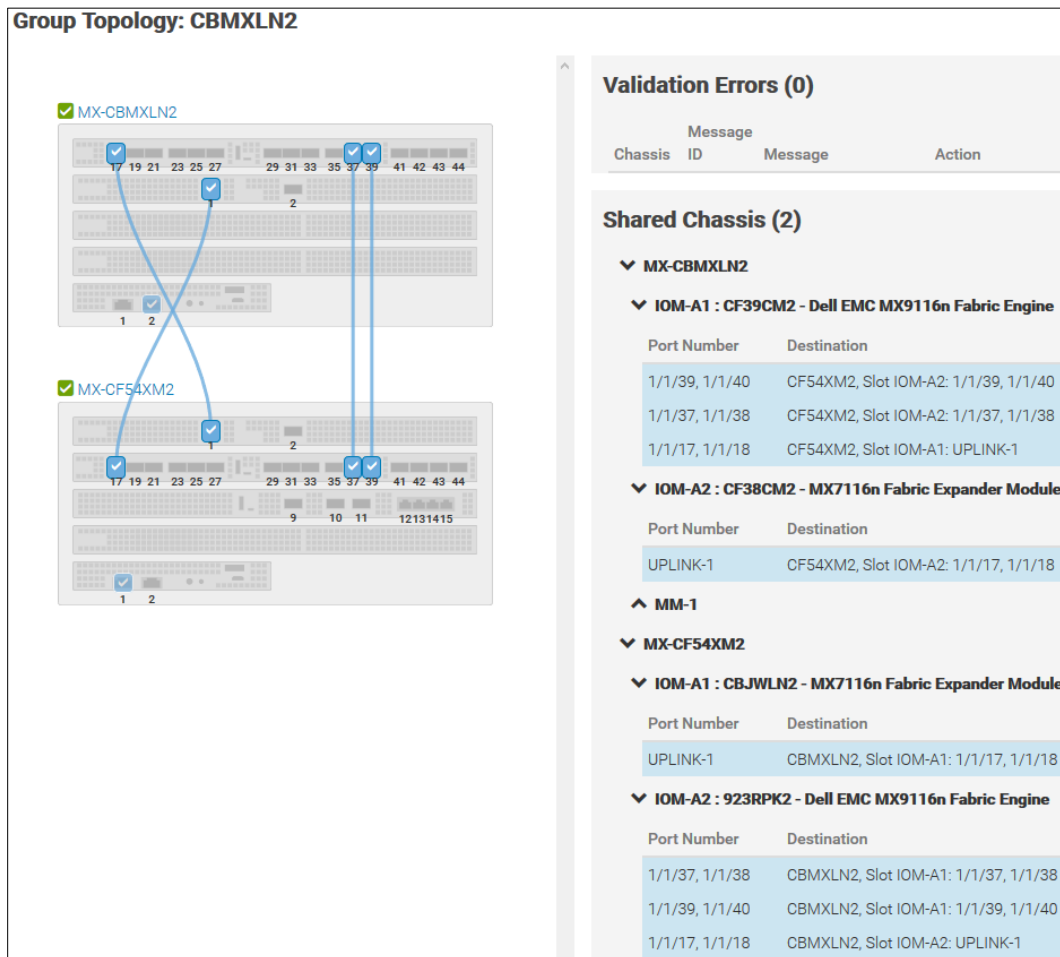


Figure 45 SmartFabric cabling

The Group Topology page shows the MX9116n FSE and MX7116n FEM connections and displays any validation errors. On the MX9116n FSEs, ports 1/1/17-18 are used to connect to the MX7116n FEMs. Ports 1/1/37-40 are used for the VLTi.

## 5.1.2 Show the SmartFabric status

The overall health of the SmartFabric is displayed as follows:

1. Open the OME-M console.
2. From the **Navigation** menu, click **Devices** and then click **Fabric**.
3. Click the fabric name, for example, **SmartFabric1**, to expand the details of the fabric.

The overall status of the fabric is in the upper left corner of the page as shown in Figure 46.



Figure 46 Fabric status details

The left pane of the **Overview** tab lists **Uplinks**, **Switches**, **Servers**, and **ISL Links**. Click the **Switches** link to view the switch health status as shown in Figure 47.

Overview		Topology				
Uplinks	<b>Switches</b>					
<a href="#">Switches</a>	HEALTH	POWER STATE	SERVICE TAG	CHASSIS	SLOT	MODEL
Servers	✓ Ok	On	CF39CM2	MX-CBMXLN2	IOM-A1	Dell EMC MX9116n Fabric Engine
ISL Links	✓ Ok	On	923RPK2	MX-CF54XM2	IOM-A2	Dell EMC MX9116n Fabric Engine

Figure 47 SmartFabric switch status

Click the **Servers** link to view the server health status as shown in Figure 48.

Overview		Topology						
Uplinks	Servers							
Switches	HEALTH	POWER STATE	NAME	SERVICE TAG	CHASSIS	SLOT	MODEL	
Servers	✔ Ok	On	Sled-1	ST0000C	MX-CBMXLN2	Sled-1	PowerEdge MX740c	
ISL Links	✔ Ok	On	Sled-3	1S34MN2	MX-CBMXLN2	Sled-3	PowerEdge MX740c	
	✔ Ok	On	Sled-1	ST00000	MX-CF54XM2	Sled-1	PowerEdge MX840c	
	✔ Ok	On	Sled-3	1S35MN2	MX-CF54XM2	Sled-3	PowerEdge MX740c	

Figure 48 SmartFabric server status

Select the **Topology** tab to view uplinks and fabric connections. Figure 49 shows the VLT port channel connection. Uplink01 is connected to the Cisco ACI vPC using ports 1/1/41-1/1/42 on each MX9116n FSE. The VLTi connection between the two MX9116n FSEs is also shown.

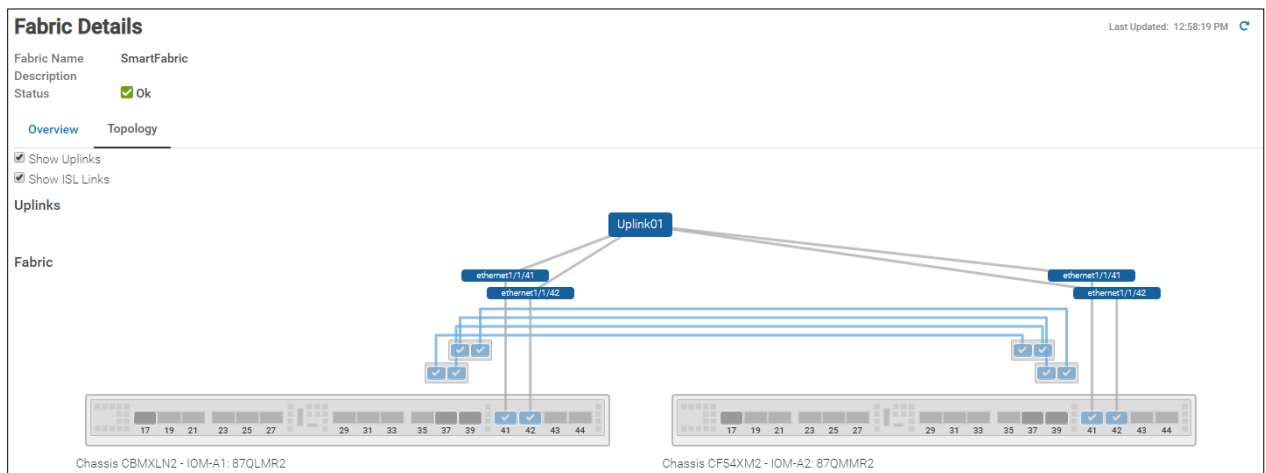


Figure 49 Uplink and VLTi (ISL) connections

The connection details display in the table at the bottom of the **Topology** page as shown in Figure 50.

^ Chassis CF54XM2					
^ IOM-A2: IOM-A2 (87QMMR2) - MX9116n Fabric Engine					
PORT NUMBER	OPERATIONAL STATUS	PORT CONFIGURATION	PORT ROLE	UPLINK NAME	DESTINATION
ethernet1/1/39	Up	NoBreakout	ISL		CBMXMLN2, Slot IOM-A1: ethernet1/1/39
ethernet1/1/40	Up	NoBreakout	ISL		CBMXMLN2, Slot IOM-A1: ethernet1/1/40
ethernet1/1/42	Up	NoBreakout	Uplink	Uplink01	Switch:4c776df1ee7b Port: Eth1/49
ethernet1/1/37	Up	NoBreakout	ISL		CBMXMLN2, Slot IOM-A1: ethernet1/1/37
ethernet1/1/38	Up	NoBreakout	ISL		CBMXMLN2, Slot IOM-A1: ethernet1/1/38
ethernet1/1/41	Up	NoBreakout	Uplink	Uplink01	Switch:00be75194011 Port: Eth1/49
^ Chassis CBMXMLN2					
^ IOM-A1: IOM-A1 (87QLMR2) - MX9116n Fabric Engine					
PORT NUMBER	OPERATIONAL STATUS	PORT CONFIGURATION	PORT ROLE	UPLINK NAME	DESTINATION
ethernet1/1/42	Up	NoBreakout	Uplink	Uplink01	Switch:4c776df1ee7d Port: Eth1/51
ethernet1/1/38	Up	NoBreakout	ISL		CF54XM2, Slot IOM-A2: ethernet1/1/38
ethernet1/1/37	Up	NoBreakout	ISL		CF54XM2, Slot IOM-A2: ethernet1/1/37
ethernet1/1/41	Up	NoBreakout	Uplink	Uplink01	Switch:00be75194013 Port: Eth1/51
ethernet1/1/39	Up	NoBreakout	ISL		CF54XM2, Slot IOM-A2: ethernet1/1/39
ethernet1/1/40	Up	NoBreakout	ISL		CF54XM2, Slot IOM-A2: ethernet1/1/40

Figure 50 SmartFabric topology connection details

### 5.1.3 Show port status

The OME-M console can be used to show MX9116n FSE port status, toggle administrative states, configure breakouts, MTU settings, and auto-negotiation.

1. Open the OME-M console.
2. From the **Navigation** menu, click **Devices** and then click **I/O Modules**.
3. Click an IOM name for the first MX9116n FSE, for example, **IOM-A1**. The **IOM Overview** page for that device displays.
4. On the **IOM Overview** page, click **Hardware**, and then click **Port Information**.

Figure 51 shows ports 1/1/1 and 1/1/3 are up. Ports 1/1/1 and 1/1/3 are connected to the compute sleds in the local chassis. The figure also shows the uplinks to the Cisco ACI leafs, using port channel 1, are up. It also shows the VLTi ports, using port channel 1000, are up.

PORT NUMBER	PORT NAME	PORT DESCRIPT...	↓ OPERATIONAL ST...	ADMIN STATE	CURRENT SPEED	PORT CONFIGU...	OPTICS TYPE	MTU SIZE	AUTO NEGOTIA...	PORT ROLE
ethernet1/1/1			Up	Enabled	25 Gb/s		Fixed	9216	Enabled	EndHost
ethernet1/1/3			Up	Enabled	25 Gb/s		Fixed	9216	Enabled	EndHost
port-channel1			Up	Enabled	200 Gb/s			9216	Disabled	Uplink
ethernet1/1/42			Up	Enabled	100 Gb/s		QSFP28	9216	Disabled	Uplink
ethernet1/1/41			Up	Enabled	100 Gb/s		QSFP28	9216	Disabled	Uplink
port-channel1000			Up	Enabled	400 Gb/s			9216	Disabled	ISL
ethernet1/1/37			Up	Enabled	100 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/39			Up	Enabled	100 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/40			Up	Enabled	100 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/38			Up	Enabled	100 Gb/s		QSFP28-DD	9216	Enabled	ISL

Figure 51 IOM port information

## 5.2 Validation using the MX9116n FSE CLI

The CLI commands shown in this section are available to help validate the configuration. The commands and output shown below are from the MX9116n FSE in the first chassis. The CLI output from the MX9116n FSE in the second chassis, not shown, is similar.

**Note:** The MX9116n FSE CLI is accessible using SSH. The default username and password are both `admin`.

### 5.2.1 show switch-operating-mode

Use the `show switch-operating-mode` command to display the current operating mode.

```
MX9116n-1# show switch-operating-mode
Switch-Operating-Mode : Smart Fabric Mode
```

### 5.2.2 show discovered-expanders

The `show discovered-expanders` command is only available on the MX9116n FSE and displays the MX7116n FEMs service tag attached to the MX9116n FSEs and the associated port-group and virtual slot.

```
MX9116n-1# show discovered-expanders
Service Model      Type  Chassis  Chassis-slot  Port-group  Virtual
tag              service-tag                               Slot-Id
```



```
-----
CBJWLN2   MX7116n FEM 1      CF54XM2      A1           1/1/1       71
```

### 5.2.3 show unit-provision

The **show unit-provision** command is only available on the MX9116n FSE. It displays the unit ID, name, and the state of each MX7116n FEM attached to the MX9116n FSE.

```
MX9116n-1# show unit-provision
Node ID | Unit ID | Provision Name | Discovered Name | State |
-----+-----+-----+-----+-----+
1       | 71      | CBJWLN2       | CBJWLN2         | up    |
```

(output truncated)

### 5.2.4 show vlt domain-id

The **show vlt domain-id** command validates the VLT configuration status. The role of one switch in the VLT pair is primary (not shown), and its peer switch is assigned the secondary role. The VLT domain ID of 255 is automatically configured in SmartFabric mode. The VLTi link Status and VLT Peer Status must both be up. SmartFabric automatically configures the VLTi as port channel 1000.

```
MX9116n-1# show vlt 255
Domain ID           : 255
Unit ID            : 1
Role               : secondary
Version            : 1.0
Local System MAC address : 20:04:0f:00:b8:1e
VLT MAC address    : 20:04:0f:00:b8:1e
IP address         : fda5:74c8:b79e:1::1
Delay-Restore timer : 90 seconds
Peer-Routing       : Disabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
port-channel1000   : up
```

VLT Peer Unit ID	System MAC Address	Status	IP Address	Version
2	20:04:0f:00:9d:1e	up	fda5:74c8:b79e:1::2	1.0

### 5.2.5 show vlt domain-id vlt-port-detail

The **show vlt domain-id vlt-port-detail** command shows the VLT port channel status for both VLT peers. The VLT in this example is connected to the Cisco ACI vPC. It is automatically configured in port channel 1, and it consists of two ports on each switch.

```
MX9116n-1# show vlt 255 vlt-port-detail
vlt-port-channel ID : 1
VLT Unit ID      Port-Channel      Status      Configured ports      Active ports
-----+-----+-----+-----+-----+
* 1              port-channel1      up          2                      2
  2              port-channel1      up          2                      2
```

## 5.2.6 show interface port channel summary

The **show interface port-channel summary** command shows the LAG number (VLT port channel 1 in this example), the mode, status and ports used in the port channel.

```
MX9116n-1# show interface port-channel summary
LAG      Mode      Status   Uptime      Ports
1        L2-HYBRID up        00:29:20   Eth 1/1/41 (Up)
                                                Eth 1/1/42 (Up)
```

## 5.2.7 show lldp neighbors

The **show lldp neighbors** command shows information about directly connected devices. Ports 1/1/1, 1/1/5, 1/71/1, and 1/71/3 are connected to the four compute sleds.

---

**Note:** Ports 1/71/1 and 1/71/3 are the compute sleds connected to the MX7116n FEM in the other chassis.

---

Two instances display for each port connected to a compute sled. One instance is the compute sled iDRAC. The iDRAC uses connectivity to the mezzanine card to advertise LLDP information. It includes the iDRAC name in the `Rem Host Name` column, the sled service tag and mezzanine card number-port-partition in the `Rem Port ID` column, and the iDRAC MAC address in the `Rem Chassis Id` column. The second instance is the mezzanine card itself and the MAC address of the mezzanine card port is shown.

Ports 1/1/37-1/1/40 are the VLTi interfaces for the SmartFabric. Ports 1/1/43-1/1/44 are the links in VLT port channel 1 connected to the Cisco ACI leaf switches.

```
MX9116n-1# show lldp neighbors
Loc PortID      Rem Host Name      Rem Port Id      Rem Chassis Id
-----
ethernet1/1/1  Not Advertised    f4:e9:d4:f2:6f:26    f4:e9:d4:f2:6f:26
ethernet1/1/1  MX740c-1-1-idrac  ST0000C NIC.Mezzanine.1A-1-1  d0:94:66:2d:b3:f4
ethernet1/1/3  Not Advertised    24:6e:96:9c:e5:da    24:6e:96:9c:e5:da
ethernet1/1/3  MX740c-1-3-idrac  1S34MN2 NIC.Mezzanine.1A-1-1  d0:94:66:29:ff:27
ethernet1/1/37 MX9116n-2        ethernet1/1/37      20:04:0f:00:9d:1e
ethernet1/1/38 MX9116n-2        ethernet1/1/38      20:04:0f:00:9d:1e
ethernet1/1/39 MX9116n-2        ethernet1/1/39      20:04:0f:00:9d:1e
ethernet1/1/40 MX9116n-2        ethernet1/1/40      20:04:0f:00:9d:1e
ethernet1/1/41 Leaf1             Eth1/51             00:be:75:19:40:13
ethernet1/1/42 Leaf2             Eth1/51             4c:77:6d:f1:ee:7d
ethernet1/71/1 Not Advertised    f4:e9:d4:f2:6f:da    f4:e9:d4:f2:6f:da
ethernet1/71/1 MX840c-2-1-idrac  ST00000 NIC.Mezzanine.1A-1-1  d0:94:66:2d:b5:2c
ethernet1/71/3 Not Advertised    24:6e:96:9c:e5:48    24:6e:96:9c:e5:48
ethernet1/71/3 MX740c-2-3-idrac  1S35MN2 NIC.Mezzanine.1A-1-1  d0:94:66:29:fa:f4
```

## 5.2.8 show qos system

The **show qos system** command displays the QoS configuration applied to the system. The command is useful to verify the service policy created automatically by the SmartFabric deployment.

```
MX9116n-1# show qos system
Service-policy (input): PM_VLAN
ETS Mode : off
```

## 5.2.9 show policy-map

Using the service policy from **show qos system**, the **show policy-map** command displays QoS policy details including class maps and QoS group settings. The QoS group values should match those configured for each VLAN. See Section 2.7 in the [Dell EMC PowerEdge MX SmartFabric Configuration and Troubleshooting Guide](#) for more information on QoS groups.

```
MX9116n-1# show policy-map
Service-policy (application) input: policy-iscsi
Service-policy (qos) input: PM_VLAN
  Class-map (qos): CM1611
    set qos-group 5
  Class-map (qos): CM1612
    set qos-group 5
  Class-map (qos): CM1613
    set qos-group 5
  Class-map (qos): CM1614
    set qos-group 3
  Class-map (qos): CM1615
    set qos-group 3
  Class-map (qos): CM1616
    set qos-group 3
```

## 5.2.10 show class-map

The **show class-map** displays details for all the configured class maps. For example, the association between CM1611 and VLAN 1611 is shown.

```
MX9116n-1# show class-map
Class-map (application): class-iscsi
Class-map (qos): class-trust
Class-map (qos): CM1611(match-any)
  Match: mac vlan 1611
Class-map (qos): CM1612(match-any)
  Match: mac vlan 1612
Class-map (qos): CM1613(match-any)
  Match: mac vlan 1613
Class-map (qos): CM1614(match-any)
  Match: mac vlan 1614
Class-map (qos): CM1615(match-any)
  Match: mac vlan 1615
Class-map (qos): CM1616(match-any)
  Match: mac vlan 1616
```

## 5.3 SmartFabric Services – Troubleshooting commands

The following commands allow user to view various SmartFabric Services configuration information. These commands can also be used as troubleshooting purpose on SmartFabric OS10.

These commands are available in OS10.5.0.1

### 5.3.1 show smartfabric cluster

The **show smartfabric cluster** command is used to see if node is part of the cluster. This displays the cluster information of the node such as node role, service, virtual IP address, and the node domain. It can also be used to verify role of the node as either Backup or Master.

```
MX9116n-1# show smartfabric cluster
```

```
-----  
CLUSTER DOMAIN ID : 97  
VIP                : fde1:53ba:e9a0:de14:0:5eff:fe00:197  
ROLE               : BACKUP  
SERVICE-TAG      : 87QLMR2  
-----
```

### 5.3.2 show smartfabric cluster member

The **show smartfabric cluster member** command is used to see the member details of the cluster. This displays the cluster member information such as service-tag, IP address, status, role, type of each node, and the service tag of the chassis where the node belongs.

```
MX9116n-1# show smartfabric cluster member
```

```
Service-tag  Status  Role  Type  Chassis-Service-Tag  Chassis-Slot  
-----  
87QLMR2     ONLINE  BACKUP  MX9116n  CBMXLN2              A1  
87QMMR2     ONLINE  MASTER  MX9116n  CF54XM2              A2
```

### 5.3.3 show smartfabric details

The **show smartfabric details** command is used to see the all configured fabric details. This displays which nodes are part of the fabric, status of the fabric, and the design type associated with the fabric.

```
MX9116n-1# show smartfabric details
```

```
-----  
Name          : SmartFabric1  
Description   :  
ID           : b16b835e-9c46-4c2b-b1ed-a11269bdea3eDesign  
Type         : 2xMX9116n_Fabric_Switching_Engines_in_different_chassis  
Validation Status: VALID  
VLTi Status  : VALID  
Placement Status : VALID  
Nodes       : 87QLMR2, 87QMMR2  
-----
```

### 5.3.4 show smartfabric uplinks

The **show smartfabric uplinks** command is used to verify the uplinks configured across the nodes in the fabric. This displays name, description, id, media type, native vlan, configured interfaces, and network profile associated with the fabric.

```
MX9116n-1# show smartfabric uplinks
```

```
-----  
Name                : Uplink01  
Description          :  
ID                  : ffa4bdfd-fd4a-4301-877a-860c93f9df39  
Media Type          : ETHERNET  
Native Vlan         : 1  
Untagged-network    :  
Networks            : ec1c6d5e-3945-41c1-92d2-371e5215c911  
Configured-Interfaces : 87QLMR2:ethernet1/1/41, 87QLMR2:ethernet1/1/42,  
87QMMR2:ethernet1/1/41, 87QMMR2:ethernet1/1/42  
-----
```

### 5.3.5 show smartfabric nodes

The **show smartfabric nodes** command is used to view the details of the nodes that are part of the cluster. This command helps the user to view the status of a node and the chassis details of the node that it belongs to.

```
MX9116n-1# show smartfabric nodes
```

```
-----  
Service-Tag   Type   Status           Mode           Chassis-Service   Chassis-Slot  
Tag  
-----  
CBJWLN2      MX7116n           NOT-APPLICABLE  CF54XM2        A1  
87QLMR2      MX9116n   ONLINE          FABRIC          CBMXLN2         A1  
CF38CM2      MX7116n           NOT-APPLICABLE  CBMXLN2        A2  
87QMMR2      MX9116n   ONLINE          FABRIC          CF54XM2        A2  
7WGQXC2      MX5108n   ONLINE          FULL-SWITCH     CBMXLN2        B1  
-----
```

### 5.3.6 show smartfabric networks

The **show smartfabric networks** command displays all of the network profile information such as the ID, type, QoS priority, and VLAN.

```
MX9116n-1# show smartfabric networks
```

```
-----  
Name           Type           QosPriority      Vlan  
-----  
web            GENERAL_PURPOSE  SILVER          1614  
db             GENERAL_PURPOSE  SILVER          1616  
VLAN001       GENERAL_PURPOSE  BRONZE          1  
app           GENERAL_PURPOSE  SILVER          1615  
vMotion       VM_MIGRATION    PLATINUM        1612  
ESXi_Mgmt     HYPERVERSOR_MANAGEMENT  PLATINUM        1611  
vSAN          STORAGE_DATA_REPLICATION  PLATINUM        1613  
-----
```

## 5.4 Cisco ACI validation

### 5.4.1 Verify vPC configuration

Verify the vPC connection from the Cisco ACI fabric to the Dell MX SmartFabric VLT, as shown in Figure 52, that it is up and properly configured to allow the designated VLANs and EPGs. This is done as follows:

1. In the APIC GUI, click **Fabric > Inventory > Pod name > Leaf name > Interfaces > vPC Interfaces** and drill down to the applicable port channel vPC policy group as shown in Figure 52.
2. Verify that the port channel shows as **lACP-active** and that the **Oper State** shows as **Up**.

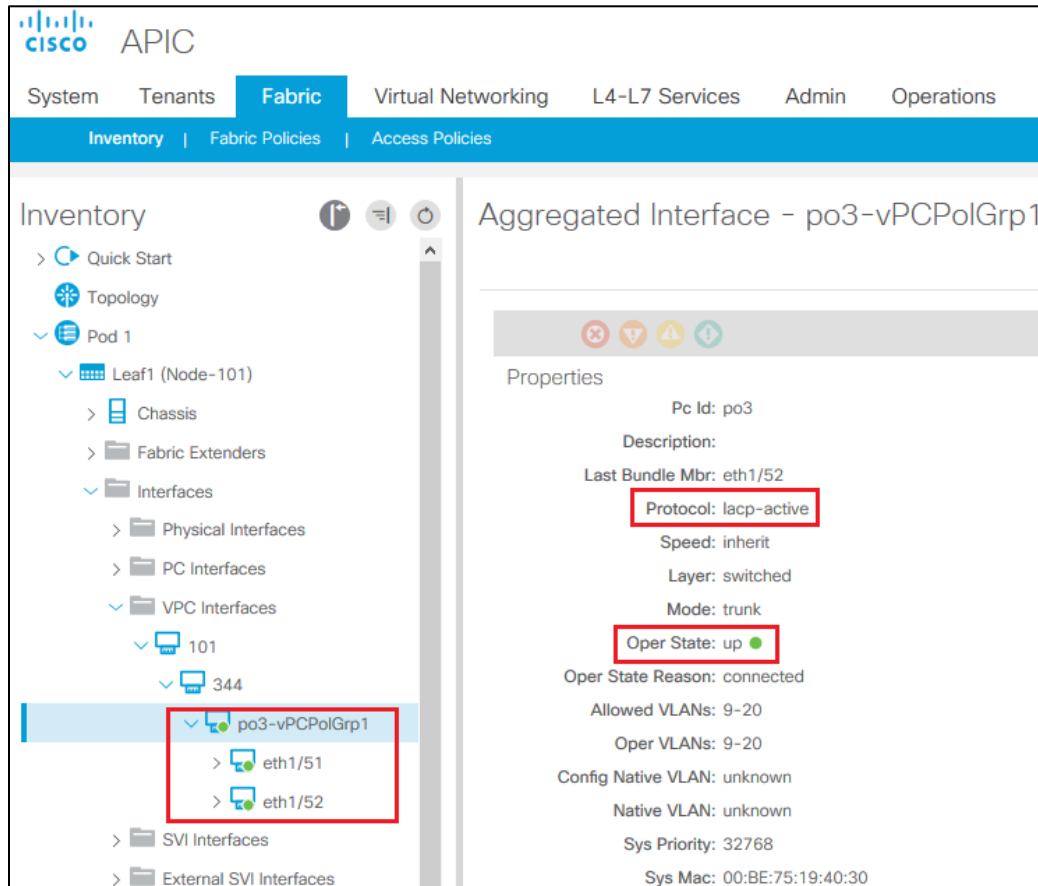


Figure 52 Cisco ACI vPC port channel and interfaces

3. Verify that all of the leaf switch interfaces in the vPC, for example, eth1/51-52, are listed beneath the port channel and are also **Up**.
4. With the port channel/vPC interface policy group selected in the left pane, click **VLANs** at the top of the right pane as shown in Figure 53.

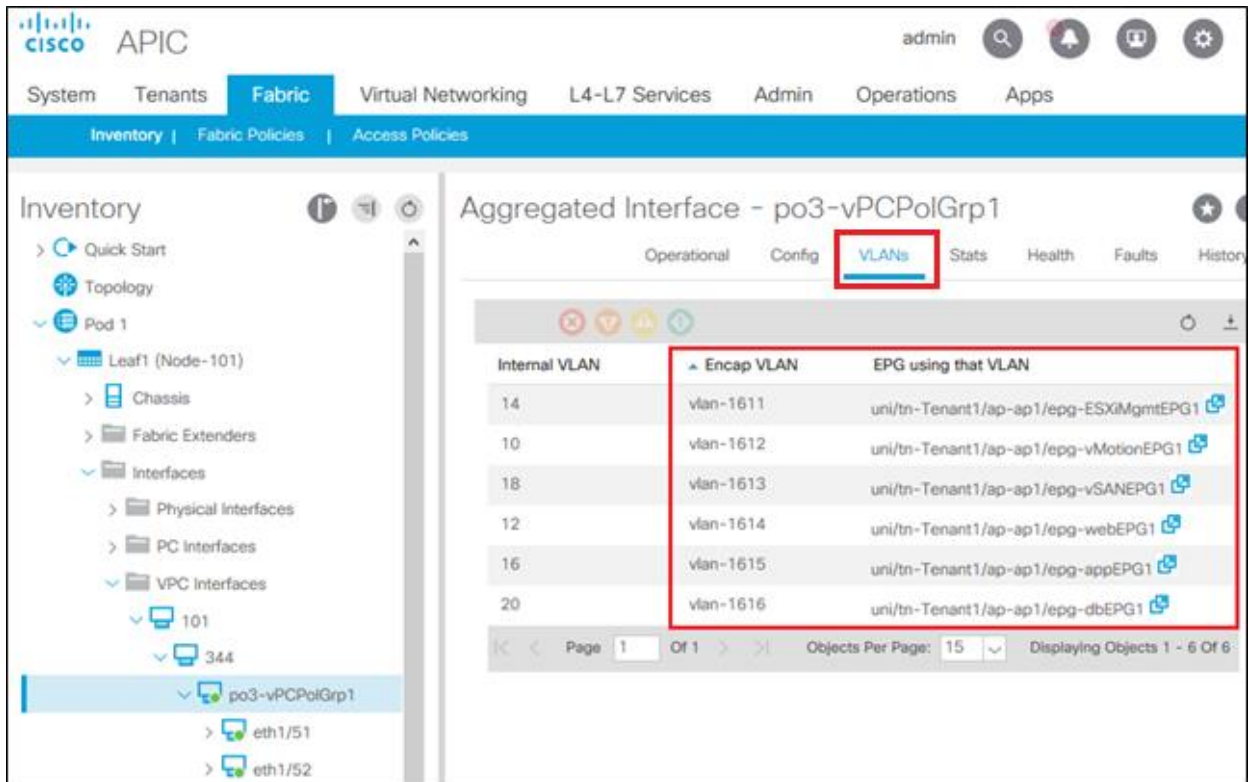


Figure 53 Cisco ACI vPC port channel VLANs and EPGs

5. Verify that the port channel includes all required VLANs, and that the EPGs are mapped to the correct VLANs.
6. Repeat the steps in this section for the remaining leaf switch.

## 5.4.2 Verify physical interface configuration

The physical, host-connected, interfaces in the validated environment are those connected directly to the PowerEdge R730xd servers as shown in Figure 4.

Verify the physical interfaces from the Cisco ACI fabric to the servers are up and properly configured to allow designated VLANs and EPGs. To verify the configuration, perform the following steps:

1. In the APIC GUI, go to **Fabric > Inventory > Pod 1 > Leaf name > Interfaces > Physical Interfaces** as shown in Figure 54.

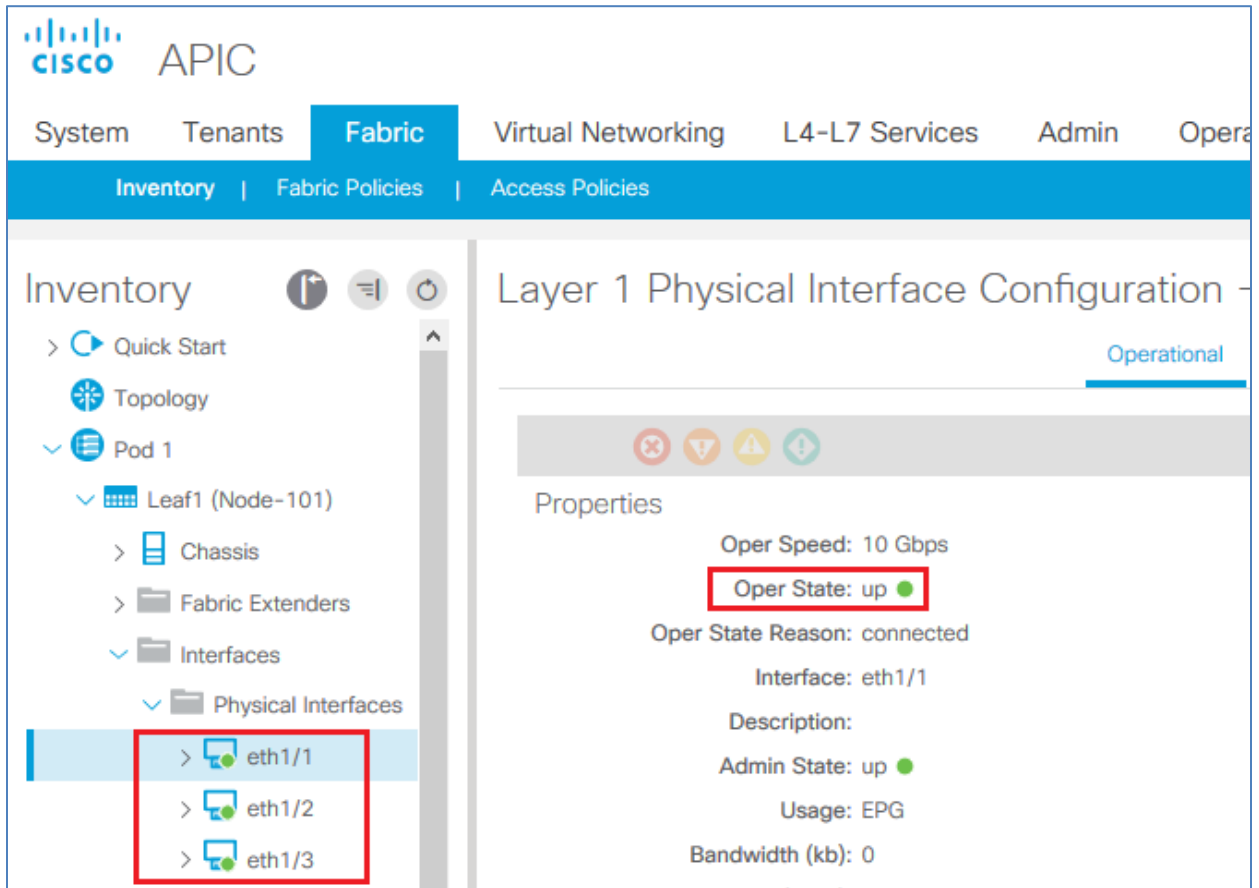


Figure 54 Cisco ACI physical interfaces

2. Verify that the required interfaces, for example, eth1/1-3, show an **up** status.
3. With an interface selected in the left navigational panel, click the **VLANs** tab in the navigation window as shown in Figure 55.

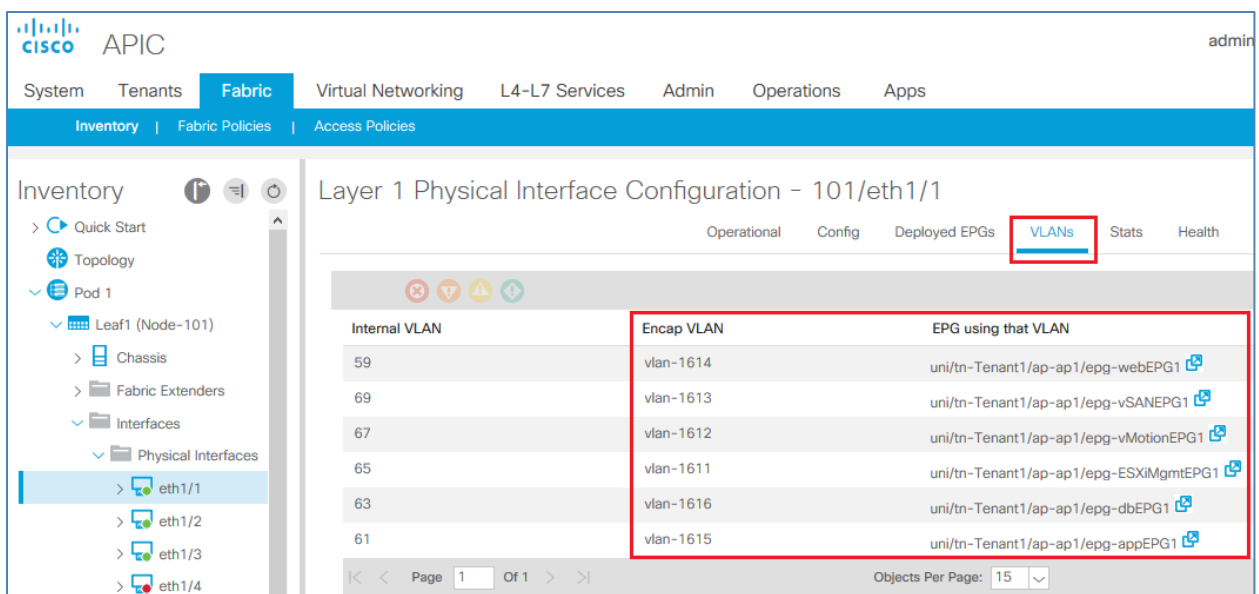


Figure 55 Cisco ACI interface VLANs and EPGs



4. Verify that the interface includes all required VLANs and EPGs. Repeat the steps for the remaining interfaces as needed.
5. Repeat the steps in this section for the remaining leaf switch.

### 5.4.3 Verify ACI learning endpoints

To verify that the ACI is learning endpoints, perform the following steps:

1. In the APIC GUI, go to Tenants > Tenant name > Application Profiles > Application Profile name > Application EPGs > and select an Application EPG.
2. Click the **Operational** tab in the navigation window as shown in Figure 56.

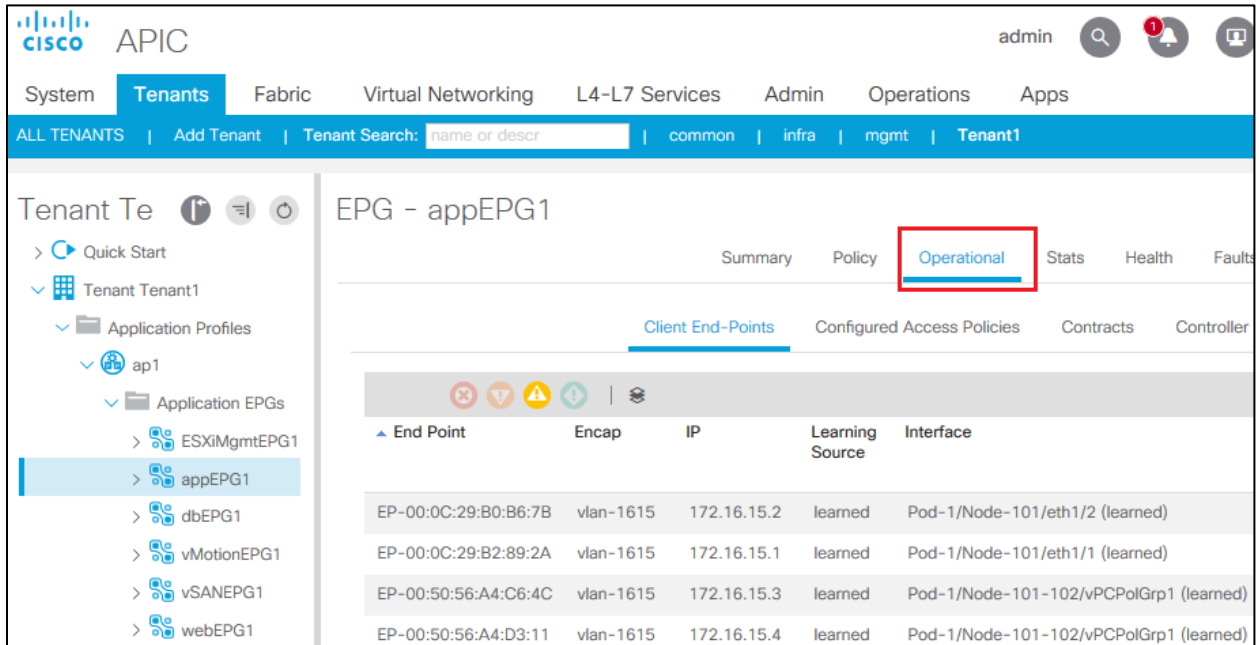


Figure 56 Cisco ACI endpoints in appEPG1

3. Review the listing of the learned endpoints for the selected EPG along with the VLAN, IP address, and interface.
4. Repeat the steps in this section for the remaining Application EPGs.

### 5.4.4 Verify ACI VMM domain integration

To verify ACI vCenter domain integration:

1. In the APIC GUI, Go to **Virtual Networking > VMM Domains > VMware > VDS-ACI**, click the **Operational** option on the top right corner. The vCenter-server and its details such as number of hypervisors and virtual machines, displays.

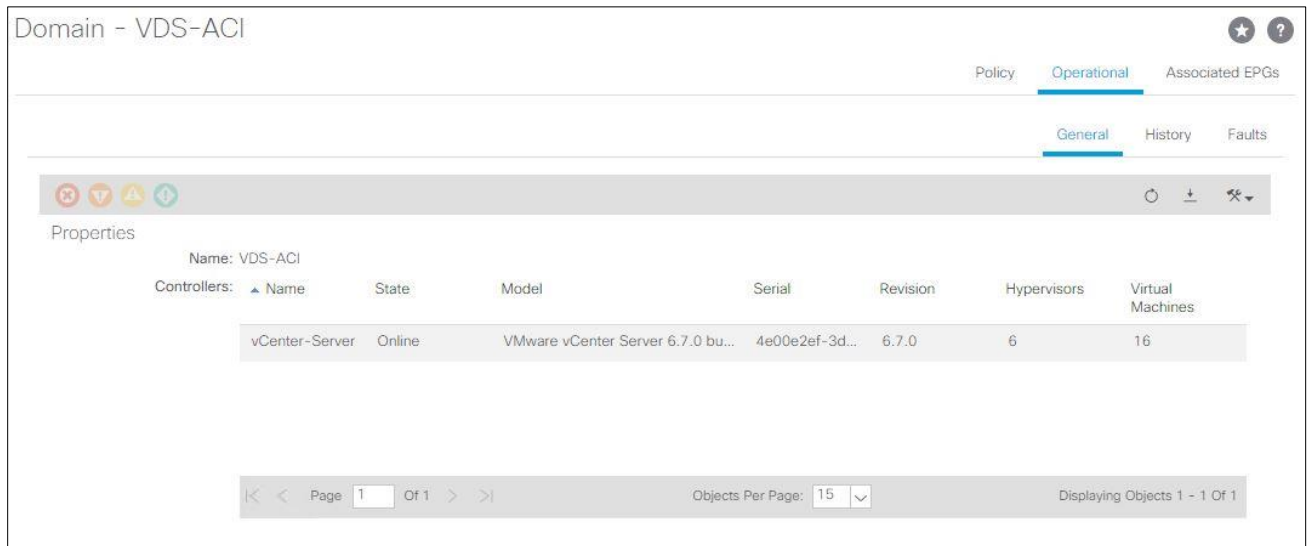


Figure 57 vCenter-Server

2. Select **Associated EPGs** to show the associated EPGs to vCenter Domain.

EPG	Tenant	Application Profile	Deployment Immediacy	Resolution Immediacy	Allow Micro-Segmentation	Vlan Mode	Switching Mode
APP-EPG	Customer-TN1	APP-TN1	On Demand	Immediate	False	Dynamic	native
DB-EPG	Customer-TN1	APP-TN1	On Demand	Immediate	False	Dynamic	native
Test-99	Customer-TN1	APP-TN1	On Demand	Immediate	False	Dynamic	native
TEST-EPG	Customer-TN1	APP-TN1	On Demand	Immediate	False	Dynamic	native
WEB-EPG	Customer-TN1	APP-TN1	On Demand	Immediate	False	Dynamic	native
VMware-MGMT	mgmt	APP-MGMT	Immediate	Immediate	False	Static	native
VMware-vMotion	mgmt	APP-MGMT	Immediate	Immediate	False	Static	native
VMware-VSAN	mgmt	APP-MGMT	Immediate	Immediate	False	Static	native

Figure 58 Associated EPGs to vCenter Domain

3. For more information about vCenter server and its associated credentials, go to **Virtual Networking > VMM Domains > VMware > VDS-ACI > Controllers > vCenter-Server**. This shows the **Datacenter, Management EPG**, and Associated Credential details.

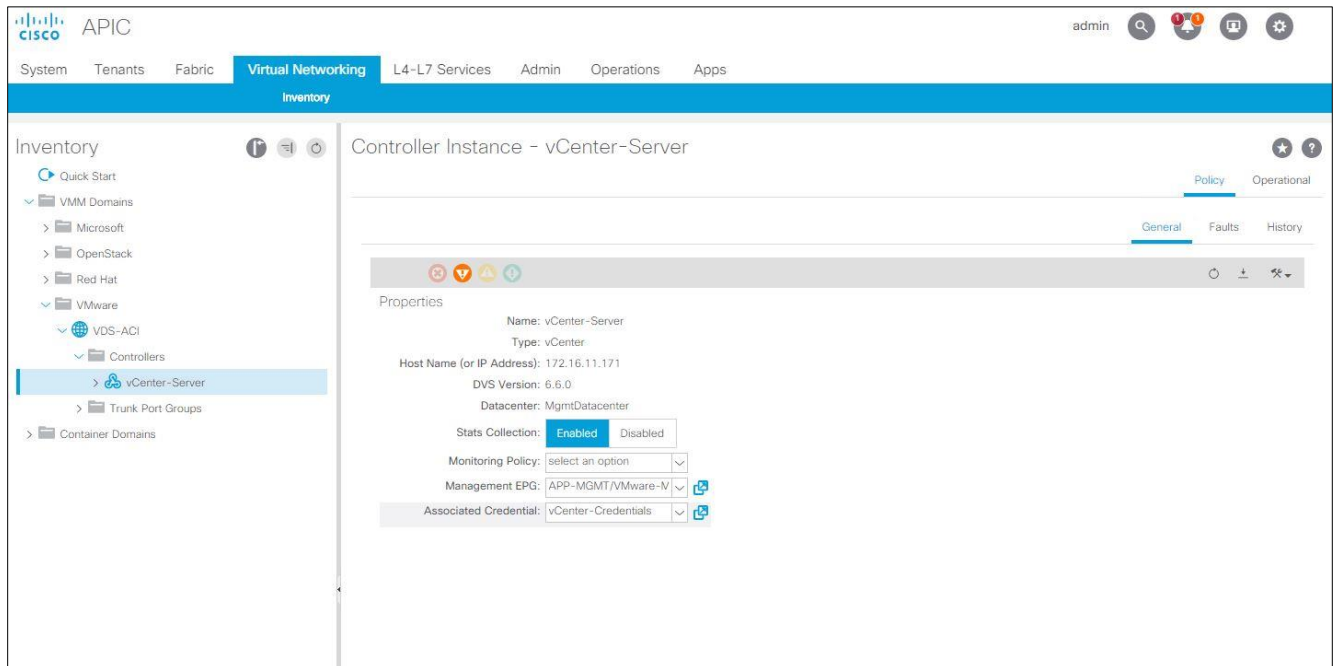


Figure 59 vCenter-Server Detail

4. Choose **vCenter-Server**, then **DVS-VDS-ACI** to see the details about Distributed Virtual Switch.

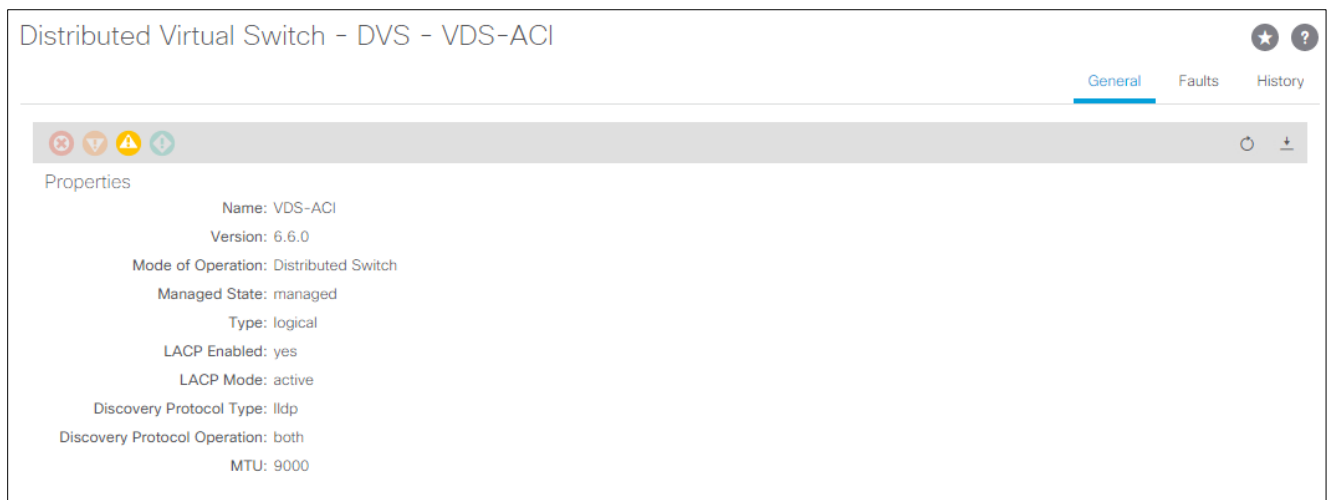


Figure 60 Distributed Virtual Switch (DVS) details

## 5.5 Verify connectivity between VMs

In ACI, by default, communication flows freely within EPGs, but not between EPGs. To enable inter-EPG communication, contracts are configured on the APIC. This example is configured for unrestricted inter-EPG communication as shown in steps 4.2.18 through 4.2.20 in the Section 4.2.

Connectivity is verified by pinging between the VMs as shown in Figure 61. Since inter-EPG communication is allowed using configured contracts, all VMs can ping all other VMs in the topology.

Figure 61 shows the VM named app-01, located in a rack server, successfully pinging the VMs named web-03 and db-04, which are located on MX compute sleds.

```
root@app-01:/#  
root@app-01:/# ping web-03  
PING web-03 (172.16.14.3) 56(84) bytes of data.  
64 bytes from web-03 (172.16.14.3): icmp_seq=1 ttl=63 time=0.509 ms  
64 bytes from web-03 (172.16.14.3): icmp_seq=2 ttl=63 time=0.468 ms  
^C  
--- web-03 ping statistics ---  
2 packets transmitted, 2 received, 0% packet loss, time 999ms  
rtt min/avg/max/mdev = 0.468/0.488/0.509/0.030 ms  
root@app-01:/# ping db-04  
PING db-04 (172.16.16.4) 56(84) bytes of data.  
64 bytes from db-04 (172.16.16.4): icmp_seq=1 ttl=62 time=0.621 ms  
64 bytes from db-04 (172.16.16.4): icmp_seq=2 ttl=62 time=0.461 ms  
64 bytes from db-04 (172.16.16.4): icmp_seq=3 ttl=62 time=0.550 ms
```

Figure 61 Verifying connectivity between VMs

## A Hardware supported in this document

This section covers the rack-mounted networking switches supported by the examples in this guide.

For detailed information about the hardware components related to the MX platform, see the [Dell EMC PowerEdge MX Networking Architecture Guide](#).

### A.1 Dell EMC PowerSwitch S3048-ON management switch

The Dell EMC PowerSwitch S3048-ON is a 1-Rack Unit (RU) switch with forty-eight 1GbE BASE-T ports and four 10GbE SFP+ ports. In this document, one S3048-ON supports out-of-band (OOB) management traffic for all examples.



Figure 62 Dell EMC PowerSwitch S3048-ON

### A.2 Dell EMC Networking MX9116n Fabric Switching Engine (FSE)

The Dell EMC Networking MX9116n Fabric Switching Engine (FSE) is a scalable, high-performance, low latency 25 GbE switch purpose-built for the PowerEdge MX platform. In addition to sixteen internal 25GbE ports, MX9116n FSE also provides two 100GbE QSFP28 ports, two 100 GbE QSFP28 unified ports and twelve 2x100GbE QSFP-28 Double Density (DD) ports. In this document, two MX9116n FSEs are used as IO modules connected to the Nexus leaf switches.



Figure 63 Dell EMC Networking MX9116n Fabric Switching Engine

### A.3 Dell EMC Networking MX5108n Ethernet switch

The Dell EMC Networking MX5108n Ethernet switch is targeted at small PowerEdge MX7000 deployments of one or two chassis. While not a scalable switch, it still provides high-performance and low latency with a non-blocking switching architecture. In addition to eight internal 25 GbE ports, the MX5108n provides one 40GbE QSFP+ port, two 100GbE QSFP28 ports, four 10GbE RJ45 BASE-T ports. In this document, two MX5108n may be used as IO modules connected to the Nexus leaf switches.



Figure 64 Dell EMC Networking MX5108n Ethernet switch

## A.4 Cisco Nexus C93180YC-EX

The Cisco Nexus C93180YC-EX switch is a 1-RU switch with forty-eight 1/10/25GbE ports and six 40/100GbE ports. A pair of Cisco Nexus C93180YC-EX switches is used as Cisco ACI leaf switches in the example in this guide.

## A.5 Cisco Nexus C9336-PQ

The Cisco Nexus C9336-PQ switch is a 2-RU switch with thirty-six 40GbE QSFP+ ports. One Cisco Nexus C9336-PQ switch is used as a Cisco ACI spine switch in the example in this guide.

## B Validated components

The following tables include the hardware, software, and firmware used to configure and validate the environment mentioned in this document.

### B.1 Dell EMC PowerSwitch

Table 7 Dell EMC PowerSwitch and OS version

Qty	Item	OS Version
1	Dell EMC PowerSwitch S3048-ON OOB management switch	10.4.1.2

### B.2 Dell EMC PowerEdge MX7000 chassis and components

Table 8 Dell EMC PowerEdge MX7000 chassis and components

Qty	Item	Version
2	Dell EMC PowerEdge MX7000 chassis	-
4	Dell EMC PowerEdge MX740c sled	-
4	Dell EMC PowerEdge M9002m modules (2 per chassis)	1.10.00
2	Dell EMC Networking MX9116n FSE (1 per chassis)	10.5.0.1
2	Dell EMC Networking MX7116n FEM (1 per chassis)	-
2	Dell EMC Networking MX5108n Ethernet switch	10.5.0.1

Table 9 MX740c sled details

Qty per sled	Item	Version
2	Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz	-
12	16GB DDR4 DIMMs (192GB total)	-
1	Boot Optimized Storage Solution (BOSS) S1 Controller w/ 1x120GB SATA SSD	2.6.13.3011
1	PERC H730P MX	25.5.5.0005
2	600GB SAS HDD	-
1	Intel(R) Ethernet 2x25GbE XXV710 mezzanine card or QLogic 2x25GbE QL41232HMKR mezzanine card	18.5.17 (Intel) or 14.07.07 (QLogic)
-	BIOS	1.6.11
-	iDRAC with Lifecycle Controller	3.22.22.22
	VMware ESXi (Dell EMC Customized)	6.7.0

## B.3 Cisco ACI components

Table 10 Cisco ACI components

Qty	Item	Version
1	Cisco APIC	4.0(3d)
1	Cisco Nexus C9336-PQ spine switch	n9000-14.0(3d)
2	Cisco Nexus C93180YC-EX leaf switches	n9000-14.0(3d)



## C Technical resources

[Dell EMC Networking Guides](#)

[Dell EMC PowerEdge MX IO Guide](#)

[Dell EMC PowerEdge MX Network Architecture Guide](#)

[Dell EMC PowerEdge MX SmartFabric Deployment Video](#)

[Dell EMC PowerEdge MX SmartFabric Deployment with Cisco ACI Video](#)

[MX Port-Group Configuration Errors Video](#)

[MX Port-Group Configuration Video](#)

[Dell EMC OpenManage Enterprise-Modular Edition User's Guide v1.00.01](#)

[OS10 Enterprise Edition User Guide for PowerEdge MX IO Modules Release 10.4.0E R3S](#)

[Dell EMC SmartFabric OS10 User Guide](#)

[Manuals and documents for Dell EMC PowerEdge MX7000](#)

[Manuals and documents for Dell EMC PowerSwitch MX9116n](#)

[Manuals and documents for Dell EMC PowerSwitch S3048-ON](#)

[Dell EMC OME-M v1.00.01 for PowerEdge MX7000 Chassis User's Guide](#)

[Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#)

[Management Networks for Dell EMC Networking](#)

## D Support and feedback

### Contacting Technical Support

Support Contact Information

Web: <http://www.dell.com/support>

Telephone: USA: 1-800-945-3355

### Feedback for this document

Dell EMC encourages readers to provide feedback on the quality and usefulness of this publication by sending an email to [Dell\\_Networking\\_Solutions@Dell.com](mailto:Dell_Networking_Solutions@Dell.com)