

Dell EMC PowerEdge MX and Cisco Application Centric Infrastructure Integration Guide

H18013.2

Abstract

This integration guide provides the steps for integrating the Dell EMC PowerEdge MX platform with the Cisco Application Centric Infrastructure (ACI) environment.

Dell Technologies Solutions

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Introduction.....	6
Overview.....	6
Dell SmartFabric OS10.....	7
VMware vCenter integration - OpenManage Network Integration.....	7
Process flow checklist.....	8
Chapter 2: Requirements.....	11
PowerEdge MX requirements.....	11
Physically cable MX7000 chassis and upstream switches.....	11
Create multichassis management group.....	11
Application Centric Infrastructure.....	11
Spanning Tree Protocol guidelines.....	12
Discovery protocol guidelines.....	12
Chapter 3: SmartFabric Connections to Cisco ACI Leaf Switches.....	13
Validated environment.....	13
Cisco APIC configuration.....	15
Create VLAN pool.....	15
Create physical domain.....	16
Create Attachable Access Entity Profile.....	17
Create port channel policy.....	18
Create Spanning Tree Interface Policy.....	18
Create LLDP Interface policy.....	19
Create Multicast Protocol interface policy.....	20
Create VPC Interface Policy Group.....	22
Create Leaf Access Port Policy Group.....	23
Create Leaf Interface Profile.....	24
Create VPC Domain policy.....	25
Create VPC Explicit Protection Group.....	26
Create Leaf Profile.....	26
Create Tenant.....	27
Create VRF.....	28
Create Bridge Domain.....	29
Create Application Profile.....	30
Create Application EPGs.....	31
Configure Access Entity Profile with EPGs and VLANs.....	32
Create vCenter domain for Cisco ACI and Virtual Machine Manager (VMM) domain integration.....	32
Create Contract filter.....	35
Create Contract.....	36
Apply Contract to VRF.....	37
SmartFabric deployment.....	38
Defining VLANs.....	39
LLDP setting for SmartFabric.....	39
Create SmartFabric.....	40

Create Uplink.....	40
Deploy Server.....	41
Create server template.....	41
Add VLANs to server templates.....	42
Deploy server templates.....	42
Configure vCenter.....	42
SmartFabric connected with MX5108n Ethernet switch and Cisco ACI leaf switches.....	44
Chapter 4: Validating the Configuration.....	46
MX validation using OME-M console.....	46
Configuration validation CLI commands for this guide.....	46
show lldp neighbors.....	46
SmartFabric Services troubleshooting commands.....	47
show smartfabric uplinks.....	47
show smartfabric networks.....	47
Cisco ACI validation.....	47
Verify VPC configuration.....	47
Verify physical interface configuration.....	49
Verify ACI endpoint learning.....	50
Verify ACI VMM domain integration.....	51
Verifying connectivity between VMs.....	52
Chapter 5: Troubleshooting.....	54
Troubleshooting LLDP.....	54
Verify Fabric Management Address in LLDP Message Option is Enabled.....	54
Verify LLDP on Cisco ACI.....	55
Verify LLDP on VMware vSphere Distributed Switch in VMware vCenter.....	56
Appendix A: Full Switch Mode Example	58
Full Switch mode.....	58
Ethernet switch configuration.....	58
Appendix B: Hardware and Software Versions.....	62
Dell EMC switches.....	62
Dell PowerSwitch S3048-ON.....	62
Dell EMC Networking MX9116n FSE.....	62
Dell EMC Networking MX5108n Ethernet switch.....	62
Cisco switches.....	63
Cisco Nexus C93180YC-EX.....	63
Cisco Nexus C9336-PQ.....	63
Validated components and software versions.....	63
Dell EMC PowerSwitch.....	63
Dell EMC PowerEdge MX7000 chassis and components	63
MX740c sled.....	64
MX840c sled.....	64
VMware components.....	64
Cisco ACI components.....	65
Appendix C: Documentation and Support.....	66

Dell Technologies documentation..... 66

OME-M and OS10 compatibility and documentation..... 66

Support and feedback..... 67

Introduction

Overview

The vision at Dell Technologies is to be the essential technology company from the edge, to the core, and to the cloud. Dell Technologies ensures modernization for today's applications and the emerging cloud-native world. Dell EMC Networking 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 Technologies strategy enables business transformations that maximize the benefits of collaborative software and standards-based hardware, including lowered costs, flexibility, freedom, and security. Dell Technologies provides further customer enablement through validated deployment guides that demonstrate these benefits while maintaining a high standard of quality, consistency, and support.

The Dell EMC PowerEdge MX platform is a unified, high-performance data center infrastructure. It 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 (SFS), 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 the 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 multichassis management group and that the reader has a basic understanding of the PowerEdge MX platform.

For a general overview of PowerEdge MX networking concepts, SmartFabric Services (SFS), Full Switch mode, and the Scalable Fabric Architecture see the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

NOTE: For an overview of hardware components supported in this document, see [Appendix B](#).

Dell SmartFabric OS10

The networking market is transitioning from a closed, proprietary stack to open hardware supporting various operating systems. Dell SmartFabric OS10 is designed to allow multilayered disaggregation of the network functionality. While OS10 contributions to Open Source provide users with freedom and flexibility to pick their own third-party networking, monitoring, management, and orchestration applications; SmartFabric OS10 bundles an industry-hardened networking stack featuring standard Layer 2 and Layer 3 protocols over a well-accepted CLI interface.

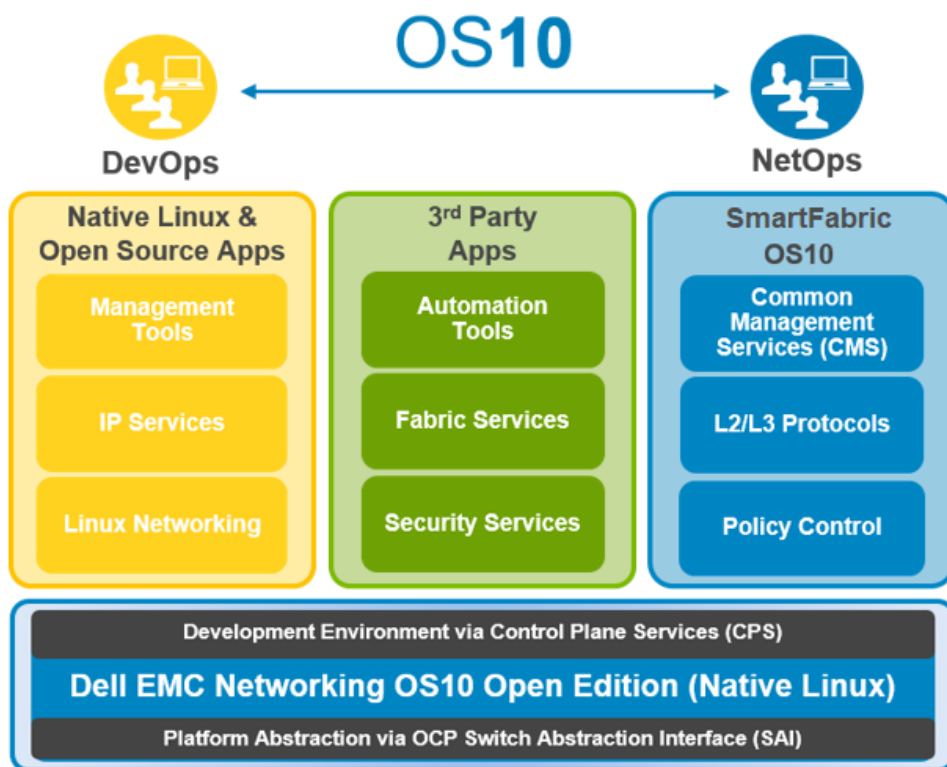


Figure 2. Dell SmartFabric OS10 high-level architecture

VMware vCenter integration - OpenManage Network Integration

Dell OpenManage Network Integration (OMNI) is an external plug-in for VMware vCenter that is designed to complement SmartFabric Services (SFS) by integrating with VMware vCenter to perform fabric automation. With the release of OMNI 2.0, this integration is extended to SFS that runs on PowerEdge MX. This integration automates VLAN changes that occur in VMware vCenter and propagates those changes into the related SFS instances running on the MX platform as shown in the following figure.

The combination of OMNI and Cisco ACI vCenter integration creates a fully automated solution. OMNI and the Cisco APIC recognize changes in vCenter and automatically propagate the changes to the MX SmartFabric and ACI fabric respectively. This allows a VLAN change to be made in vCenter, and it will flow through the entire solution without any manual intervention.

For more information about OMNI, see the SmartFabric Services for OpenManage Network Integration User Guide on the [Dell OpenManage Network Integration for VMware vCenter](#) documentation page.

NOTE: OMNI 2.0 and 2.1 only support VLAN automation with one uplink per SmartFabric.

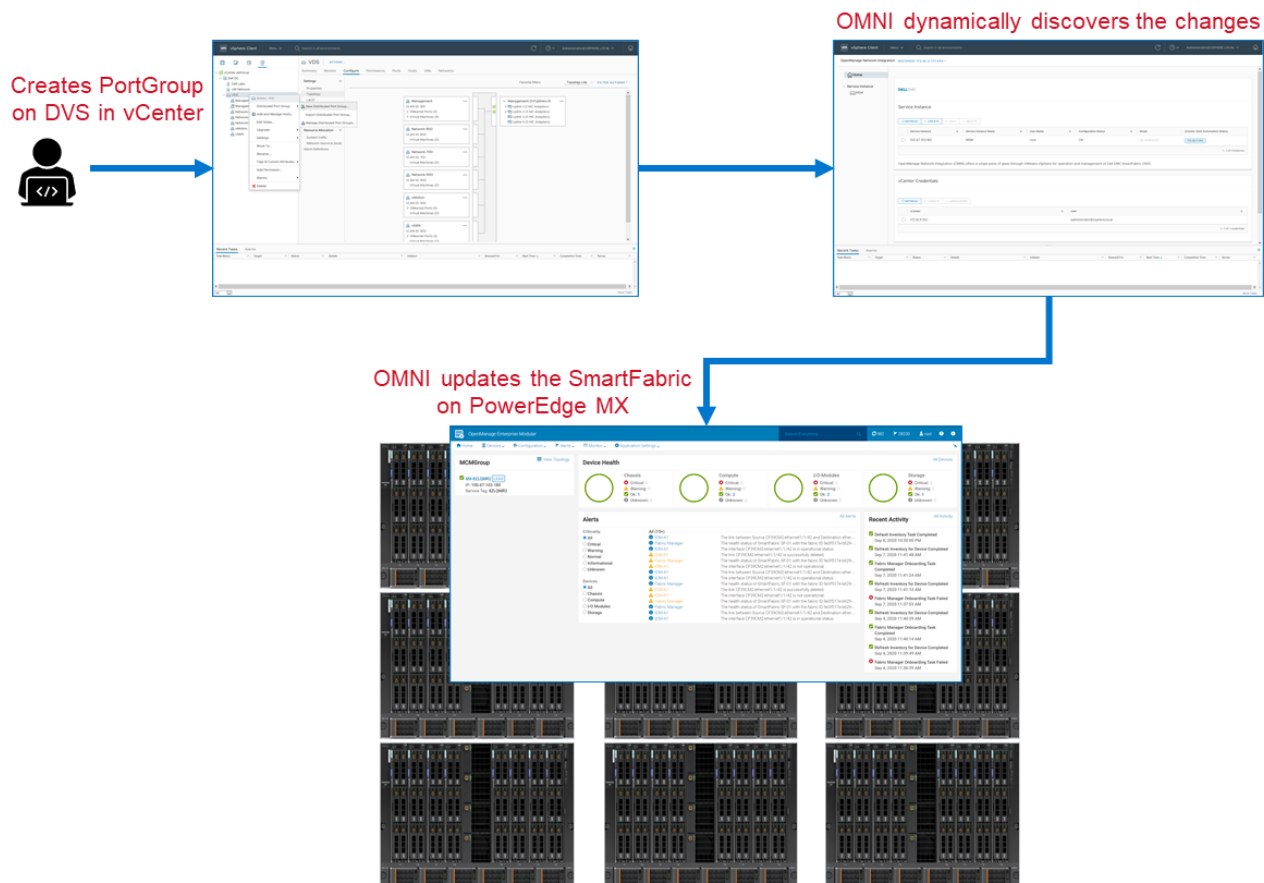


Figure 3. OMNI integration workflow

Process flow checklist

This guide is used with other documents to configure the validated PowerEdge MX SmartFabric and Cisco ACI environment that is shown in [Figure 4](#).

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 in the table below.

Table 1. Step reference table and checklist


☑	Step	Description	Reference	Where to implement
<input type="checkbox"/>	1	Physically cable the MX Chassis and upstream switches	Dell EMC PowerEdge MX Networking Deployment Guide	Hardware
<input type="checkbox"/>	2	Create multichassis management group	This document: Create multichassis management group	OME-M
<input type="checkbox"/>	3	Initial deployment of APIC and Nexus leaf and spine switches	Fabric Initialization and switch discovery	APIC
<input type="checkbox"/>	4	Create VLAN pool	This document: Create VLAN pool	APIC
<input type="checkbox"/>	5	Create physical domain	This document: Create physical domain	APIC

Table 1. Step reference table and checklist (continued)

☒	Step	Description	Reference	Where to implement
<input type="checkbox"/>	6	Create attachable access entity profile	This document: Create attachable access entity profile	APIC
<input type="checkbox"/>	7	Create Port Channel Policy	This document: Create port channel policy	APIC
<input type="checkbox"/>	8	Create Spanning Tree Interface policy	This document: Create Spanning Tree Interface policy	APIC
<input type="checkbox"/>	9	Create LLDP Interface policy	This document: Create LLDP Interface policy	APIC
<input type="checkbox"/>	10	Create Miscabling protocol interface policy	This document: Create Miscabling protocol interface policy	APIC
<input type="checkbox"/>	11	Create vPC interface policy group	This document: Create VPC Interface policy group	APIC
<input type="checkbox"/>	12	Create Leaf Access Port policy group	This document: Create Leaf Access Port policy group	APIC
<input type="checkbox"/>	13	Create Leaf Interface profile	This document: Create Leaf Interface profile	APIC
<input type="checkbox"/>	14	Create VPC Domain policy	This document: Create VPC Domain policy	APIC
<input type="checkbox"/>	15	Create vPC Explicit Protection Group	This document: Create VPC Explicit Protection Group	APIC
<input type="checkbox"/>	16	Create Leaf Profile	This document: Create Leaf Profile	APIC
<input type="checkbox"/>	17	Create Tenant	This document: Create Tenant	APIC
<input type="checkbox"/>	18	Create VRF	This document: Create VRF	APIC
<input type="checkbox"/>	19	Create Bridge Domain	This document: Create Bridge Domain	APIC
<input type="checkbox"/>	20	Create Application Profile	This document: Create Application Profile	APIC
<input type="checkbox"/>	21	Create Application EPGs	This document: Create Application EPGs	APIC
<input type="checkbox"/>	22	Configure Access Entity Profile with EPGs and VLANs	This document: Configure Attachable Access Entity Profile with EPGs and VLANs	APIC
<input type="checkbox"/>	23	Create vCenter domain for Cisco ACI and Virtual Machine Manager (VMM) domain integration:	This document: Create vCenter domain for Cisco ACI and Virtual Machine Manager domain integration	APIC
<input type="checkbox"/>	24	Create Contract Filter	This document: Create Contract Filter	APIC
<input type="checkbox"/>	25	Create Contract	This document: Create Contract	APIC
<input type="checkbox"/>	26	Apply Contract to VRF	This document: Apply Contract to VRF	APIC
<input type="checkbox"/>	27	Define VLANs	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	28	Create SmartFabric	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	29	Create Uplink	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	30	Create Server Template	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	31	Add VLANs to server template	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	32	Deploy Server Template	Dell EMC PowerEdge MX Networking Deployment Guide	OME-M
<input type="checkbox"/>	33	Create a Data Center	VMware document: Create a Data Center	vCenter

Table 1. Step reference table and checklist (continued)

☒	Step	Description	Reference	Where to implement
<input type="checkbox"/>	34	Create a Cluster	VMware document: Create a Cluster	vCenter
<input type="checkbox"/>	35	Configure a Cluster	VMware document: Configure a Cluster	vCenter
<input type="checkbox"/>	36	Add a Hosts	VMware document: Add a Host	vCenter
<input type="checkbox"/>	37	Create a Virtual Machine	VMware document: Create a Virtual Machine	vCenter
<input type="checkbox"/>	38	Create VDS and Set Up Networking	VMware document: Setting up Networking with vSphere Distributed Switches	vCenter

 **NOTE:** For more information about the configuration of VMware vSphere, see the [Organizing your Inventory](#) section within the VMware vSphere Product Documentation page.

Requirements

PowerEdge MX requirements

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

Configuration of SmartFabric on PowerEdge MX with Cisco ACI makes the following assumptions:

- All MX7000 chassis and management modules are cabled correctly and in a multichassis management group.
- PowerEdge and Cisco ACI platforms are healthy.
- VLTi cables between MX switches have been connected.
- OpenManage Enterprise Modular is at version 1.20.10 or later, and SmartFabric OS10 is at version 10.5.1.6 or later.

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

Physically cable MX7000 chassis and upstream switches

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

- For Management Module cabling, see *Direct from Development - PowerEdge MX7000 Chassis Management Network Cabling* on the [Documentation](#) tab of the PowerEdge MX7000 support site.
- For information and requirements on cabling a PowerEdge MX Scalable Fabric, MX switches to the upstream network, and VLTi connections, see the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

Create multichassis management group

For deployment that uses more than one MX chassis, the chassis must be in a multichassis management (MCM) group. See the *Dell OpenManage Enterprise-Modular Edition for PowerEdge MX7000 Chassis User's Guide* on the [Documentation](#) page of the support site for information about how to create the MCM group.

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

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 the table below.

Table 2. APIC leaf and spine node IDs and names

Node ID	Node name
101	Leaf1
102	Leaf2

Table 2. APIC leaf and spine node IDs and names (continued)

Node ID	Node name
201	Spine1

The networks used are shown in the following table 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

Spanning Tree Protocol guidelines

ACI switches do not actively participate in Spanning Tree Protocol (STP). ACI switches forward spanning tree Bridge Protocol Data Units (BPDUs) across EPGs on which they are received. The spanning tree links are peer-to-peer (P2P), which does not cause loops until ACI acts as a hub for BPDUs. To avoid loops on switches connected to ACI, you can share the spanning tree link. You can also create the spanning tree interface policy to change spanning tree protocol settings on ACI Fabric.

In a PowerEdge MX environment integrated with the Cisco ACI, the best practice is to use **Ethernet - No Spanning Tree** as the uplink type for the MX SmartFabric. This disables STP on the MX I/O modules.

When using the **Legacy Ethernet** uplink in a PowerEdge MX SmartFabric environment, disable STP on the MX I/O modules by running the `spanning-tree disable` command globally on the IOM command-line interface on both switches before creating a Legacy Ethernet uplink.

NOTE: For information about creating an uplink in a PowerEdge MX environment, see the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

NOTE: See the [Cisco Application Centric Infrastructure Design Guide White Paper](#) for additional infrastructure design information.

Discovery protocol guidelines

ACI switches discover neighbor devices, hosts, and VMs through discovery protocols. Link Layer Discovery Protocol (LLDP) or Cisco Discovery Protocol (CDP) can be implemented for use in the PowerEdge MX environment integrated with Cisco ACI.

It is recommended to disable LLDP in the distributed virtual switch (vDS) while creating a SmartFabric. After the SmartFabric is established, globally enable LLDP in the APIC fabric or enable under each interface. Enable LLDP in the vCenter vDS. Using LLDP does not impact VM discovery or network traffic between the vDS and ACI.

Alternatively, CDP can be used by selecting CDP on both ACI and vDS. LLDP should remain disabled on the MX SmartFabric.

NOTE: When a storage area network protocol (for example, FCoE) is configured, use CDP as a discovery protocol on ACI and vCenter while LLDP remains disabled on the MX SmartFabric.

SmartFabric Connections to Cisco ACI Leaf Switches

This chapter covers the deployment of a PowerEdge MX SmartFabric solution connected to a Cisco ACI environment.

The validated Cisco ACI environment includes a pair of Nexus C93180YC-EX switches as leafs connected to a single Nexus C9336-PQ switch as the spine using 40 GbE connections. 100 GbE MX9116n FSE switches are connected to the C93180YC-EX leafs. The physical connections are shown in the figure below.

CAUTION: The connection of an MX switch directly to the ACI spine is not supported.

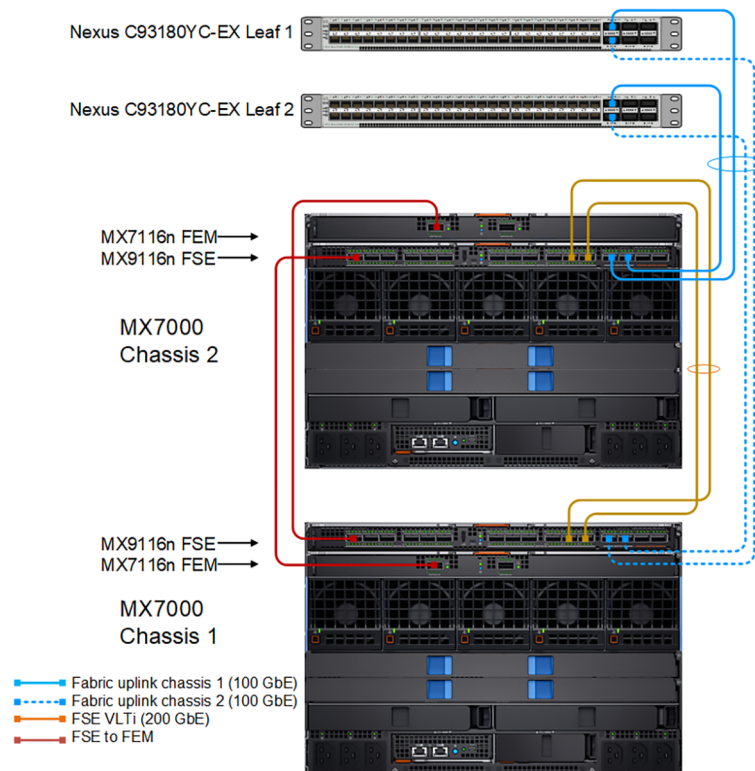


Figure 4. PowerEdge MX connected to Cisco ACI leaf switches

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

Validated environment

In this scenario, two MX7000 chassis are joined to an existing Cisco ACI environment. The MX solution consists of two MX9116n FSEs, two MX7116n Fabric Expander Modules (FEMs), one MX7000 chassis, three MX740c compute sleds, and one MX840c compute sled.

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

All devices in the validated environment are connected as shown in the following figure:

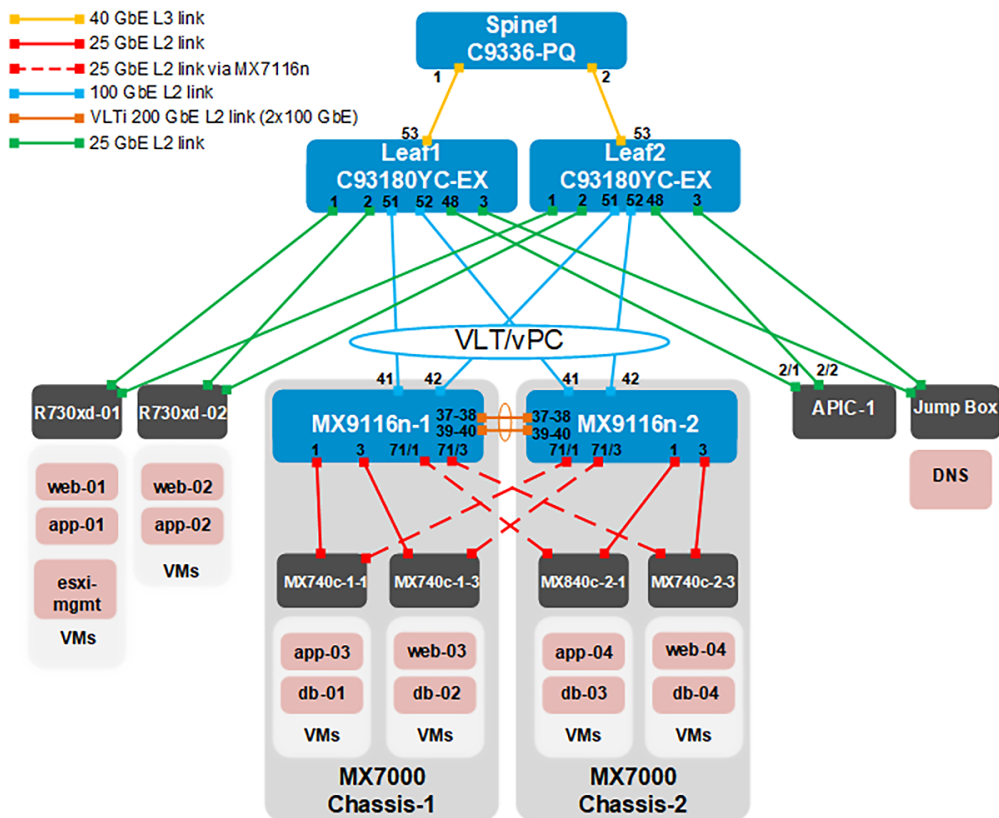


Figure 5. Validated SmartFabric and ACI environment

While a production ACI environment has multiple Application Policy Infrastructure Controllers (APICs), this example uses a single APIC (APIC-1).

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 PowerEdge servers, follow the instructions provided in the [VMware vSphere ESXi 6.7.x Dell EMC PowerEdge Servers Installation Instructions and Important Information Guide](#).

VMs named web, app, and db on the ESXi hosts are running Ubuntu Linux. A third R730xd server is used 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 MGMT on R730xd-03, as shown in the figure above.

The environment uses the six networks that are shown in the following table:

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 vSphere vMotion and vSAN networks are configured in this example, their use is out of the scope of this guide.

VMs in the validated environment use the IP addresses shown in the following table:

Table 5. VM IP addresses

VM name	VLAN name	IP address
MGMT	ESXi_Mgmt	172.16.11.50
web01-web04	web	172.16.14.1-4
app01-app04	app	172.16.15.1-4
db01-db04	db	172.16.16.1-4

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, 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](#), along with the corresponding bridge domain, and application EPG names used in the APIC configuration.


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

The following steps were performed in the Cisco APIC UI for the environment shown in [Figure 4](#).

Create VLAN pool

To create a VLAN pool, perform the following steps:

1. Go to **Fabric > Access Policies > Pools > VLAN**.
2. From the **VLAN** screen, right-click **VLAN** and select **Create VLAN Pool**.
3. In the **Name** field, enter **VLANPool1**.
4. From the **Allocation mode** option, select **Dynamic Allocation**. In this example, Dynamic allocation mode is used to enable APIC to choose VLANs from the pool.

 **NOTE:** 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:

Description:

Allocation Mode: ☒ Dynamic Allocation ☐ Static Allocation

Encap Blocks:

VLAN Range	Allocation Mode	Role
[1611-2000]	Dynamic Allocation	External or On the wire en...

Cancel

Submit

Figure 6. Create VLAN pool

- From the **Encap Blocks** field, click the **Add (+)** icon.
- In the **VLAN Range** fields, enter **1611** and **2000** as shown in the figure below.
- From the **Allocation Mode** field, select **Dynamic Allocation**.
- For the **Role**, select **External or On the wire encapsulations**.

Create Ranges



Specify the Encap Block Range

Type: VLAN

Range: -
Integer Value Integer Value

Allocation Mode: ☒ Dynamic Allocation ☐ Inherit allocMode from parent ☐ Static Allocation

Role: ☒ External or On the wire encapsulations ☐ Internal

Cancel

OK

Figure 7. VLAN range

- Click **OK** and then **Submit**.

Create 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 **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 in Create a VLAN pool).
5. Click **Submit**.

Create Physical Domain

Specify the domain name and the VLAN Pool

Name: physDomain1

Associated Attachable Entity Profile:

select a value

VLAN Pool:

VLANPool1(dynamic)

Security Domains:

Select	Name	Description
<input type="checkbox"/>	vCenter-Credentials	

Cancel

Submit

Figure 8. Create Physical Domain

Create 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 **Attachable Access Entity Profiles** and select **Create Attachable Access Entity Profile**.
3. In the **Name** field, enter **AEP1**.
4. In the **Domains** field, click the **Add (+)** icon.
5. Select **physDomain1** (created in **Create a physical domain**) and then click **Update**.
6. Click **Next** and then **Finish**.

Create Attachable Access Entity Profile

STEP 1 > Profile

1. Profile

2. Association To Interfaces

Specify the name, domains and infrastructure encaps

Name:

Description:

Enable Infrastructure VLAN: ☐

Domains (VMM, Physical or External) To Be Associated To Interfaces:

Domain Profile	Encapsulation
Physical Domain - physDomain1	from:vlan-1611 to:vlan-2000

Figure 9. Create Attachable Access Entity Profile screen

Create port channel policy

To create a port channel policy:

1. Go to **Fabric > Access Policies > Policies > Interface > Port Channel**.
2. Right-click **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**.

Create Port Channel Policy

Specify the Port Channel Policy

Name:

Description:

Alias:

Mode:

Not Applicable for FC PC

Control:

Figure 10. Create Port Channel Policy screen

Create Spanning Tree Interface Policy

The Spanning Tree Interface Policy defines a common configuration that is configured globally and can be deployed to one or multiple interfaces. Because uplinks from the MX network do not have Spanning Tree Protocol (STP) enabled, this policy acts as an extra layer of protection against possible STP issues.

NOTE: Enabling the Spanning Tree Interface Policy is not required; however it is recommended.

To create a Spanning Tree Interface Policy, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Interface > Spanning Tree Interface**.
2. Right-click **Spanning Tree Interface** and select **Create Spanning Tree Interface Policy**.
3. In the **Name** field, enter **ACI-STP**.
4. From the **Interface Controls** option, check the **BPDU filter enabled** and **BPDU Guard enabled** boxes.
5. Click **Submit**.

Create Spanning Tree Interface Policy

Define the STP Interface Policy

Name: ACI-STP

Description: optional

Alias:

Interface controls: ☒ BPDU filter enabled
☒ BPDU Guard enabled

Cancel Submit

Figure 11. Create Spanning Tree Interface Policy screen

Create LLDP Interface policy

NOTE: If using CDP in lieu of LLDP, the LLDP policy is not required. Creating an LLDP policy on ACI is optional and is required only if LLDP is enabled on MX SmartFabric and vCenter vDS.

To create the LLDP Interface policy, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Interface > LLDP Interface**.
2. Right click **LLDP Interface** and select **Create LLDP Interface policy**.
3. In the **Name** field, enter **ACI-LLDP**.
4. Keep all other settings at **default**.
5. Click **Submit**.

Figure 12. Create LLDP Interface policy screen

Create Miscabling Protocol interface policy

Miscabling Protocol, or MCP, detects a loop from connected external devices. MCP disables the interfaces on which ACI receives its own packets. MCP policies are enabled by default on interfaces, however the policies are not enabled globally, and will not take effect until they are enabled globally. MCP is not implemented until it is enabled globally. For best practice, enable the MCP policy globally and on all interfaces.

NOTE: In this example deployment, only once instance of MCP is running.

Enable MCP policy globally

1. Go to **Fabric > Access Policies > Policies > Global > MCP Instance Policy default**.
2. Click **MCP Instance Policy default**.
3. Change **Admin State** to **Enabled**.
4. Click the **Enable MCP PDU per VLAN Controls** box.
5. Enter the unique key for the current ACI fabric in the **Key** field, and then reenter the key in the **Confirm Key** field. The key uniquely identifies the MCP packets within this Fabric.

NOTE: To change the key, go to **Fabric > Access Policies > Policies > Global** and right-click the **MCP Instance Policy default**. In the field provided, enter the **New Key** and **New Key Confirmation** in the fields provided.
6. Click to place a check in the **Loop Protection Action** box to select the **Port disable** option. The MCP policy is enabled globally.

Properties

Name: default

Description: optional

Admin State: Disabled **Enabled**

Controls: ☒ Enable MCP PDU per VLAN

Key:

Confirm Key:

Loop Detect Multiplication Factor: 3

Loop Protection Action: ☒ Port Disable

Initial Delay (sec): 180

Transmission Frequency (sec): 2 (msec): 0

Show Usage Reset Submit

Figure 13. Enable MCP Instance Policy screen

Create MCP Interface policy

To create the Mis-Cabling Protocol Interface policy, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Interface > MCP Interface**.
2. Right-click **MCP Interface** and select **Create Mis-cabling Protocol Interface policy**.
3. In the **Name** field, enter **ACI-MCP**.
NOTE: Keep each of the other settings at the default selection.
4. Click **Submit**.

Create Mis-cabling Protocol Interface Policy

Specify the MCP Interface Policy Properties

Name: ACI-MCP

Description: optional

Admin State: ☐ Disabled ☒ Enabled

Cancel Submit

Figure 14. Create Mis-Cabling Protocol Interface policy screen

Create VPC Interface Policy Group

When interfaces are configured in VPC, an interface policy group needs to be created. A 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 **VPC Interface** and select **Create VPC Interface Policy Group**.
3. In the **Name** field, enter **vPCPolGrp1**.
4. From the **MCP Policy** drop-down, select **ACI-MCP** (enabled in [Create MCP policy](#)).
5. From the **LLDP Policy** drop-down, select **ACI-LLDP** (created in [Create LLDP policy](#)).
6. From the **STP Interface Policy** drop-down, select **ACI-STP** (created in [Create Spanning Tree Interface policy](#)).
7. From the **Attached Entity Profile** drop-down, select **AEP1** (created in [Create Attachable Access Entity profile](#)).
8. From the **Port Channel Policy** drop-down, select **LACPPol1** (created in [Create Port Channel Policy](#)).
9. Click **Submit**.

Create VPC Interface Policy Group

Specify the Policy Group identity

Name: vPCPolGrp1

Description: optional

Link Level Policy: select a value

CDP Policy: select a value

MCP Policy: ACI-MCP

CoPP Policy: select a value

LLDP Policy: ACI-LLDP

STP Interface Policy: ACI-STP

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

Slow Drain Policy: select a value

MACsec Policy: select a value

Attached Entity Profile: AEP1

Port Channel Policy: LACPPol1

Monitoring Policy: select a value

Storm Control Interface Policy: select a value

NetFlow Monitor Policies:

NetFlow IP Filter Type

NetFlow Monitor Policy

Cancel

Submit

Figure 15. Create VPC Interface Policy Group screen

Create Leaf Access Port Policy Group

To create a leaf access port policy group, perform the following steps:

1. Go to **Fabric > Access Policies > Interfaces > Leaf Interfaces > Policy Groups > Leaf Access Port**.
2. Right-click **Leaf Access Port** and select **Create Leaf Access Port Policy Group**.
3. In the **Name** field, enter **LeafHostPortGrp1**.
4. From the **MCP Policy** drop-down, select **ACI-MCP** (enabled in [Create MCP policy](#)).
5. From the **LLDP Policy** drop-down, select **ACI-LLDP** (created in [Create LLDP policy](#)).
6. From the **STP Interface Policy** drop-down, select **ACI-STP** (created in [Create Spanning Tree Interface policy](#)).
7. From the **Attached Entity Profile** drop-down, select **AEP1** (created in [Create attachable access entity profile](#)).
8. Click **Submit**.

Create Leaf Access Port Policy Group

Name: LeafHostPortGrp1

Description: optional

Link Level Policy: select a value

CDP Policy: select a value

MCP Policy: ACI-MCP

CoPP Policy: select a value

LLDP Policy: ACI-LLDP

STP Interface Policy: ACI-STP

Storm Control Interface Policy: select a value

L2 Interface Policy: select a value

Port Security Policy: select a value

Egress Data Plane Policing Policy: select a value

Ingress Data Plane Policing Policy: select a value

Monitoring Policy: select a value

Priority Flow Control Policy: select a value

Fibre Channel Interface Policy: select a value

PoE Interface Policy: select a value

Slow Drain Policy: select a value

MACsec Policy: select a value

802.1x Port Authentication Policy: select a value

DWDM Policy: select a value

Attached Entity Profile: AFP1

Cancel Submit

Figure 16. Create Leaf Access Port Policy Group screen

Create Leaf Interface Profile

Once the vPC interface policy group and leaf access port policy group are created to bundle the interfaces, the interfaces need to be added to the policy groups. To add the interfaces to the policy groups, a 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 **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 17. Create Leaf Interface Profile screen

5. From the **Create Access Port Selector** screen, perform the following steps:
 - a. In the **Name** field, enter **LeafHostSel1**.
 - b. From the **Interface IDs**, enter **1/1-3**.

NOTE: These ports are connected directly to the Dell PowerEdge R730xd servers.
 - c. From the **Interface Policy Group** drop-down, select **LeafHostPortGrp1** (created in [Create Leaf Access Port Policy Group](#)).
 - 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 18. Create Access Port Selector screen

- f. **LeafvPCSel1** contains vPC interfaces **1/51-52**. The ports on the Nexus leaf switches are vPC ports, connected to the Dell EMC Networking MX9116n FSEs. Associate it to **vPCPolGrp1** (created in [Create VPC Interface Policy Group](#)) and click **OK**.
- g. Click **Submit**.

Create Access Port Selector

Specify the selector identity

Name: LeafvPCSel1

Description: optional

Interface IDs: 1/51-52
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: vPCPolGrp1

Figure 19. Create Access Port Selector for vPC interfaces screen

Create VPC Domain policy

To create a VPC domain policy, perform the following steps:

1. Go to **Fabric > Access Policies > Policies > Switch > VPC Domain**.
2. Right-click **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: vPCDom1

Description: optional

Peer Dead Interval: 200

Figure 20. Create VPC Domain Policy screen

Create VPC Explicit Protection Group

To create a VPC explicit protection group, perform the following steps:

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 in [Create VPC domain policy](#)).
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 21. Create VPC Explicit Protection Group screen

Create Leaf Profile

To create a leaf profile, perform the following steps:

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

Create Leaf Profile ? ×

STEP 1 > Profile

Specify the profile Identity

Name: LeafProf1

Description: optional

Leaf Selectors:

Name	Blocks	Policy Group
LeafSel1	101,102	

Previous Cancel Next

Figure 22. Create Leaf Profile screen

- Click **Next**.
- From the **Interface Selector Profiles**, select **LeafIntProf1** (created in [Create leaf interface profile](#)), then click **Finish**. Leaf 101 and 102 display in the **Leaf Profile** shown in the figure below:

Create Leaf Profile ? ×

STEP 2 > Associations

Select the interface/module selector profiles to associate

Interface Selector Profiles:

Select	Name	Description
<input type="checkbox"/>	IntProPorts_...	
<input type="checkbox"/>	IntProPorts_50	
<input type="checkbox"/>	Leaf101Profil...	GUI Interface Selector Generated PortP Profile: Leaf101Profile
<input type="checkbox"/>	Leaf102Profil...	
<input checked="" type="checkbox"/>	LeafIntProf1	

Module Selector Profiles:

Select	Name	Description
--------	------	-------------

Previous Cancel Finish

Figure 23. Choose Interface selector profile screen

Create Tenant

To create a tenant, perform the following steps:

- Go to **Tenants > Add Tenant**.

2. In the **Name** field, enter **Customer-TN1**.
3. Click **Submit**.

Create Tenant

Specify tenant details

Name:

Alias:

Description: optional

Tags:
enter tags separated by comma

GUID:

Provider	GUID

Monitoring Policy:

Security Domains:

Name	Description

Figure 24. Create Tenant screen

Create VRF

Virtual Routing and Forwarding (VRF), or private networks, are a unique Layer 3 forwarding and application policy domain. Private networks contain Bridge domains. To create a VRF, perform the following steps:

1. Go to **Tenants > Customer-TN1 > Networking > VRFs**.
2. Right-click **VRFs** and select **Create VRF**.
3. In the **Name** field, enter **VRF1**.
4. For the **Policy Control Enforcement Preference**, select **Unenforced**.
5. Click to clear 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:

Policy Control Enforcement Direction:

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:

Create A Bridge Domain: ☐

Figure 25. Create VRF screen

Create Bridge Domain

Layer 2 forwarding domain within the Fabric is a bridge domain. A bridge domain is linked to a private network and can have multiple subnets.

NOTE: See Table 3 as needed to complete the steps in this section.

To create bridge domains for each VLAN, perform the following steps:

1. Click **Tenants > Customer-TN1 > Networking > Bridge Domains**.
2. Right-click **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 in [Create VRF](#)), and click **Next**.

Create Bridge Domain

STEP 1 > Main

1. Main

Specify Bridge Domain for the VRF

Name: webBD1

Alias:

Description: optional

Tags:
 enter tags separated by comma

Type: fc regular

Advertise Host Routes: ☐

VRF: VRF1

Forwarding: Optimize

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

IGMP Snoop Policy: select a value

Figure 26. Create Bridge Domain screen

- Next to the **Subnets** listing, click the **Add (+)** icon.
- In the **Gateway IP** field, enter **172.16.14.254/24** for the address and mask for the bridge domain.
NOTE: Leave the remaining values at their default settings.
- Click **OK**, **Next**, and then click **Finish**.
- Repeat the steps in this section as needed for each VLAN.
NOTE: The additional bridge domains created in this example are **appBD1**, **dbBD1**, **ESXiMgmtBD1**, **vMotionBD1**, and **vSANBD1**.

Create Subnet

Specify the Subnet Identity

Gateway IP: 172.16.14.254/24
 address/mask

Treat as virtual IP address: ☐

Make this IP address primary: ☐

Scope: ☒ Private to VRF
 ☐ Advertised Externally
 ☐ Shared between VRFs

Description: optional

Subnet Control: ☐ No Default SVI Gateway
 ☐ Querier IP

L3 Out for Route Profile: select a value

Route Profile: select a value

ND RA Prefix policy: select a value

Cancel

OK

Figure 27. Create Subnet screen

Create Application Profile

To create an application profile, perform the following steps:

- Go to **Tenants > Customer-TN1 > Application Profiles**.
- Right-click **Application Profiles** and select **Create Application Profile**.
- 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 28. Create Application Profile screen

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: See Table 3 for the required network information.

1. Click **Tenants > Customer-TN1 > Application Profiles > ap1 > Application EPGs**.
2. Right-click **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 a separate EPG for each of the remaining bridge domains using the EPG names provided in Table 3: appEPG1, dbEPG1, ESXiMgmtEPG1, vMotionEPG1, and vSANEPG1.

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 29. Create Application EPG

Configure Access Entity Profile with EPGs and VLANs

To configure the access entity profile with EPGs and VLANs, perform the following steps.

NOTE: See [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 in [Create Attachable Access Entity Profile](#)).

Create Attachable Access Entity Profile

STEP 1 > Profile

Specify the name, domains and infrastructure encaps

Domain Profile	Encapsulation
Physical Domain - physDomain1	from:vlan-1611 to:vlan-2000

Figure 30. Create Attachable Access Entity Profile

3. At the bottom of the 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**.

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

Tenant	Application Profile	EPG	Encap	Primary Encap	Mode
Customer-TN1	ap1	webEPG1	vlan-1614		Trunk

Figure 31. Attach AEP to EPGs and bridge domains

Create vCenter domain for Cisco ACI and Virtual Machine Manager (VMM) domain integration

By creating vCenter domain, the user provides a bridge between vCenter and ACI. After creating the domain, the user can see the VMs in the **Created EPGs** area of the Cisco APIC. vCenter Domain also creates Distributed virtual switch (DVS) on APIC to contain the port groups related to EPGs.

NOTE: The name of the Datacenter created in APIC under vCenter domain must be same as the Datacenter name in vCenter mentioned in vCenter configuration overview.

To create a VMware vCenter domain, perform the following steps.

1. Click **Virtual Networking > VMM Domains**.
2. Right-click **VMware** and choose **Create vCenter Domain**.
3. In **Virtual Switch Name** field, enter **VDS-ACI**.
4. From the **Virtual Switch**, select **VMware vSphere Distributed Switch**.
5. From the **Associated Attachable Entity Profile** menu, select **AEP-1**.
6. 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: VDS-ACI

Virtual Switch: VMware vSphere Distributed Switch Cisco AVS Cisco AVE

Associated Attachable Entity Profile: AEP1

Delimiter:

Enable Tag Collection: ☐

Access Mode: Read Only Mode Read Write Mode

Endpoint Retention Time (seconds): 0

VLAN Pool: VLANPool1(dynamic)

Security Domains:

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

vCenter Credentials:

Profile Name	Username	Description
vCenter-Credent...	administrator@vcsa.local	

Cancel Submit

Figure 32. 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 the **Password**, then click **OK**.

Create vCenter Credential



Specify account profile

Name:	<input type="text" value="vCenter-Credential"/>
Description:	<input type="text" value="optional"/>
Username:	<input type="text" value="administrator@vcsa.local"/>
Password:	<input type="password" value="....."/>
Confirm Password:	<input type="password" value="....."/>

Cancel

OK

Figure 33. Create vCenter Credential

8. Next to the **vCenter** listing, click the **Add (+)** icon to add the vCenter Controller.

- In the **Name** field, enter **vCenter**.
- Enter **Host Name (or IP Address)** as per the configuration.
- In the **Datacenter** field, enter **ACI_DC**.
- Associate **vCenter-Credentials** created in the previous step 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.

Add vCenter Controller



Specify controller profile

vCenter Controller

Name:	<input type="text" value="vCenter"/>
Host Name (or IP Address):	<input type="text" value="172.16.11.50"/>
DVS Version:	<input type="text" value="vCenter Default"/>
Stats Collection:	<input checked="" type="radio"/> Disabled <input type="radio"/> Enabled
Datacenter:	<input type="text" value="ACI_DC"/>
Management EPG:	<input type="text" value="select an option"/>
Associated Credential:	<input type="text" value="vCenter-Credential"/>

Cancel

OK

Figure 34. Add vCenter Controller

- Select the **Port Channel Mode**, **vSwitch Policy**, and **NetFlow Exporter Policy** as per configuration. In this example, **LLDP** is selected as **vSwitch Policy**.
- Click **Submit**.

Create vCenter Domain

Specify vCenter domain users and controllers

Profile Name	Username	Description
vCenter-Credential	administrator@vcsa...	

Name	IP	Type	Stats Collection
vCenter	172.16.11.50	vCenter	Disabled

Port Channel Mode:

vSwitch Policy: ☐ CDP ☒ LLDP ☐ Neither

NetFlow Exporter Policy:

Figure 35. Create vCenter Domain after adding vCenter

- Once Submitted, click **vSwitch Policy**, then under the **Port Channel Policy** drop-down menu, select **LACPPol1** (created in [Create Port Channel Policy](#)); and under the **LLDP Policy** drop-down menu, select **ACI-LLDP** (created in [Create LLDP Policy](#)).
- Click **Submit**.

Domain - VDS-ACI

Properties

Port Channel Policy:

LLDP Policy:

CDP Policy:

NetFlow Exporter Policy:

Enhanced Lag Policy

Name	Mode
------	------

Figure 36. vSwitch Policy configurations

Create Contract filter

Contracts are necessary in order to communicate between EPGs. To create a contract filter, perform the following steps.

1. Go to **Tenants > Customer-TN1 > Contracts > Filters**.
2. Right-click **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 **IP** as **EtherType**.
 - c. Leave remaining items at their defaults, click **Update**, and then **Submit**.

Create Filter

Specify the Filter Identity

Name:

Alias:

Description:

Tags: enter tags separated by comma

Entries:

Name	Alias	EtherType	ARP Flag	IP Protocol	Match Only Fragments	Stateful	Source Port / Range	Destination Port / Range	TCP Session Rules	
							From	To	From	To
Allow		IP		unspecified	False	False				

Figure 37. Create Filter

Create Contract

A contract provides a way to control traffic flow within the ACI fabric between EPGs. To create a 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

Specify Identity Of Contract

Name:

Alias:

Scope:

QoS Class:

Target DSCP:

Description:

Tags: enter tags separated by comma

Subjects:

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

Cancel Submit

Figure 38. 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 in [Create Contract filter](#)).

Create Contract Subject

Name:

Alias:

Description:

Target DSCP:

Apply Both Directions: ☒

Reverse Filter Ports: ☒

Filter Chain

L4-L7 Service Graph:

QoS Priority:

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

Cancel OK

Figure 39. Create Contract Subject

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

Apply Contract to VRF

To apply the contract to the VRF, perform the following steps.

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

Policy Operational Faults History

General Subject Labels EPG-Any Labels

Properties

Match Type: AtleastOne

Preferred Group Member: Disabled Enabled

Provided Contracts:

Name	Tenant	Type	QoS Class	Match Type	State
AllowAllContract1	Customer-TN1	Contract	Unspecified	AtleastOne	formed

Consumed Contracts:

Name	Tenant	Type	QoS Class	State
AllowAllContract1	Customer-TN1	Contract	Unspecified	formed

Figure 40. 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 that when a tagged frame (for example, VLAN 1611) enters the ACI fabric, ACI knows that it belongs to the ESXi Management EPG and treats it accordingly.

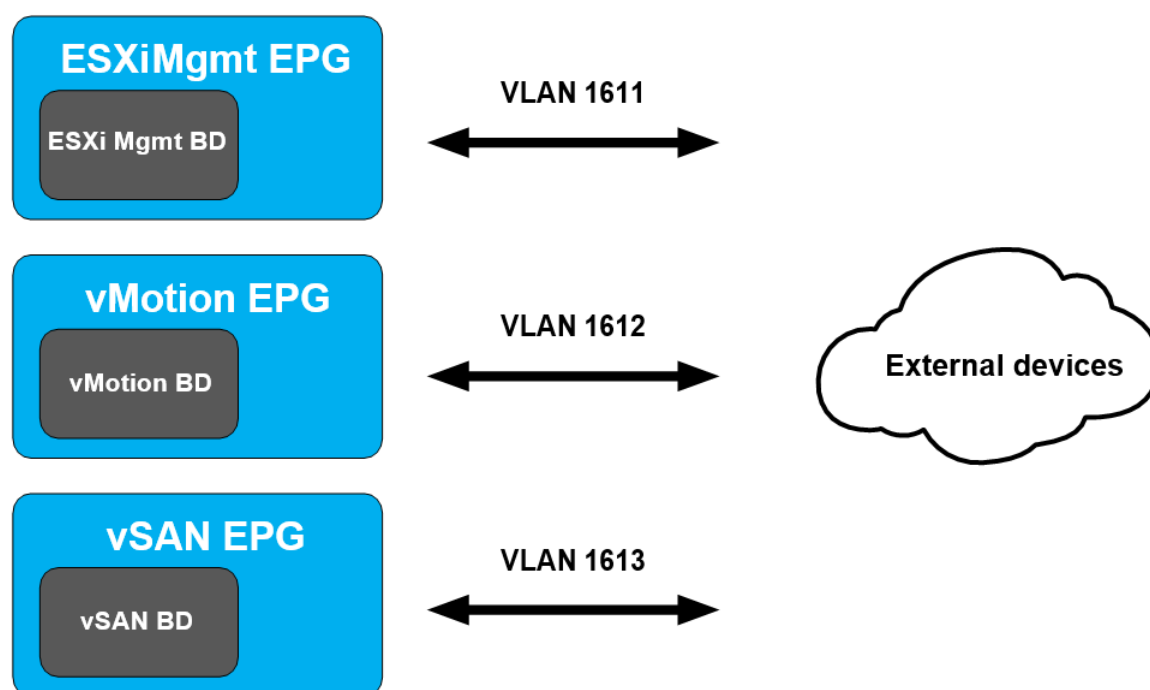


Figure 41. Bridge domains associated with EPGs mapped to external VLANs

SmartFabric deployment

This section covers configuration of PowerEdge MX with Cisco ACI in SmartFabric mode. If using Full Switch mode, see [Appendix A](#) for Full Switch mode example. Once complete, go to the [Server deployment](#) section.

This section provides the details used to deploy the SmartFabric that is used in the example provided in this guide.

Defining VLANs

The VLAN settings used during the SmartFabric deployment for this environment are shown in the following table.

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 the *Define VLANs* section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

NOTE: For information about network type and QoS group settings, see the *Networks and Automated QoS* section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

The configured VLANs for this example are shown in the following figure.

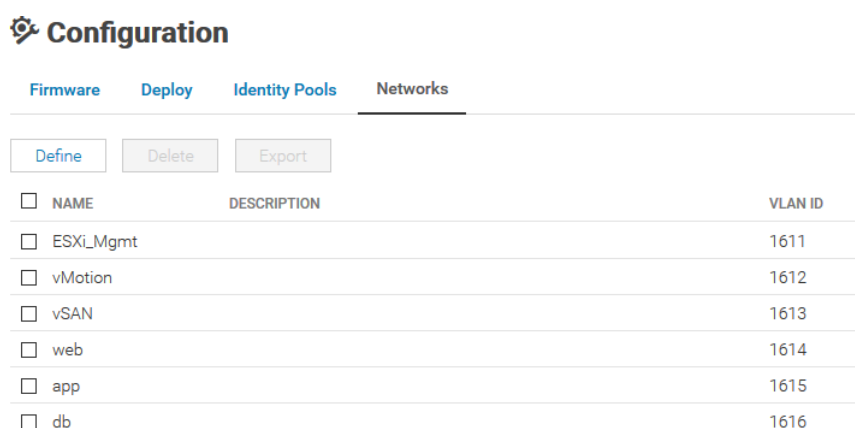


Figure 42. Defined VLANs

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, select the checkbox next to **Include Fabric Management Address in LLDP Messages** on the **Create Fabric** screen, as shown in the following figure, during deployment.

NOTE: Without the Include Fabric Management Address in LLDP Messages feature enabled, the ACI fabric cannot discover the complete network topology.

Create Fabric ? X

Description ✓

Design

Summary

Name

SmartFabric

Description

☒ Include Fabric Management Address in LLDP Messages i

Step 1 of 3

Next

Cancel

Figure 43. Enabling LLDP in SmartFabric

i **NOTE:** LLDP must be enabled on ACI and vCenter. To enable LLDP on ACI, see [Create vCenter Domain for Cisco ACI and Virtual Machine Manager \(VMM\) domain integration](#). To enable LLDP under vDS on vCenter, see the [Enable Link Layer Discovery Protocol on a vSphere Distributed Switch](#) article.

i **NOTE:** When a storage area network protocol (for example, FCoE) is configured, use CDP as a discovery protocol on ACI and vCenter while LLDP remains disabled on the MX SmartFabric.



After creating the SmartFabric (see [Create SmartFabric](#)) and creating the uplink (see [Create Uplink](#)), the VMs display in the APIC under the **Tenants** tab after configuring vCenter. Select the **Tenant** and click **Networking** to view the network topology.

i **NOTE:** If VMs are not present in APIC after creating the SmartFabric with this feature enabled, toggle the server facing switch ports down and back up for the ESXi servers hosting the VMs.

Create SmartFabric

To create a SmartFabric, perform the steps in the *Create the SmartFabric* section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

The figure below shows the new SmartFabric object.

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

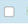

All Devices Chassis Compute I/O Modules Storage Fabric						
Add Fabric		Delete				
HEALTH	FABRIC	DESCRIPTION	SWITCH COUNT	COMPUTE COUNT	UPLINK COUNT	
	SmartFabric		2	4	 0	

Figure 44. SmartFabric after deployment before uplinks are created

Create Uplink

If the port speed or breakout configuration of the ports used in the uplink need to be changed, make those changes before creating the uplink. See the *Configure uplink port speed or breakout* section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#) and make those changes before creating the uplinks.

i **NOTE:** A port breakout was not used in this example.

Ethernet - No Spanning Tree uplink creation

Dell Technologies recommends creating an Ethernet – No Spanning Tree uplink and not the legacy Ethernet uplink. To create the uplink from the switch to the Cisco ACI leafs, see the *Create Ethernet – No Spanning Tree uplink* section in the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

Legacy Ethernet uplink

PowerEdge MX versions prior to 1.20.00 and OS10.5.0.7 did not support the Ethernet – No Spanning Tree uplink type. It is recommended to be running at least PowerEdge MX 1.20.00 and OS10.5.0.7 and to not use the legacy Ethernet uplink. If that is not possible, follow these steps instead of the steps described above.

To create a legacy Ethernet uplink, see the *Create Ethernet uplink* section in the [Dell EMC PowerEdge MX Networking Deployment Guide](#).


If running the legacy Ethernet uplink, perform the following steps before creating the legacy Ethernet uplink in OpenManage Enterprise - Modular.

- 1. Check if any Spanning tree protocol is enabled by running the following command.

```
MX9116n-1# show spanning-tree brief
```

- 2. If any spanning tree protocol is running in this MX environment with Cisco ACI, run the following command to disable spanning tree:

```
MX9116n-1# spanning-tree disable
```

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 the following figure.





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

Figure 45. SmartFabric status after uplink is created

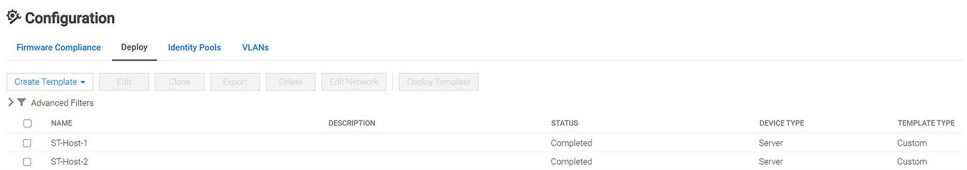
Deploy Server

Create server template

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

-  **NOTE:** For the hardware used in this example, two templates were created.
-  **NOTE:** To create a server template, follow the steps in the Create a server template section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

The templates created for this example are shown in the following figure.



NAME	DESCRIPTION	STATUS	DEVICE TYPE	TEMPLATE TYPE
ST-Host-1		Completed	Server	Custom
ST-Host-2		Completed	Server	Custom

Figure 46. Server templates created

Add VLANs to server templates

After successfully creating server templates, associate each template with the appropriate VLANs. See the *Associate server template with networks* section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#) for the steps necessary.

NOTE: If running in Full Switch mode, VLAN assignment is done manually using the OS10 CLI. See [Appendix A](#) for more information.

Select VLAN 🔍 ✕

Available VLANs:

<input type="checkbox"/>	NAME	VLAN ID
<input type="checkbox"/>	VLAN 1	1

Selected VLANs:

<input type="checkbox"/>	NAME	VLAN ID
<input type="checkbox"/>	ESXi-MGMT	1611
<input type="checkbox"/>	vMotion	1612
<input type="checkbox"/>	vSAN	1613
<input type="checkbox"/>	Web	1614
<input type="checkbox"/>	App	1615
<input type="checkbox"/>	DB	1616

>>

<<

Finish

Cancel

Figure 47. VLANs added to server template

Deploy server templates

To deploy the server templates, complete the steps in the Deploy a server template section of the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

Configure vCenter

The existing ACI environment has two PowerEdge R730xd rack servers connected to the ACI leafs. The rack servers and MX compute sleds are in a vSphere cluster named **Compute**.

After the SmartFabric and uplink is deployed, the rack servers and MX compute sleds can be added to vCenter. For details on how to create a data center, cluster, and virtual machines, and how to add hosts to vCenter, see the [VMware Organizing Your Inventory](#) documentation.

For information on creating vSphere Distributed Switch (vDS) and configuring networking for VDS, see the [Setting up Networking with vSphere Distributed Switches](#) section within the [VMware vSphere Product Documentation](#).

The MX compute sleds can now communicate with the rack servers and the MGMT vCenter instance. The MX compute sleds are joined to the vCenter cluster by an administrator as shown in the following figure.

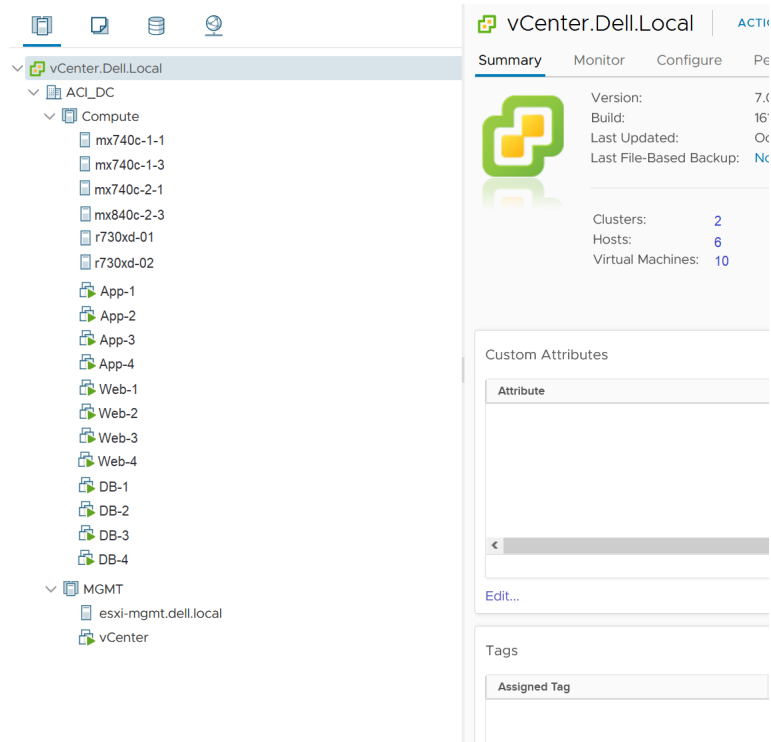


Figure 48. Hosts and VMs used in the validated environment in a single vSphere cluster

A VDS named **VDS-ACI** along with six distributed port groups, one for each VLAN, are used as shown in the following figure.

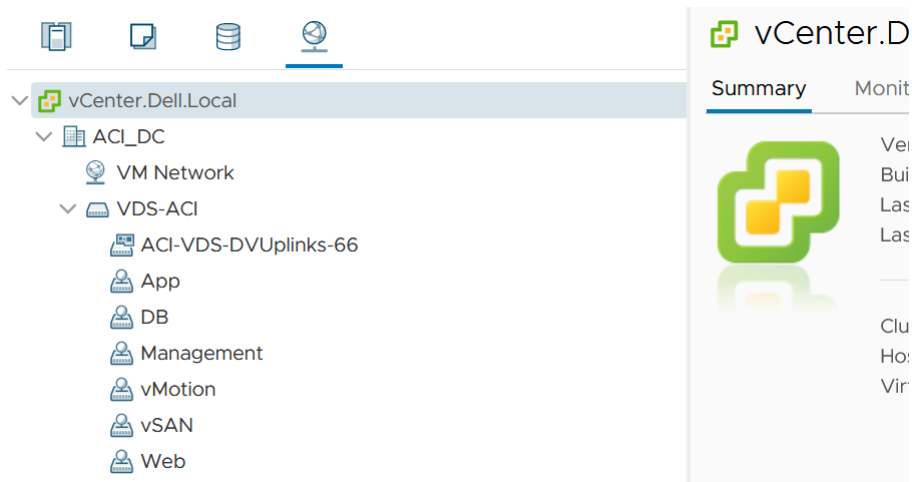


Figure 49. VDS and port groups used in the validated environment

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 the [VMware Validated Design Documentation](#).

App - Edit Settings

General

Advanced

VLAN

Security

Teaming and failover

Traffic shaping

Monitoring

Miscellaneous

Load balancing

Network failure detection

Notify switches

Failback

Failover order ⓘ

Route based on physical NIC load

Link status only

Yes

Yes

↑ ↓

Active uplinks

Uplink 1

Uplink 2

Uplink 3

Uplink 4

Standby uplinks

Unused uplinks

CANCEL

OK

Figure 50. Load balancing method

Detailed vCenter configuration is beyond the scope of this document. For more information about vCenter configuration, see the [VMware vSphere Documentation](#).

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 multichassis 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 the following figure.

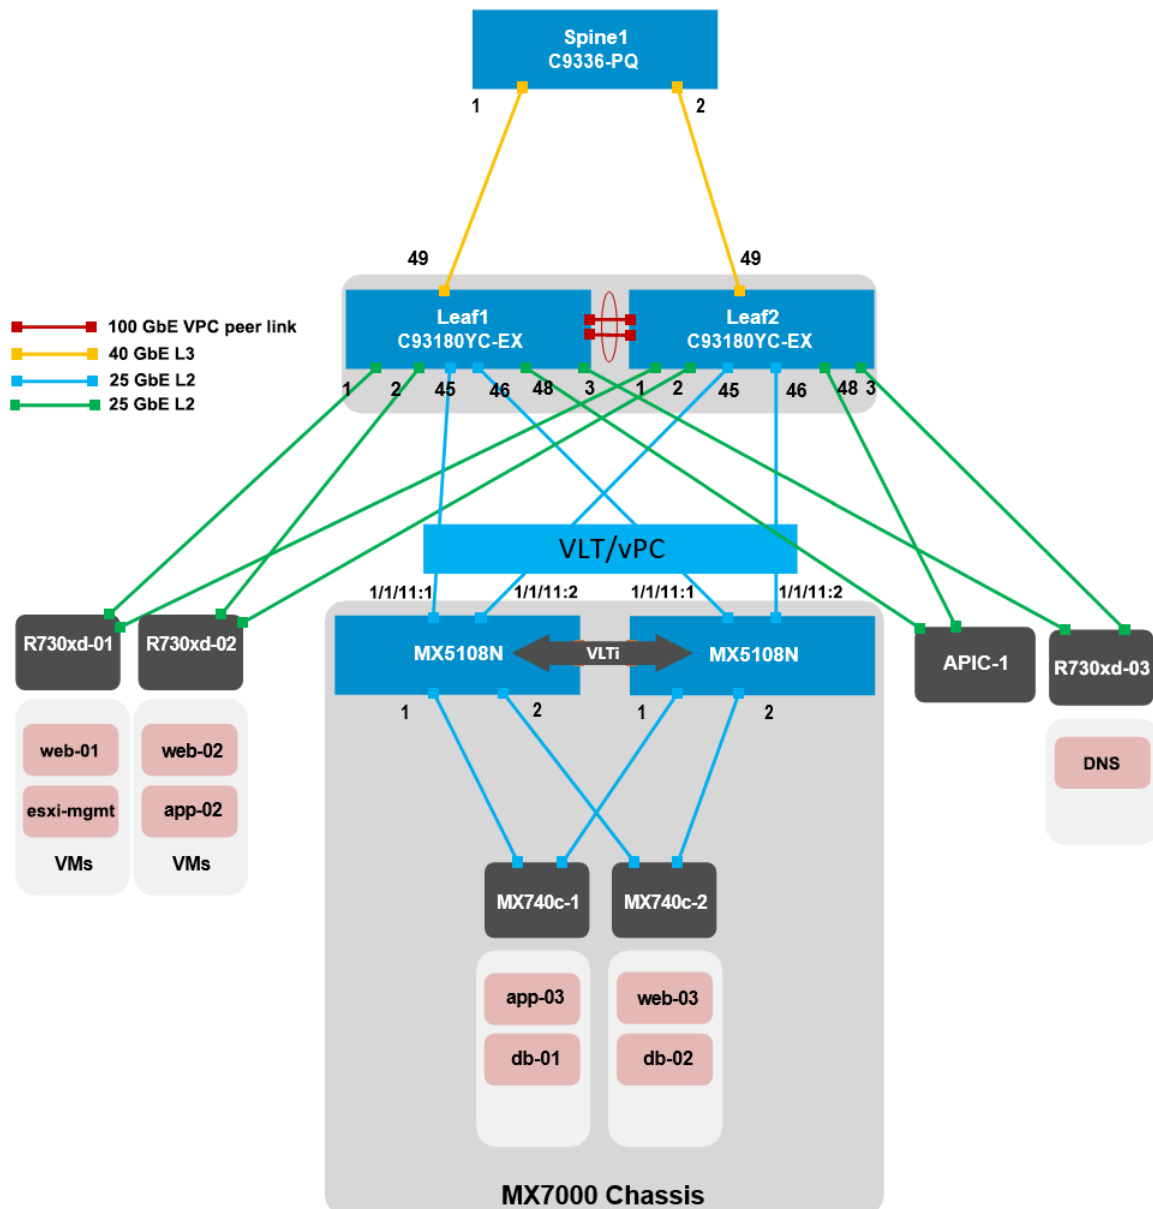


Figure 51. SmartFabric and ACI environment using MX5108n Ethernet switches

The SmartFabric creation and APIC configuration steps are the same as mentioned earlier in this guide (sections [Cisco APIC configuration](#) through [Configure vCenter](#)). Refer to these sections to deploy the ACI infrastructure on the MX7000 Chassis in SmartFabric mode using MX5108n switches.

Validating the Configuration

MX validation using OME-M console

For basic validation of the MX environment and the health of the SmartFabric, see the *SmartFabric Deployment Validation* section in the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

Configuration validation CLI commands for this guide

The CLI command shown in this section is available to help validate the configuration. The command and output shown below is 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**.

For more information about OS10 validation CLI commands, see the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

show lldp neighbors

The `show lldp neighbors` command shows information about devices directly connected to the MX switch. Ports 1/1/1, 1/1/3, 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 1/1/41- 1/1/42 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
```

SmartFabric Services troubleshooting commands

The following commands allow the user to view SmartFabric Services configuration information. These commands can also be used for troubleshooting.

These commands are available in OS10.5.0.1 and later.

 **NOTE:** For more information about SmartFabric Services troubleshooting commands, see the [Dell EMC PowerEdge MX Networking Deployment Guide](#).

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. The configured interface shown below shows the interfaces that connect the MX switches to the ACI system.

```
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
-----
```

show smartfabric networks

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

The output below shows each of the VLANs that were created for this environment.

```
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 HYPERVISOR_MANAGEMENT    PLATINUM    1611
vSAN      STORAGE_DATA_REPLICATION PLATINUM    1613
-----
```

Cisco ACI validation

The following sections show how to validate the ACI portion of this reference architecture.

Verify VPC configuration

Verify the VPC connection from the Cisco ACI fabric to the Dell MX SmartFabric uplink, as shown in the following figure, that it is up and properly configured to the designated VLANs and EPGs. To do this, perform the following steps.

1. In the APIC UI, click **Fabric** > **Inventory** > **Pod name** > **Leaf name** > **Interfaces** > **VPC Interfaces** and scroll down to the applicable port channel VPC policy group as shown in the following figure.

2. Verify that the port channel shows as **lACP-active** and that the **Oper State** shows as **up**.

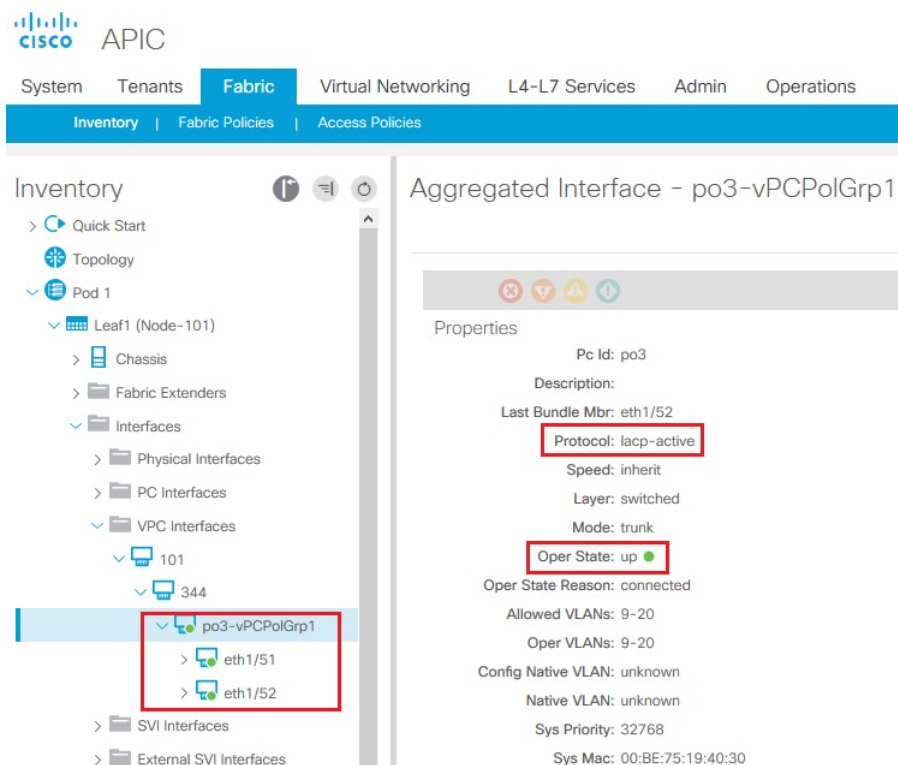


Figure 52. Cisco ACI VPC port channel and interfaces

3. Verify that all of the leaf switch interfaces in the VPC, **eth1/51-52** for example, 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 the following figure.
5. Verify that the port channel includes all required VLANs, and that the EPGs are mapped to the correct VLANs.

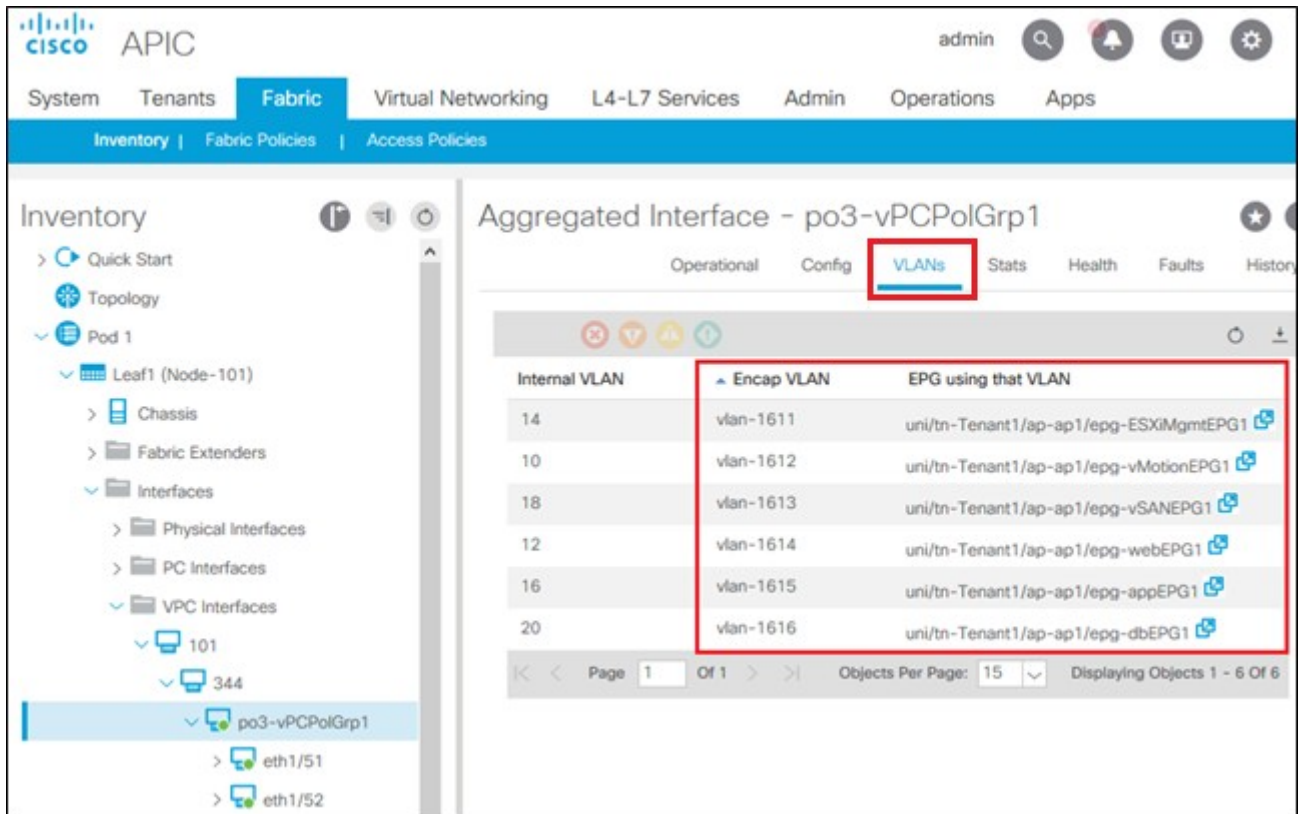


Figure 53. Cisco ACI VPC port channel VLANs and EPGs

Repeat the steps in this section for the remaining leaf switch.

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

1. In the APIC UI, go to **Fabric > Inventory > Pod 1 > Leaf name > Interfaces > Physical Interfaces** as shown in the following figure.

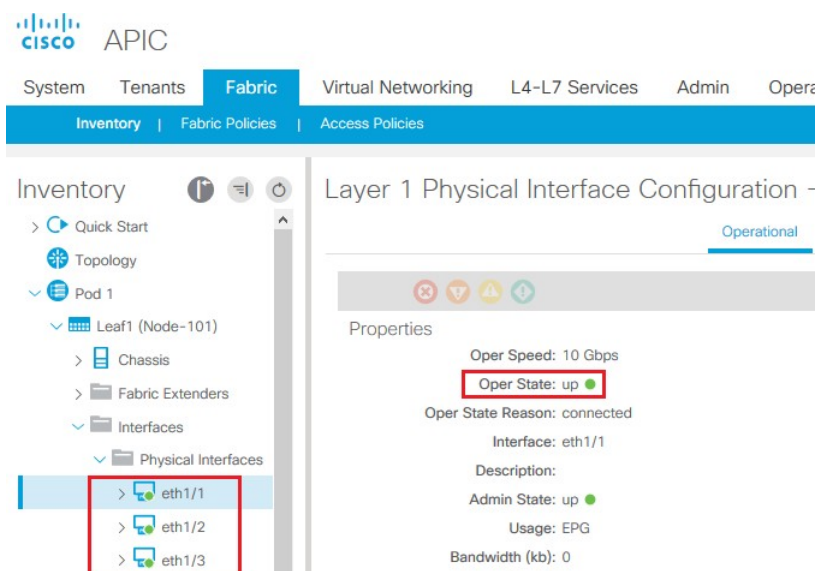


Figure 54. Cisco ACI physical interfaces

2. Verify that the required interfaces, **eth1/1-3** for example, 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 the following figure.
4. Verify that the interface includes all required VLANs and EPGs.

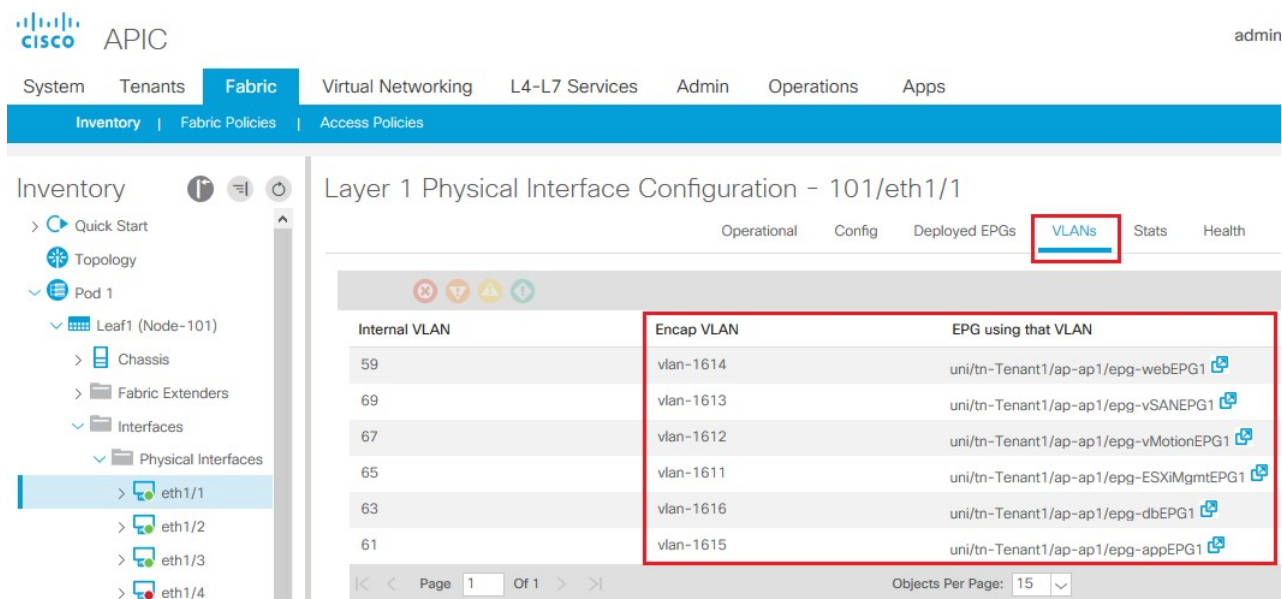


Figure 55. Cisco ACI interface VLANs and EPGs

Repeat the steps in this section for the remaining interfaces as needed, and for the remaining leaf switch.

Verify ACI endpoint learning

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

1. In the APIC UI, go to **Tenants > Tenant name > Application Profiles > Application Profile name > Application EPGs** and select an **EPG**. In this deployment example, **Web EPG** is selected.
2. Click the **Operational** tab in the navigation window as shown in the following figure.
3. Review the listing of the learned endpoints for the selected EPG along with the learning source, IP address, and interface.

NOTE: When a storage area network protocol (for example, FCoE) is configured, use CDP as a discovery protocol on ACI and vCenter while LLDP remains disabled on the MX SmartFabric.

In this deployment example, the PowerEdge MX environment Cisco ACI, and VMware vSphere devices discover neighbors through LLDP. Configure LLDP on all of the devices. When configured correctly, the output displays as shown in the figure below:

EPG - Web-EPG1

Summary Policy Operational Stats Health Faults History						
Hosts VMs						
Client End-Points Configured Access Policies Contracts Controller End-Points Learned End-Points						
100						
End Point	MAC	IP	Learning Source	Hosting Server	Reporting Controller Name	Interface
Web-1	00:50:56:A9:E0:3B	172.16.14.10	learned vmm	100.67.105.51	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
Web-2	00:50:56:A9:A5:C8	172.16.14.11	learned vmm	100.67.105.52	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
EP-00:50:56:6F:56:EE	00:50:56:6F:56:EE	172.16.14.50	learned vmm	100.67.105.50	---	Pod-1/Node-101/eth1/2 (learned,vmm) Pod-1/Node-102/eth1/2 (vmm)
EP-00:50:56:6C:AD:5A	00:50:56:6C:AD:5A	172.16.14.52	learned vmm	100.67.105.52	---	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
EP-00:50:56:61:16:34	00:50:56:61:16:34	172.16.14.51	learned vmm	100.67.105.51	---	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...

Figure 56. Cisco ACI endpoints in Web EPG

Repeat the steps in this section for the remaining Application EPGs.

NOTE: If the output does not display as shown in the figure above, go to the [Troubleshooting LLDP](#) section.

Verify ACI VMM domain integration

To verify ACI vCenter domain integration, perform the following steps.

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

Domain - VDS-ACI

Policy

Operational

Associated EPGs

General

History

Faults

Properties

Name: VDS-ACI

Controllers:

Name

State

Model

Serial

Revision

Hypervisors

Virtual Machines

vCenter-Server

Online

VMware vCenter Server 6.7.0 bu...

4e00e2ef-3d...

6.7.0

6

16

Page 1

Of 1

Objects Per Page: 15

Displaying Objects 1 - 1 Of 1

Figure 57. vCenter server

- 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**.

This shows the Datacenter, Management EPG, and Associated Credential details.

Figure 59. vCenter server details

Verifying 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. The configuration used in this guide is to allow unrestricted inter-EPG communication, as detailed in the [Create Contract filter](#), [Create Contract](#), and [Apply Contract to VRF](#) sections.

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

The following figure shows the VM named app-01, in a rack server, successfully pinging the VMs named web-03 and db-04, which are 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 60. Verifying connectivity between VMs

Troubleshooting

Troubleshooting LLDP

Devices use Link Layer Discovery Protocol (LLDP) to discover and establish a neighbor relationship with other devices in the same network that has LLDP enabled. LLDP is a Layer 2 protocol which uses the Logical Link Control (LLC) services to receive and transmit information to and from other LLDP agents in the network.

In the deployment example with a PowerEdge MX environment, Cisco ACI, and VMware vSphere, all the devices discover neighbors through LLDP.

If LLDP is not configured correctly, incomplete information is visible when viewing the Client End-Points for each EPG in the ACI UI.

To access the Client End-Points screen, perform the following steps:

1. Log in to the Cisco APIC and go to **Tenants > Customer-TN1 > Application Profiles > ap1 > Application EPGs**
2. Select any of the EPGs listed, then select **Operational**.

NOTE: In this example, the Web EPG is selected.

If the PowerEdge MX compute sled hosts in the **Client End-Points** tab show learned and not learned VMM as shown in the figure below, go to step 3.

EPG - Web-EPG1

End Point	MAC	IP	Learning Source	Hosting Server	Reporting Controller Name	Interface
Web-1	00:50:56:A9:E0:3B	172.16.14.10	learned vmm	100.67.105.51	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (learned)
Web-2	00:50:56:A9:A5:C8	172.16.14.11	learned vmm	100.67.105.52	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (learned)
EP-00:50:56:6F:56:EE	00:50:56:6F:56:EE	172.16.14.50	learned vmm	100.67.105.50	---	Pod-1/Node-101/eth1/2 (learned,vmm) Pod-1/Node-102/eth1/2 (vmm)
EP-00:50:56:6C:AD:5A	00:50:56:6C:AD:5A	172.16.14.52	learned	---	---	Pod-1/Node-101-102/vPCPolGrp1 (learned)
EP-00:50:56:61:16:34	00:50:56:61:16:34	172.16.14.51	learned	---	---	Pod-1/Node-101-102/vPCPolGrp1 (learned)

Figure 61. Client End-Points Output on APIC when LLDP is not correctly configured

3. To verify that the LLDP configuration is correct across the different components of the solution, ensure that the **Fabric Management Address** in the **LLDP Message options** listing is enabled on the MX switches.
4. Verify the LLDP on the Cisco ACI.
5. Verify the LLDP on VMware vSphere Distributed Switch (vDS) on vCenter.

Verify Fabric Management Address in LLDP Message Option is Enabled

SmartFabric mode

To verify that Fabric Management Address is in LLDP Protocol Data Units (PDUs) in SmartFabric mode, perform the following steps on the OME-M Console.

1. Open the OME-M console and select **Devices > Fabric**.

2. Select the SmartFabric to verify, then select **Edit**.
3. Ensure that the checkbox next to the **Include Fabric Management Address in LLDP Messages** option is selected, then click **Finish**.

Edit Fabric ? X

Description	Name	Description
	SmartFabric	

☒ Include Fabric Management Address in LLDP Messages i

Step 1 of 1

Finish
Cancel

Figure 62. Enable LLDP in SmartFabric Mode

Full Switch mode

To verify that Fabric Management Address is in LLDP Protocol Data Units (PDUs) in Full Switch mode, run the `show running-configuration` command on the OS10 CLI. Ensure that the output includes the `lldp management-addr-tlv ipv4 virtual-ip` command as one of the listed commands. If it is not configured, perform the below steps.

i **NOTE:** This must be configured on both switches within the fabric.

1. SSH to the OS10 CLI on the switch.
2. Run the following command on all the IOMs within the fabric:

```
MX9116n-1# configure terminal
```

3. If the switch running the configuration shows that the LLDP is not enabled, run the `lldp enable` command:

```
MX9116n-1<config># lldp enable
```

4. Run the `lldp management-addr-tlv ipv4 virtual-ip` command to include the Fabric Management Address in LLDP messages:

```
MX9116n-1<config># lldp management-addr-tlv ipv4 virtual-ip
```

5. Exit and save the configuration by running the following commands:


```
MX9116n-1<config># exit
MX9116n-1<config># write memory
```

Verify LLDP on Cisco ACI

Confirm that the LLDP Interface Policy has been created as show in the [Create LLDP Interface policy](#) section.

If the vCenter domain has not been created, create the vCenter domain with LLDP as vSwitch Policy as mentioned in the [Create vCenter domain for Cisco ACI and Virtual Machine Manager \(VMM\) domain integration](#) section. If the vCenter domain is already created, perform the following steps to verify if the LLDP policy has been created and applied.

1. Log in to the Cisco APIC, select the **Virtual Networking** tab, then select **Inventory**.
2. Expand the **VMM Domains** and **VMware** options.
3. Select the **Distributed Switch** that was created.
4. From the **Policy** tab, select the **vSwitch Policy** option.
5. From the **LLDP policy** drop-down menu, ensure that **LLDP policy** is selected.

 **NOTE:** The ACI-LLDP policy that was created in the **Create LLDP policy**, is shown in the figure below.

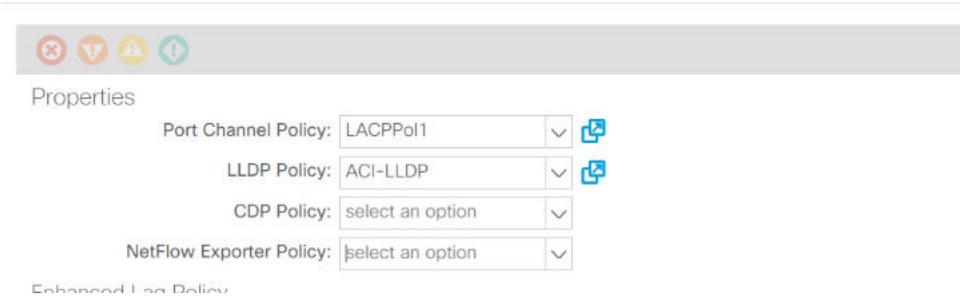


Figure 63. LLDP Policy Assigned to vSwitch Policy in Cisco ACI

Verify LLDP on VMware vSphere Distributed Switch in VMware vCenter

- 1. Log in to VMware vCenter and select the **Networking** tab.
- 2. Expand the **vSphere Data Center** section and right-click the **Distributed Switch** listing.
- 3. From the **Settings** option, click **Edit Settings** and select **Advanced**.
- 4. Within the **Link Layer policy** drop-down menu, select **Link Layer Discovery Protocol**, and from the **Operation** listing, select **Both** as shown in figure below.

VDS-ACI - Edit Settings

General

Advanced

MTU (Bytes)

9000

Multicast filtering mode

Basic

Discovery protocol

Type

Link Layer Discovery Protocol

Operation

Both

Administrator contact

Name

Other details

CANCEL

OK

Figure 64. Enable LLDP on VMware vSphere Distributed Switch under vCenter

Once LLDP is enabled on all devices, and the Fabric Management Address is in the LLDP messages throughout the PowerEdge MX environment, hosts and VMs are learned through the VMM domain integration. This can be verified within the **Client End-Points** section within the EPGs in the Cisco APIC as shown in figure below.

Hosts

VMs

Client End-Points

Configured Access Policies

Contracts

Controller End-Points

Learned End-Points

End Point	MAC	IP	Learning Source	Hosting Server	Reporting Controller Name	Interface
Web-1	00:50:56:A9:E0:3B	172.16.14.10	learned vmm	100.67.105.51	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
Web-2	00:50:56:A9:A5:C8	172.16.14.11	learned vmm	100.67.105.52	vCenter	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
EP-00:50:56:6F:56:EE	00:50:56:6F:56:EE	172.16.14.50	learned vmm	100.67.105.50	---	Pod-1/Node-101/eth1/2 (learned,vmm) Pod-1/Node-102/eth1/2 (vmm)
EP-00:50:56:6C:AD:5A	00:50:56:6C:AD:5A	172.16.14.52	learned vmm	100.67.105.52	---	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...
EP-00:50:56:61:16:34	00:50:56:61:16:34	172.16.14.51	learned vmm	100.67.105.51	---	100.67.105.221 (vmm) Pod-1/Node-101-102/vPCPolGrp1 (lear...

Figure 65. Client End-Points Output on APIC when LLDP Option is Enabled on PowerEdge MX

Full Switch Mode Example

The Dell EMC Networking MX9116n Fabric Switching Engine (FSE) operates in one of two modes:

- Full Switch mode (Default) – All switch-specific SmartFabric OS10 capabilities are available.
- SmartFabric mode – Switches operate as a Layer 2 I/O aggregation fabric and are managed through the Open Manage Enterprise-Modular (OME-M) console.

This section describes the Full Switch mode configuration for the deployment example in this guide.

Full Switch mode

In Full Switch mode, all SmartFabric OS10 features and functions supported by the hardware are available to the user. The MX switch will operate the same way as any other SmartFabric OS10 switch. Configuration is primarily done through the CLI; however, the following items can be configured or managed using the OME-M user interface:

- Initial switch deployment: Configure hostname, password, SNMP, NTP, and so on
- Set ports administratively up or down, configure MTU
- Monitor health, logs, alerts, and events
- Update the SmartFabric OS10 software
- View physical topology
- Switch power management

Full Switch mode is typically used when a wanted feature or function is not available when operating in SmartFabric Services mode. For more information about Dell EMC SmartFabric OS10 operations, see the *Dell EMC SmartFabric OS10 User Guide* on the [Documentation](#) page of the support site.

Ethernet switch configuration

This section outlines example configuration commands issued to the Dell EMC Networking MX9116n switches. The switches start at their factory default setting. See the *Dell EMC SmartFabric OS10 User Guide* on the [Documentation](#) page of the support site for more information.

To configure the MX9116n switches, perform the following steps.

1. Set the switch hostname and management IP address. Configure Spanning Tree and run the command to Include Fabric Management Address in LLDP messages.
2. Configure VLT between the switches.
3. Configure VLANs.
4. Configure port channels to connect to the upstream switches.
5. Configure the ports connected to the upstream switches and compute the sleds.
6. Configure the uplink-state group.

Use the following commands to configure the hostname, management interface, default gateway, spanning tree and required LLDP settings.

NOTE: LLDP must be enabled on ACI and vCenter. To enable LLDP on ACI, see the [Create vCenter domain for Cisco ACI and VMM domain integration](#) section. To enable LLDP under vDS on vCenter, see the [Enable Link Layer Discovery Protocol on a vSphere Distributed Switch](#) article.

MX9116n-A1	MX9116n-A2
<code>configure terminal</code>	<code>configure terminal</code>

MX9116n-A1	MX9116n-A2
<pre>hostname MX9116n-A1 interface mgmt 1/1/1 no ip address dhcp no shutdown ip address 100.67.105.221/24 management route 0.0.0.0/0 100.67.105.254 spanning-tree disable lldp management-addr-tlv ipv4 virtual-ip</pre>	<pre>hostname MX9116n-A2 interface mgmt 1/1/1 no ip address dhcp no shutdown ip address 100.67.105.222/24 management route 0.0.0.0/0 100.67.105.254 spanning-tree disable lldp management-addr-tlv ipv4 virtual-ip</pre>

Configure VLT between switches using the following commands. VLT configuration involves setting a discovery interface range and discovering the VLT peer in the VLT.

NOTE: The default speed of port groups 1/1/10 through 1/1/12 is 100g-2x (QSFP28-DD 200 GbE connection) which means all the ports from 1/1/35 through 1/1/40 are all 100 GbE ports. Only port groups 1/1/11 and 1/1/12 (ports 1/1/37 through 1/1/40) are used in this example.

MX9116n-A1	MX9116n-A2
<pre>interface range ethernet1/1/37-1/1/40 description VLTi no shutdown no switchport vlt-domain 1 backup destination 100.67.105.222 discovery-interface ethernet 1/1/37-1/1/40</pre>	<pre>interface range ethernet1/1/37-1/1/40 description VLTi no shutdown no switchport vlt-domain 1 backup destination 100.67.105.221 discovery-interface ethernet 1/1/37-1/1/40</pre>

Configure the required VLANs on each switch as described in the [Define VLANs](#) section, and also define the VRRP (Virtual Routing Redundancy Protocol) groups under each VLAN to provide gateway access.

MX9116n-A1	MX9116n-A2
<pre>interface vlan1611 description ESXi-MGMT no shutdown ip address 172.16.11.252/24 vrrp-group 11 virtual-address 172.16.11.254 interface vlan1612 description vMotion no shutdown ip address 172.16.12.252/24 vrrp-group 12 virtual-address 172.16.12.254 interface vlan1613 description vSAN no shutdown ip address 172.16.13.252/24 vrrp-group 13 virtual-address 172.16.13.254 interface vlan1614 description web no shutdown ip address 172.16.14.252/24</pre>	<pre>interface vlan1611 description ESXi-MGMT no shutdown ip address 172.16.11.253/24 vrrp-group 11 virtual-address 172.16.11.254 interface vlan1612 description vMotion no shutdown ip address 172.16.12.253/24 vrrp-group 12 virtual-address 172.16.12.254 interface vlan1613 description vSAN no shutdown ip address 172.16.13.253/24 vrrp-group 13 virtual-address 172.16.13.254 interface vlan1614 description web no shutdown ip address 172.16.14.253/24</pre>

MX9116n-A1	MX9116n-A2
<pre> vrp-group 14 virtual-address 172.16.14.254 interface vlan1615 description app no shutdown ip address 172.16.15.252/24 vrp-group 15 virtual-address 172.16.15.254 interface vlan1616 description db no shutdown ip address 172.16.16.252/24 vrp-group 16 virtual-address 172.16.16.254 </pre>	<pre> vrp-group 14 virtual-address 172.16.14.254 interface vlan1615 description app no shutdown ip address 172.16.15.253/24 vrp-group 15 virtual-address 172.16.15.254 interface vlan1616 description db no shutdown ip address 172.16.16.253/24 vrp-group 16 virtual-address 172.16.16.254 </pre>

Configure the port channel that connects to the upstream switches. The LACP protocol is used to create the dynamic LAG. Trunk ports allow tagged VLANs to traverse the trunk link. In this example, the trunk is configured to allow all VLANs. Configure upstream ports and ports connected to compute sleds.

MX9116n-A1	MX9116n-A2
<pre> interface port-channel1 description "LACP to ACI" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 vlt-port-channel 1 interface ethernet1/1/41 description "Connection to ACI switches" no shutdown no switchport channel-group 1 mode active flowcontrol receive off interface ethernet1/1/42 description "Connection to ACI switches" no shutdown no switchport channel-group 1 mode active flowcontrol receive off interface ethernet 1/1/1 description "Chassis1-Sled1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 flowcontrol receive off interface ethernet 1/1/3 description "Chassis1-Sled2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 flowcontrol receive off </pre>	<pre> interface port-channel1 description "LACP to ACI" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 vlt-port-channel 1 interface ethernet1/1/41 description "Connection to ACI switches" no shutdown no switchport channel-group 1 mode active flowcontrol receive off interface ethernet1/1/42 description "Connection to ACI switches" no shutdown no switchport channel-group 1 mode active flowcontrol receive off interface ethernet 1/1/1 description "Chassis2-Sled1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 flowcontrol receive off interface ethernet 1/1/3 description "Chassis2-Sled2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 1611-1616 flowcontrol receive off </pre>

Configure and enable the uplink-state group for upstream and downstream connections. When the upstream links go down, the switch disables the downstream links.

MX9116n-A1	MX9116n-A2
<pre>uplink-state-group 1 enable downstream ethernet1/1/1-1/1/3 upstream port-channel1</pre>	<pre>uplink-state-group 1 enable downstream ethernet1/1/1-1/1/3 upstream port-channel1</pre>

Hardware and Software Versions

Dell EMC switches

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 Deployment Guide](#).

The following table lists the switches used and the role of each switch in this example. The following sections describe each Dell EMC switch in greater detail.

Table 7. Dell EMC switches overview

Qty	Model number	Role
1	Dell EMC PowerSwitch S3048-ON	Supports out-of-band (OOB) management traffic for all examples
2	Dell EMC Networking MX9116n FSE	Used as IO modules connected to the Nexus leaf switches
2	Dell EMC Networking MX5108n	Used as IO modules connected to the Nexus leaf switches

Dell PowerSwitch S3048-ON

The Dell PowerSwitch S3048-ON is a 1U switch with forty-eight 1 GbE BASE-T ports and four 10 GbE SFP+ ports.




Figure 66. PowerSwitch S3048-ON

Dell EMC Networking MX9116n 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 25 GbE ports, the MX9116n FSE also provides two 100 GbE QSFP28 ports, two 100 GbE QSFP28 unified ports, and twelve 2x 100 GbE QSFP-28 Double Density (DD) ports.



Figure 67. Dell EMC Networking MX9116n Fabric Switching Engine

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 40 GbE QSFP+ port, two 100 GbE QSFP28 ports, and four 10 GbE RJ45 BASE-T ports.

These ports provide a combination of network uplink, VLT interconnect (VLTi), or FCoE connectivity.

NOTE: The MX5108n supports FCoE Initialization Protocol (FIP) Snooping Bridge (FSB) mode but does not support NPG or direct-attach FC capabilities.

The Dell EMC PowerEdge MX7000 supports up to four MX5108n Ethernet switches in Fabric A or Fabric B, or both.



Figure 68. Dell EMC Networking MX5108n Ethernet switch

Cisco switches

Cisco Nexus C93180YC-EX

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

Cisco Nexus C9336-PQ

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

Validated components and software versions

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

Dell EMC PowerSwitch

Table 8. Dell EMC PowerSwitch and OS version

Qty	Item	OS Version
1	Dell EMC PowerSwitch S3048-ON OOB management switch	10.5.0.7

Dell EMC PowerEdge MX7000 chassis and components

Table 9. Dell EMC PowerEdge MX7000 chassis and components

Qty	Item	Version
2	Dell EMC PowerEdge MX7000 chassis	1.20.10
3	Dell EMC PowerEdge MX740c sled	-
1	Dell EMC PowerEdge MX840c sled	-
4	Dell EMC PowerEdge M9002m modules (2 per chassis)	1.20.10
2	Dell EMC Networking MX9116n FSE (1 per chassis)	10.5.1.6
2	Dell EMC Networking MX7116n FEM (1 per chassis)	N/A
2	Dell EMC Networking MX5108n Ethernet switch	10.5.1.6

MX740c sled

Table 10. MX740c sled details

Qty per sled	Item	Version
2	Intel(R) Xeon(R) Silver 4114 CPU @ 2.20 GHz	-
12	16 GB DDR4 DIMMs (192 GB total)	-
1	Boot Optimized Storage Solution (BOSS) S1 Controller w/ 1 x 120 GB SATA SSD	2.6.13.3011
1	PERC H730P MX	25.5.5.0005
2	600 GB SAS HDD	-
1	Intel(R) Ethernet 2x 25 GbE XXV710 mezzanine card or QLogic 2x 25 GbE QL41232HMKR mezzanine card	19.5.12 (Intel) or 15.15.11 (QLogic)
-	BIOS	2.8.1
-	iDRAC with Lifecycle Controller	4.22.00.00
-	VMware ESXi (Dell EMC Customized)	6.7.0

MX840c sled

Table 11. MX840c sled details

Qty per sled	Item	Version
2	Intel(R) Xeon(R) Gold 5220 CPU @ 2.20 GHz	-
12	16 GB DDR4 DIMMs (192 GB total)	-
1	Boot Optimized Storage Solution (BOSS) S1 Controller w/ 1 x 120 GB SATA SSD	2.6.13.3011
1	PERC H730P MX	25.5.5.0005
2	600 GB SAS HDD	-
1	Intel(R) Ethernet 2x 25 GbE XXV710 mezzanine card or QLogic 2x 25 GbE QL41232HMKR mezzanine card	19.5.12 (Intel) or 15.15.11 (QLogic)
-	BIOS	2.8.1
-	iDRAC with Lifecycle Controller	4.22.00.00
-	VMware ESXi (Dell EMC Customized)	6.7.0

VMware components

Table 12. VMware components

Item	Version
VMware vCenter Server Appliance (VCSA)	6.7.0 (8169921)
VMware-VMvisor-Installer (ESXi)	6.7.0 (9484548)
VMware Remote Console (VMRC)	10.0.3 (9300449)
VMware Virtual Distributed Switch (vDS)	6.6.0

Cisco ACI components

Table 13. 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)

Documentation and Support

Dell Technologies documentation

The following Dell Technologies documentation provides additional and relevant information. Access to these documents may depend on your log in credentials. If you do not have access to a document, contact your Dell Technologies representative.

- [Dell EMC Networking Guides](#)
- [Dell EMC PowerEdge MX IO Guide](#)
- [Interactive Demo: OpenManage Enterprise Modular for MX solution management](#)
- [Dell EMC PowerEdge Networking Deployment Guide](#)
- [Dell EMC PowerEdge MX SmartFabric Deployment Video](#)
- [Dell EMC PowerEdge MX SmartFabric Deployment with Cisco ACI Video](#)
- [Dell EMC PowerEdge MX VMware ESXi with SmartFabric Services Deployment Guide](#)
- [SmartFabric Services for PowerEdge MX Port-Group Configuration Errors Video](#)
- [SmartFabric Services for PowerEdge MX Port-Group Configuration Video](#)
- [Manuals and documents for Dell EMC SmartFabric OS10](#)
- [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](#)
- [Manuals and Documents for Dell EMC PowerEdge MX740c](#)
- [Manuals and Documents for Dell EMC PowerEdge MX840c](#)
- [Manuals and Documents for Dell EMC PowerEdge R730xd](#)

OME-M and OS10 compatibility and documentation

This section includes the compatibility matrix of OME-M and OS10 and provides links to OME-M and OS10 user guides and release notes for all versions.

OME-M and OS10 compatibility

OME-M version	OS10 version
1.10.00	10.5.0.1
1.10.20	10.5.0.5
1.20.00	10.5.0.7, 10.5.9
1.20.10	10.5.1.6, 10.5.1.7, 10.5.1.9
1.30.00	10.5.2.3 (factory only), 10.5.2.4, 10.5.2.6
1.30.10	10.5.2.6
1.40.00, 1.40.10, 1.40.20	10.5.3.1
2.00.00	10.5.4.1

OME-M and OS10 documentation

The following OME-M documents are available on the [Documentation tab of the PowerEdge MX7000 support site](#).

- Dell OpenManage Enterprise-Modular Edition for PowerEdge MX7000 Chassis User's Guide
- Dell OpenManage Enterprise-Modular Edition for PowerEdge MX7000 Chassis Release Notes

The following OS10 documents are available on the [Documentation tab of the SmartFabric OS10 Software support site](#).

- Dell SmartFabric OS10 User Guide
- SmartFabric OS10 Release Notes for PowerEdge MX

Support and feedback

For technical support, go to <https://www.dell.com/support> or call (USA) 1-800-945-3355.

Dell Technologies and the authors of this document welcome your feedback on the solution and the solution documentation. Contact the Dell Technologies Solutions team by [email](#).