

# Dell EMC PowerEdge MX Series Fibre Channel Storage Network Deployment with Ethernet IOMs

## Abstract

This document provides the deployment steps for configuring PowerEdge MX-series networking switches in SmartFabric mode to attach external FC storage. Examples include FCoE FIP Snooping Bridge (FSB), NPIV Proxy Gateway (NPG), and Direct Connect of FC storage devices.

April 2019

## Revisions

Date	Description
April 2019	Added appendix A.8, appendix E, port breakout screen shots, minor corrections and additions.
November 2018	Initial release

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

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

© 2018-2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

# Table of contents

Revisions.....	2
1 Introduction.....	6
2 Hardware overview.....	8
2.1 Dell EMC PowerEdge MX7000 chassis .....	8
2.2 Featured PowerEdge MX I/O modules used in this document .....	9
2.2.1 Dell EMC Networking MX9116n FSE .....	9
2.2.2 Dell EMC Networking MX5108n Ethernet switch .....	10
2.3 Dell EMC PowerEdge MX9002m module .....	11
2.4 Dell EMC PowerEdge MX740c.....	11
2.5 Rack-mounted networking switches .....	12
2.5.1 Dell EMC Networking S4148U-ON.....	12
2.5.2 Dell EMC Networking Z9100-ON.....	13
2.5.3 Dell EMC Networking Z9264F-ON .....	13
2.5.4 Dell EMC Networking S3048-ON .....	13
2.6 Dell EMC Unity 500F storage array.....	13
3 IOM operating modes.....	15
3.1 SmartFabric mode details.....	15
4 Server preparation.....	17
4.1 Reset server CNAs to factory defaults .....	17
4.2 Configure NIC partitioning on CNAs.....	17
5 Scenario 1: Connect MX9116n FSE to Fibre Channel storage .....	19
5.1 Define VLANs .....	21
5.2 Create identity pools.....	21
5.3 Create the SmartFabric .....	22
5.4 Breakout ports .....	23
5.5 Define uplinks .....	24
5.6 Create an FCoE server template.....	24
5.6.1 Associate server template with a VLAN and Identity Pool .....	26
5.6.2 Deploy a server template.....	26
6 Scenario 2: Connect MX5108n to Fibre Channel storage .....	28
6.1 Option 1: Connect to S4148U-ON in F_port mode.....	28
6.2 Option 2: Connect to S4148U-ON in NPG mode .....	28
6.3 Deployment steps.....	29

6.3.1	Define VLANs .....	29
6.3.2	Create identity pools .....	30
6.3.3	Create the SmartFabric .....	31
6.3.4	Breakout ports .....	32
6.3.5	Define uplinks .....	33
6.3.6	Create an FCoE server template.....	34
6.3.6.1	Associate server template with a VLAN and Identity Pool.....	35
6.3.6.2	Deploy a server template .....	35
<b>7</b>	<b>Scenario 3: Boot from SAN .....</b>	<b>37</b>
7.1	Configure NIC Boot Device .....	37
7.2	Configure BIOS settings .....	38
7.3	Connect FCoE LUN.....	38
7.4	Set up install media connection.....	39
7.5	Set up OS driver install media using Lifecycle Controller.....	39
<b>8</b>	<b>Verify configuration.....</b>	<b>40</b>
8.1	Validate MX7000 setup using OME-M .....	40
8.1.1	Show the MX chassis group topology .....	40
8.1.2	Show the SmartFabric status .....	41
8.1.3	Show port status .....	43
8.2	Switch CLI validation .....	43
8.2.1	show switch-operating-mode.....	43
8.2.2	show discovered-expanders.....	44
8.2.3	show unit-provision .....	44
8.2.4	show lldp neighbors.....	44
8.2.5	show vlan.....	45
8.2.6	show vlt <i>domain id</i> .....	45
8.2.7	show spanning-tree brief .....	45
8.2.8	show lldp dcbx interface ethernet <i>interface_number</i> .....	46
8.2.9	show lldp dcbx interface ethernet <i>interface_number</i> pfc detail .....	47
8.2.10	show lldp dcbx interface ethernet <i>interface_number</i> ets detail.....	48
8.2.11	show qos system.....	49
8.2.12	show policy-map.....	49
8.2.13	show qos interface ethernet <i>interface_number</i> .....	50
8.2.14	show class-map.....	50
8.2.15	show fc ns switch .....	50

8.2.16	show fcoe sessions .....	51
8.2.17	show vfabric .....	51
8.2.18	show fc switch .....	52
<b>A</b>	<b>Additional information .....</b>	<b>53</b>
A.1	Factory default PowerEdge MX7000 .....	53
A.1.1	Remove the SmartFabric .....	53
A.1.2	Use RACADM to reset each chassis .....	53
A.2	Factory default OS10EE .....	54
A.3	Spanning Tree Protocol recommendations .....	54
A.4	QSFP28 double density connectors .....	55
A.5	VLAN management and automated QoS .....	56
A.6	Rate limited 32 Gb Fibre Channel .....	57
A.7	Matching FC uplink port speed .....	58
A.8	Fibre Channel optics for MX9116n .....	58
<b>B</b>	<b>Storage connection information .....</b>	<b>59</b>
B.1	Determine Unity 500F storage array FC WWPNs .....	59
B.2	CNA FCoE port WWPNs .....	60
B.3	Configure Unity FC storage .....	61
B.3.1	Create a storage pool .....	61
B.3.2	Add hosts .....	62
B.3.3	Create LUNs and configure host access .....	62
<b>C</b>	<b>Validated components .....</b>	<b>64</b>
<b>D</b>	<b>FC attachment methods defined .....</b>	<b>65</b>
D.1	F_port .....	65
D.2	NPIV Proxy Gateway (NPG) .....	65
D.3	FCoE (FSB) .....	65
<b>E</b>	<b>Configure SRIOV on blade servers .....</b>	<b>66</b>
<b>F</b>	<b>Technical resources .....</b>	<b>69</b>
<b>G</b>	<b>Support and feedback .....</b>	<b>70</b>

# 1 Introduction

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



Figure 1 Dell EMC PowerEdge MX7000 chassis

The chassis features:

- Support for two and four-socket Intel based servers with massive storage potential, up to six local drives per two-socket server, and eight local drives per four-socket server
- Support for 16-drive 12Gbps SAS storage sleds
- Comprehensive I/O options including 25Gbps Ethernet, 12Gb SAS, and 32GB FC
- Three I/O networking fabrics:
  - Fabrics A and B for Ethernet connectivity
  - Fabric C for storage connectivity
- Midplane-free design for Fabrics A and B. Fabric A and B mezzanine cards dock directly with network I/O modules (IOMs) using orthogonal connectors. By removing the backplane there are no throughput limitations, resulting in high-speed connections

This guide provides step-by-step examples for deploying the network infrastructure that is required to attach external FC storage. Detailed examples for configuring the network for connecting FCoE FIP Snooping Bridge (FSB), NPIV Proxy Gateway (NPG), and Direct Connect of FC storage devices are provided. The Dell EMC Unity storage series is used in each demonstration. The goal is to enable a user to deploy the network for use with the storage protocols outlined in the guide.

This deployment guide assumes a basic understanding of the PowerEdge MX platform. Table 1 outlines what this document is and is not.

Table 1 Dell EMC PowerEdge MX Storage Network Deployment Guide - is/is not

This guide is	This guide is not
A deployment guide with step-by-step instructions	An overview of the MX Series
A networking guide for deploying external storage	A guide for deploying client Ethernet traffic or internal MX Series storage
Focused on network IOMs MX9116n Fabric Switching Engine (FSE) and MX5108n switches	A guide for configuring the MXG610s FC switch or MX5000s SAS switch
A guide for using SmartFabric Services to manage the network	A guide for Full Switch mode configurations
A step-by-step guide for configuring all FC connections in a single zone using SmartFabric Services	A guide for configuring multiple FC zones. For this, use Full Switch mode
An instruction manual for using the MX9116n FSE to connect to FC storage	A guide for connecting multiple chassis using MX9116n FSEs and MX7116n Fabric Expander Modules (FEM)

---

**Note:** Some of the figures in this guide show MX Series IOMs connecting to top-of-rack (TOR) or spine switches to demonstrate placement of these switches within common topologies. These TOR and spine switches are grayed out in the figures, as discussions of these networks are beyond the scope of this document. Refer to the [Dell EMC PowerEdge MX SmartFabric Network Deployment Guide](#) for instructions on configuring these upstream switches.

---

## 2 Hardware overview

The examples provided in this document were validated using specified networking switches and operating systems. This section describes the hardware used for validation. See [Appendix C](#) for a complete listing of the hardware and software validated for this guide.

Steps given in each scenario throughout this guide can be applied to comparable Dell EMC Networking switch models using the same networking operating system version, or later.

### 2.1 Dell EMC PowerEdge MX7000 chassis

Figure 2 shows the front view of the PowerEdge MX7000 chassis. The chassis can have one of three control panel options for administration, up to six hot-pluggable, redundant, 3,000-watt power supplies, and up to eight compute and storage sleds. Figure 2 shows the installed components:

- One, optional touchscreen LCD panel
- Two Dell EMC PowerEdge MX740c compute sleds in slots 1 and 2
- Six blank inserts in slots 3 through 8




Figure 2 Dell EMC PowerEdge MX7000—front

The MX7000 includes three I/O fabrics. Fabric A and B for Ethernet and future I/O module connectivity, and Fabric C for FC and SAS connectivity. Each Fabric provides two slots to provide redundancy. Figure 3 shows the back of the PowerEdge MX7000 chassis. From top to bottom, the chassis is configured with:

- One Dell EMC Networking MX9116n FSE installed in fabric slot A1
- One Dell EMC Networking MX7116n FEM installed in fabric slot A2
- Two Dell EMC PowerEdge MX9002m modules installed in management slots MM1 and MM2






Figure 3 Dell EMC PowerEdge MX7000—back

For attaching the MX5108n to FC (discussed in chapter 6), the most common scenario is to replace the MX9116n FSE in slot A1 and the MX7116n FEM in slot A2 with two MX5108n switches. For other recommended slot configurations, refer to the MX IOM support matrix found in the [Dell EMC PowerEdge MX Network Architecture Guide](#).

## 2.2 Featured PowerEdge MX I/O modules used in this document

Dell EMC offers several options for connecting external storage to the MX7000. Each option involves configuring an IOM installed into the rear side of the chassis. This guide provides overviews for the featured IOMs listed below, example topologies for each, and step-by-step instructions on how to configure them. For a complete overview of all supported PowerEdge MX Network devices, see the [Dell EMC PowerEdge MX Network Architecture Guide](#).

### 2.2.1 Dell EMC Networking MX9116n FSE

The MX9116n FSE provides 16 internal 25GbE ports, four QSFP28 100GbE ports for uplinks and twelve QSFP28-Double Density (QSFP28-DD) ports. These QSFP28-DD ports provide capacity for more uplinks, VLTi links, connections to rack servers at 10GbE or 25GbE, and connectivity for up to nine more MX7000 chassis by connecting to Dell EMC Networking MX7116n Fabric Expander Modules (FEM).

Example topologies and instructions for using the MX9116n FSE to connect to FC storage are covered in this guide. The first example demonstrates direct attach of FC storage, and the second example demonstrates NPG capabilities for attaching to an FC storage network.




Figure 4 Dell EMC Networking MX9116n FSE

The following MX9116n FSE components are labeled in Figure 4.

1. Express service tag
2. Storage USB port
3. Micro-B USB console port
4. Power and indicator LEDs
5. Handle release
6. Two QSFP28 ports
7. Two QSFP28 unified ports
8. 12 QSFP28-DD ports

## 2.2.2 Dell EMC Networking MX5108n Ethernet switch

The MX5108n offers eight internal 25GbE ports, two QSFP28 and one QSFP+ interface along with four 10GBase-T ports for uplinks. The ports can be used to provide a combination of network uplink, VLT interconnects (VLTi), or for FCoE (FSB) connectivity, but does not support NPG or FC capabilities. A MX7000 chassis supports up to four MX5108n Ethernet switches in Fabric A and/or B.

Example topologies and instructions for using FCoE (FSB) on the MX5108n to connect to FC storage are covered in chapter 6. One example connects the MX5108n to a Dell EMC Networking S4148U-ON in NPG mode, while a second example connects the MX5108n to a Dell EMC Networking S4148U-ON in F\_port mode.




Figure 5 Dell EMC Networking MX5108n Ethernet switch

The following MX5108n components are labeled in Figure 5.

1. Express service tag
2. Storage USB port
3. Micro-B USB console port

4. Power and indicator LEDs
5. Module insertion/removal latch
6. One QSFP+ port
7. Two QSFP28 ports
8. Four 10GBASE-T ports

## 2.3 Dell EMC PowerEdge MX9002m module

The Dell EMC MX9002m module controls the overall chassis power and cooling. The MX9002m module hosts the OpenManage Enterprise Modular, or OME-M, console. Most of the steps in this document use the OME-M console user interface. The MX9002m module also contains Ethernet ports to allow administrator connectivity to all MX-series components.

The MX7000 supports two MX9002m modules for redundancy. Figure 6 shows a single MX9002m module and its components.




Figure 6 Dell EMC PowerEdge MX9002m module

The following MX9002m module components are labeled in Figure 6.

1. Handle release
2. Gigabit Ethernet port 1
3. Gigabit Ethernet port 2
4. ID button and health status LED
5. Power status LED

Refer to the [Dell EMC PowerEdge MX Network Architecture Guide](#) for more information on the MX9002m, including cabling chassis together to form a Multi-Chassis Management group.

## 2.4 Dell EMC PowerEdge MX740c

The PowerEdge MX740c is a two-socket, full-height, single-width compute sled with impressive performance and scalability when paired with the MX-series system. It is ideal for dense virtualization environments and can serve as a foundation for collaborative workloads. An MX7000 chassis supports up to eight MX740c compute sleds.

PowerEdge MX740c key features include:

- Single slot design

- Two CPU sockets
- 24 DIMM slots of DDR4 memory
- Up to six SAS/SATA SSD/HDD and NVMe PCIe SSDs
- Two PCIe mezzanine card slots for connecting to network Fabric A and B
- One PCIe mini-mezzanine card slots for connecting to storage Fabric C
- iDRAC9 with Lifecycle Controller



Figure 7 Dell EMC PowerEdge MX740c compute sled with six 2.5-inch SAS drives

**Note:** Each MX7000 chassis example in this document has two MX740c compute sleds installed.

## 2.5 Rack-mounted networking switches

This section covers the rack-mounted Ethernet switches discussed in this document.

### 2.5.1 Dell EMC Networking S4148U-ON

The Dell EMC Networking S4148U-ON is a 1-Rack Unit (RU) switch with 48x SFP+ ports, 2x QSFP+ ports, and 4x QSFP28 ports. In this document, two S4148U-ON supports storage traffic, and is the first of two leaf switch options.




Figure 8 Dell EMC Networking S4148U-ON

## 2.5.2 Dell EMC Networking Z9100-ON

The Dell EMC Networking Z9100-ON is a 1RU fixed form-factor 100GbE switch with 32 QSFP28 ports. Each QSFP28 port can operate at 10, 25, 40, 50, or 100 Gbps. This switch may be used as a leaf or spine switch in a Leaf-spine topology.



Figure 9 Dell EMC Networking Z9100-ON

## 2.5.3 Dell EMC Networking Z9264F-ON

The Dell EMC Networking Z9264F-ON is a 2RU fixed form-factor 100GbE multi-rate switch is optimized for non-blocking 100GbE leaf/spine fabrics and high-density 25/50GbE in-rack server and storage connections. Provides up to 64 ports of 100GbE QSFP28 or up to 128 ports of 1/10/25/40/50GbE ports using breakout cables. This switch may be used as a leaf or spine switch in a Leaf-spine topology.




Figure 10 Dell EMC Networking Z9264F-ON

## 2.5.4 Dell EMC Networking S3048-ON

The Dell EMC Networking S3048-ON is a 1-Rack Unit (RU) switch with forty-eight 1 GbE Base-T ports and four 10 GbE SFP+ ports. One S3048-ON was used to support management traffic while validating the scenarios in this document.



Figure 11 Dell EMC Networking S3048-ON

## 2.6 Dell EMC Unity 500F storage array

The Unity 500F storage platform delivers all-flash storage with up to 8PB raw capacity. It has concurrent support for NAS, iSCSI, and FC protocols. The Disk Processing Enclosure (DPE) has a 2-RU form factor, redundant Storage Processors (SPs), and supports up to twenty-five 2.5" drives. Additional 2-RU Disk Array Enclosures (DAEs) may be added providing twenty-five additional drives each. This array is used for the FC connections in this document.



Figure 12 Dell EMC Unity 500F front view



Figure 13 Dell EMC Unity 500F rear view

## 3 IOM operating modes

PowerEdge MX networking IOMs operate in one of two modes:

- Full Switch Mode – Enabled by default, all switch-specific OS10EE capabilities are available
- SmartFabric Mode – IOMs operate as Layer 2 I/O aggregation devices and are managed through the OME-M console

The examples in this guide use SmartFabric mode. Table 2 outlines the differences between the two operating modes. These differences apply to both the MX9116n FSE and the MX5108n.

Table 2 IOM operating mode differences

Full Switch mode	SmartFabric mode
All OS10 configuration changes are persistent during power cycle events.	OS10 configuration changes made using the commands listed in section 3.1 below are persistent across power cycle events. Other commands are disabled.
All switch interfaces are assigned to VLAN 1 by default and are in the same L2 bridge domain.	L2 bridging is disabled, and interfaces must join a bridge domain (VLAN) before being able to forward frames.
All configurations changes are saved in the running configuration by default. To display the current configuration, use the <code>show running-configuration</code> command.	Verify configuration changes using feature-specific show commands, such as <code>show interface</code> and <code>show vlan</code> , instead of <code>show running-configuration</code> .
Allows for single or multiple FC zones.	Allows for a single FC zone.

### 3.1 SmartFabric mode details

A SmartFabric is a logical entity consisting of a collection of physical resources, such as servers and switches, and logical resources, such as networks, templates, and uplinks. The OME-M console provides a method to manage these resources as a single unit and supports most switch configuration settings. SmartFabric mode supports the following CLI commands:

- `clock` - Configure clock parameters
- `end` - Exit to the EXEC mode
- `exit` - Exit from the current mode
- `help` - Display available commands
- `hostname` - Set the system hostname
- `interface` - Configure or select an interface
- `ip nameserver` - Configure nameserver
- `logging` - Configure system logging
- `management route` - Configure the IPV4/IPV6 management route
- `no` - Delete or disable commands in Configuration mode
- `ntp` - Configure the network time protocol
- `show switch-operating-mode` - shows current mode, SmartFabric or FullSwitch
- `snmp-server` - Configure the SNMP server

- `username` - Create or modify user credentials
- `spanning-tree`
  - `disable` - Disable spanning tree globally
  - `mac-flush-timer` - Set the time used to flush MAC address entries
  - `mode` - Enable a spanning-tree mode, such as RSTP or MST
  - `mst` - Configure multiple spanning-tree (MST) mode
  - `rstp` - Configure rapid spanning-tree protocol (RSTP) mode
  - `vlan` - Configure spanning-tree on a VLAN range

---

**Note:** The CLI `show running configuration` command does not provide complete information when in SmartFabric mode. Dell EMC recommends using other CLI show commands or the OME-M console to view switch information when using SmartFabric.

---



## 4 Server preparation

The examples in this guide use MX740c compute sleds with QLogic (model QL41262HMKR) Converged Network Adapters (CNAs) installed. CNAs are required to achieve FCoE connectivity. Use the steps below to prepare each CNA by setting them to factory defaults (if required) and configuring NPAR (NIC partitioning).

---

**Note:** iDRAC steps in this section may vary depending on hardware, software and browser versions used. See the [Installation and Service Manual](#) for your PowerEdge server for instructions on connecting to the iDRAC. From the link, select your server, then Manuals and documents.

---

### 4.1 Reset server CNAs to factory defaults

Reset the CNAs to their factory defaults using the steps in this section. Resetting CNAs to factory default is only necessary if the CNAs installed have been modified from their factory default settings.

1. From the OME-M user interface, select the server to use to access the storage.
2. Launch the server **Virtual Console**.
3. From the **Virtual Console**, select **Next Boot** then **BIOS Setup**.
4. Reboot the server.
5. From the System Setup Main Menu, select **Device Settings**.
6. From the **Device Settings** page, select the first CNA port.
7. From the **Main Configuration** page, click the **Default** button.
8. Click **Yes** to load the default settings, and then click **OK**.
9. Click **Finish**. Notice if a message indicates a reboot is required for changes to take effect.
10. Click **Yes** to save changes, then click **OK**.
11. Repeat the steps in this section for each CNA port listed on the **Device Settings** page.

If required per step 9, reboot the system and return to **System Setup** to configure NIC partitioning.

### 4.2 Configure NIC partitioning on CNAs

In this section, each QLogic CNA port is partitioned into one Ethernet and one FCoE partition.

---

**Note:** This is only done on CNA ports that carry converged traffic. In this example, these are the two 25GbE QLogic CNA ports on each server that attach to the switches internally through an orthogonal connection.

---

If the system is already in **System Setup** from the previous section, skip to step 4.

1. Using a web browser, connect to the iDRAC server and launch the **Virtual Console**.
2. From the **Virtual Console**, click **Next Boot** menu then select **BIOS Setup**.
3. Select the option to reboot the server.
4. On the System Setup Main Menu, select **Device Settings**.
5. Select the first CNA port.
6. Select **Device Level Configuration**.
7. Set the **Virtualization Mode** to **NPAR**, if not already set, and then click **Back**.
8. Select **NIC Partitioning Configuration**, **Partition 1 Configuration**, and click to set the **NIC + RDMA Mode** to **Disabled**.
9. Click **Back**.

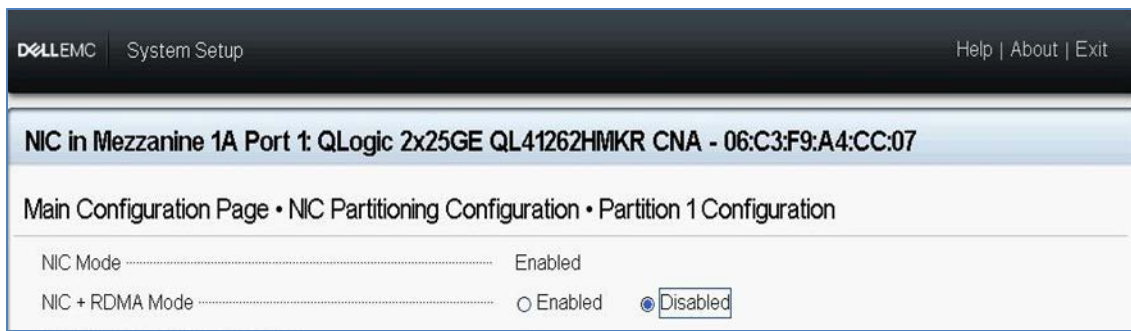


Figure 14 CNA partition 1 configuration

10. Select **Partition 2 Configuration** and set the **NIC Mode** to **Disabled**.
11. Set the **FCoE Mode** to **Enabled**, then click **Back**.

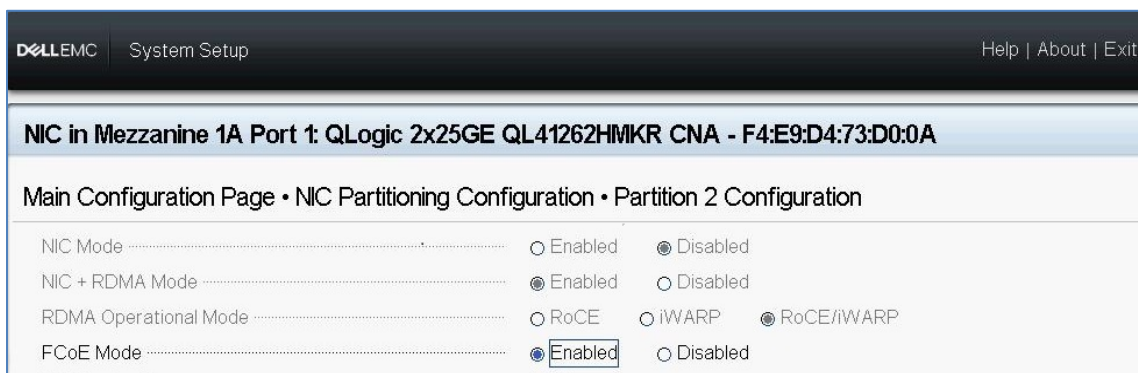


Figure 15 CNA partition 2 configuration

12. If present, select **Partition 3 Configuration** and set all modes to **Disabled**, then click **Back**.
13. If present, select **Partition 4 Configuration** and set all modes to **Disabled**, then click **Back**.
14. Click **Back**, and then **Finish**.
15. When prompted to save changes, click **Yes** and then click **OK** in the **Success** window.
16. Select the second CNA port and repeat steps in this section for port 2
17. After configuring port 2, click **Finish**, and then **Finish**.
18. Click **Yes** to exit and reboot.

---

**Note:** This server configuration may be used in later chapters to generate a template to deploy to other servers with identical hardware. When a template is not used, steps in this chapter should be repeated for each MX server sled requiring access to FC storage.

---

## 5 Scenario 1: Connect MX9116n FSE to Fibre Channel storage

This chapter discusses two methods for connecting the MX9116n FSE to the Dell EMC Unity FC storage array.

An MX7000 server sled uses an MX9116n FSE switch to either directly attach to FC storage (Figure 16), or to attach to FC storage through an FC switch (Figure 17). Each option is described below and includes steps for deployment.

Using SmartFabric, the steps to configure either option on the MX9116n FSE are virtually identical. The only difference is a single mouse click while assigning uplink ports. This step is in section 5.5.

---

**Note:** The deployment steps in this chapter assume the chassis are connected to a management network with assigned management IP addresses, and joined in a chassis group. See the [Dell EMC PowerEdge MX Network Architecture Guide](#) for the detailed configuration of these management features.

---

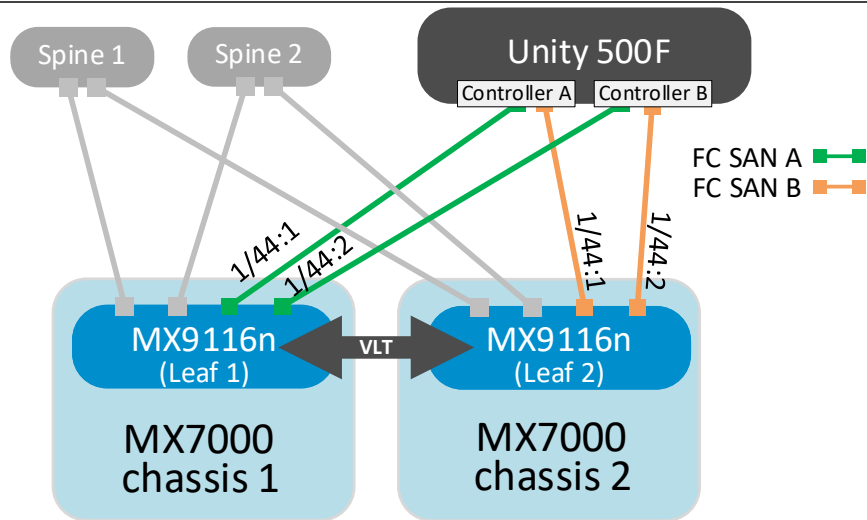


Figure 16 Fibre Channel (F\_Port) direct connect network to Dell EMC Unity

---

**Note:** See [Appendix A.8](#) for supported Gigabit Fibre Channel (GFC) transceivers on MX9116n unified ports. This section demonstrates Fibre Channel directly attaching to the Unity 500F. MX9116n FSE universal ports 44:1 and 44:2 are required for FC connections and operate in F\_PORT mode, which allows for an FC storage array to be connected directly to the MX9116n FSE.

---

**Note:** This topology is supported using MX9116n. It is not supported using the MX5108n.

---

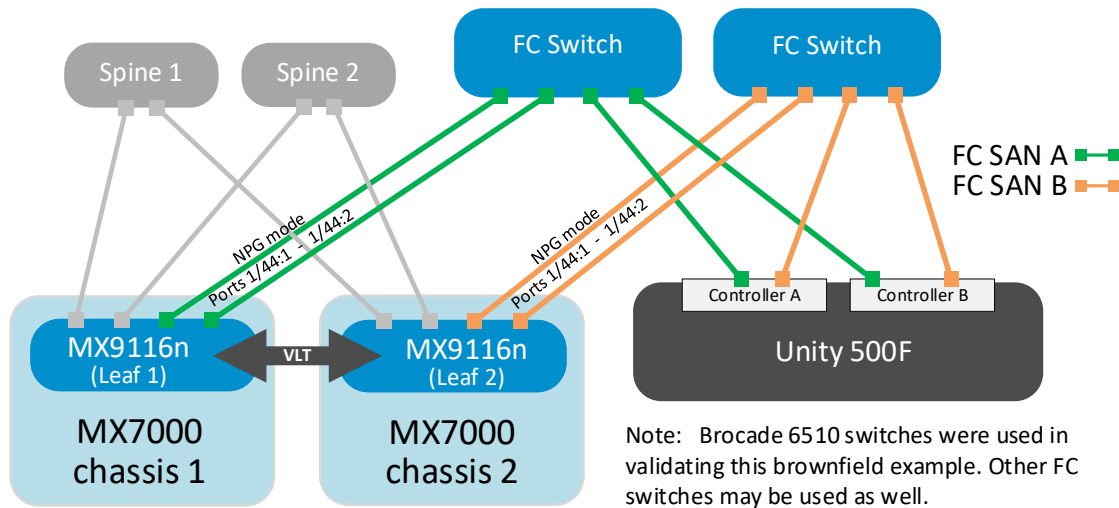


Figure 17 FC (NPG) network to Dell EMC Unity

This topology shows the storage network to provide connections from the servers in the MX7000 chassis to an FC storage array using the MX9116n FSE switches in NPG mode, going to the FC switches connecting directly into the FC array. NPG allows for a switch-based solution for larger SAN deployments. Setting the switch ports to NPG mode allows the switch to act as host to the FC switch of another vendor's FC switch. This is helpful since FC switch vendors do not generally support the connection of other vendor switches to form an FC link between their E-ports.

This scenario shows attachment to a brownfield FC switch infrastructure. Configuration of the existing FC switches is beyond the scope of this document.

---

**Note:** This topology is only supported using MX9116n. It is not supported using the MX5108n.

SmartFabric deployment consists of the six key steps. The following steps are completed using the OME-M console:

1. Define VLANs to use in the fabric.
2. Create the Identity Pools.
3. Select the switches and create the SmartFabric based on the desired physical topology.
4. Breakout uplink ports.
5. Use F\_Port to define direct uplinks from the fabric to the FC storage, or to the existing FC switch using NPG, and assign VLANs to those uplinks.
6. Create and deploy the appropriate server templates to the compute sleds.

The remainder of this chapter provides the instructions necessary for completing the steps above for SmartFabric deployment.

## 5.1 Define VLANs

To define VLANs using the OME-M console, perform the following steps:

1. Open the OME-M console.
2. From the navigation menu, click **Configuration**, then click **Networks**.
3. In the **Network** pane, click **Define** to open the **Define Network** window.
4. In the **Name** box, enter **Default**.
5. Optionally, enter a description in the **Description** box.
6. In the **VLAN ID** box, enter **1**.
7. From the **Network Type** list, select **General Purpose (Bronze)**.
8. Click **Finish**.
9. Using the VLAN attributes data from Table 3, repeat the steps in this section to create the other VLANs.

Table 3 VLAN attributes

Name	Description	Network Type	VLAN ID	SAN
Default	Default VLAN	General Purpose (Bronze)	1	N/A
FC A1	FCOE A1	Storage - FCoE	30	A
FC A2	FCOE A2	Storage - FCoE	40	B

**Note:** For information on Network Types, see A.5.

Figure 18 shows the two new VLANs.

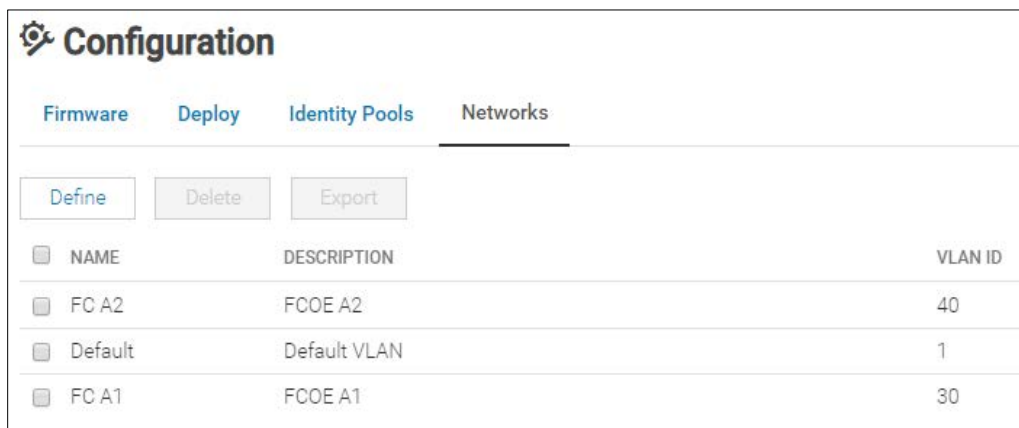


Figure 18 Defined VLAN list

## 5.2 Create identity pools

Virtual identity pools are used in conjunction with server templates to automate network onboarding of compute sleds. Follow these steps to create an ID pool:

1. Open the OME-M console.
2. From the navigation menu, click **Configure**, then click **Identity Pools**.
3. In the **Network** panel, click **Create**. The **Create Identity Pool** window displays.
4. Type Ethernet CNA into the **Pool Name** box.
5. Optionally, enter a description in the **Description** box.
6. Click **Next**.
7. Click to select the **Include Ethernet Virtual MAC Addresses** option.
8. In the **Starting MAC Address** box, type a unique MAC address (e.g. 06:3C:F9:A4:CC:00).
9. Type 255 in the **Number of Virtual MAC Identities** box, click **Next**, then click **Next** again.
10. Select the **Include FCoE Identity** option.
11. In the **Starting MAC Address** box, type a unique MAC address (e.g. 06:3C:F9:A4:CD:00).
12. Type 255 in the **Number of FCoE Identities** box.

The screenshot shows the 'Create Identity Pool' window. On the left, a sidebar lists 'Pool Information', 'Ethernet', 'iSCSI', 'FCoE', and 'Fibre Channel', each with a green checkmark. 'FCoE' is highlighted in blue. Below the sidebar, it says 'Step 4 of 5'. The main content area has a checked box for 'Include FCoE Identity' and a note: 'The WWPN and the WWNN will be generated by prefixing the MAC address with 0x2001 and 0x2000 respectively.' Below this, there are two input fields: 'Starting MAC Address' with the value '06:C3:F9:A4:CD:00' and 'Number of FCoE Identities' with the value '255'. At the bottom right, there are four buttons: 'Previous', 'Next', 'Finish', and 'Cancel'.

Figure 19 Include FCoE Identity

13. Click **Finish**, then click **Finish** again.

---

**Note:** The starting MAC address must be a locally administered unicast address.

---

## 5.3 Create the SmartFabric

To create a SmartFabric, use the OME-M console to perform the following steps:

1. Open the OME-M console.
2. From the navigation menu, click **Devices**, then click **Fabric**.
3. In the **Fabric** panel, click **Add Fabric** to open the **Create Fabric** window.
4. Enter **SmartFabric** in the **Name** box.
5. Optionally, enter a description in the **Description** box.
6. Click **Next**.
7. From the Design Type drop-down menu, select **2x M9116n Fabric Switching Engine in different chassis**.

Figure 20 Select Design Type

8. From the **Chassis-X** list, select the first MX7000 chassis to join the fabric.
9. From the **Switch-A** list, select **Slot-IOM-A1**.
10. From the **Chassis-Y** list, select the second MX7000 chassis to join the fabric.
11. From the **Switch-B** list, select **Slot-IOM-A2**, then click **Next**.
12. On the **Summary** page, verify the proposed configuration, then click **Finish**.

---

**Note:** A list of the physical cabling requirements may be printed from the Summary screen.

---

## 5.4 Breakout ports

Perform the following steps on each MX9116n FSE.

---

**Note:** Port-group 1/1/16 is used for FC connections in this example.

---

1. Open the OME-M console.
2. From the navigation menu click **Devices**, then click **I/O Modules**.
3. In the **Devices** panel, click to select the IOM to configure.
4. In the **IOM** panel click **Hardware**, then click **Port Information**.

---

**Note:** See [SmartFabric Mode - MX Port-Group Configuration Errors](#) video for more information on configuration errors.

---

5. Click the **port-group 1/1/16** check box, then click **Configure breakout**.
6. In **Configure breakout** panel, select **HardwareDefault** as the breakout type.
7. Click **Finish**.
8. To set the port group 1/1/16 to **4X16GFC**, select the **port-group 1/1/16** check box, then click **Configure breakout**.
9. In **Configure breakout** panel, select **4X16GFC** as the breakout type.
10. Click **Finish**.

---

**Note:** The Fibre Channel ports that are broken out are administratively down by default. Select the ports and click Toggle Admin State button. Click Finish to administratively set the ports to up.

---

## 5.5 Define uplinks

**Note:** The steps in this section allow you to connect to the FC storage using an F\_Port, for direct attach, or in NPG mode, for attaching to an FC switch. This is the only setting within the MX chassis that distinguishes the two configurations.

After the initial deployment, the new fabric shows **Uplink Count** as 'zero' and shows a warning (⚠️) icon. The lack of a fabric uplink results in a failed health check, which displays the (❌) icon.

To create uplinks, follow these steps:

1. Open the OME-M console.
2. From the navigation menu click **Devices**, then click **Fabric**.
3. Click the **SmartFabric** fabric name.
4. In the **Fabric Details** panel, click **Uplinks**, then click the **Add Uplinks** button.
5. From the **Add Uplinks** window, use the information in Table 4 to enter an uplink name in the **Name** box.
6. Optionally, enter a description in the **Description** box.
7. From the **Uplink Type** list, select **Type**, as defined in Table 4, then click **Next**.
8. From the **Switch Ports** list, select the FC ports as defined in Table 4.
9. From the **Tagged Networks** list, select VLAN defined in Table 4, then click **Finish**. SmartFabric creates the uplink object and the status for the fabric changes to **OK**.

Table 4 Uplink attributes

Uplink Name	Description	Type	Ports	VLAN
FCoE A1 *	FC Uplink for switch in Slot A1	FC Gateway/Direct Attach	1/1/44:1 and 1/1/44:2	30
FCoE A2 *	FC Uplink for switch in Slot A2	FC Gateway/Direct Attach	1/1/44:1 and 1/1/44:2	40

\*Select the appropriate Fibre Channel type and port for the connected uplink. Select either FC Gateway or FC direct attach.

---

**Note:** Use the guidelines provided in chapter 8 to verify the SmartFabric configuration.

---

## 5.6 Create an FCoE server template

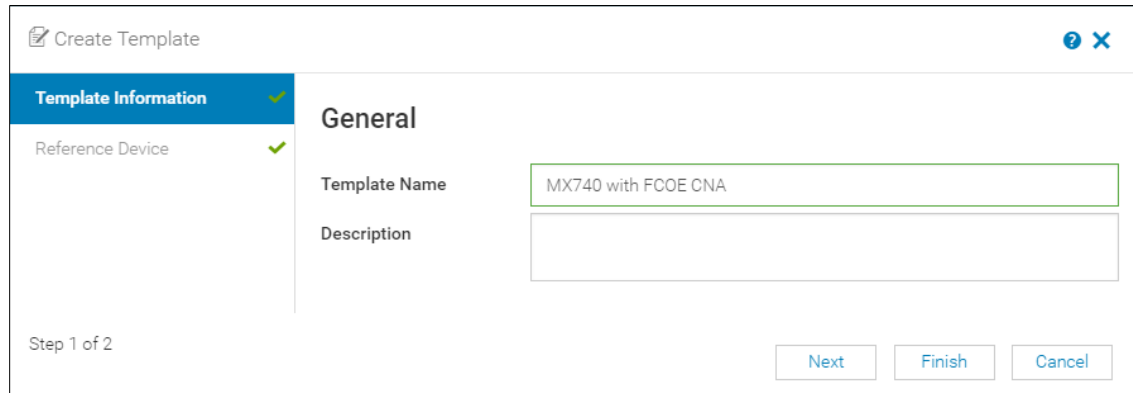
A server template contains parameters extracted from a server and allows these parameters to be quickly applied to multiple compute sleds. The server template allows an administrator to associate VLANs to compute sleds. The templates contain settings for the following categories:

- Local access configuration
- Location configuration
- Power configuration
- Chassis network configuration
- Slot configuration
- Setup configuration



To create a server template, follow these steps:

1. Open the OME-M console.
2. From the navigation menu, click **Configuration**, then click **Deploy**.
3. From the center panel, click **Create Template**, then click **From Reference Device** to open the **Create Template** window.
4. In the Template Name box, enter MX740c with FCOE CNA.



The screenshot shows a 'Create Template' dialog box. On the left, a sidebar titled 'Template Information' has a 'Reference Device' section with a green checkmark. The main area is titled 'General' and contains a 'Template Name' field with the text 'MX740 with FCOE CNA' and an empty 'Description' field. At the bottom, there are 'Next', 'Finish', and 'Cancel' buttons. The progress indicator shows 'Step 1 of 2'.

Figure 21 Create template dialog box

5. Optionally, enter a description in the **Description** box, then click **Next**.
6. In the Device Selection section, click Select Device.
7. From the **Select Devices** window, choose the server previously configured, then click **Finish**.
8. From the **Elements to Clone** list, select **iDRAC**, **System**, **NIC**, and then click **Finish**.

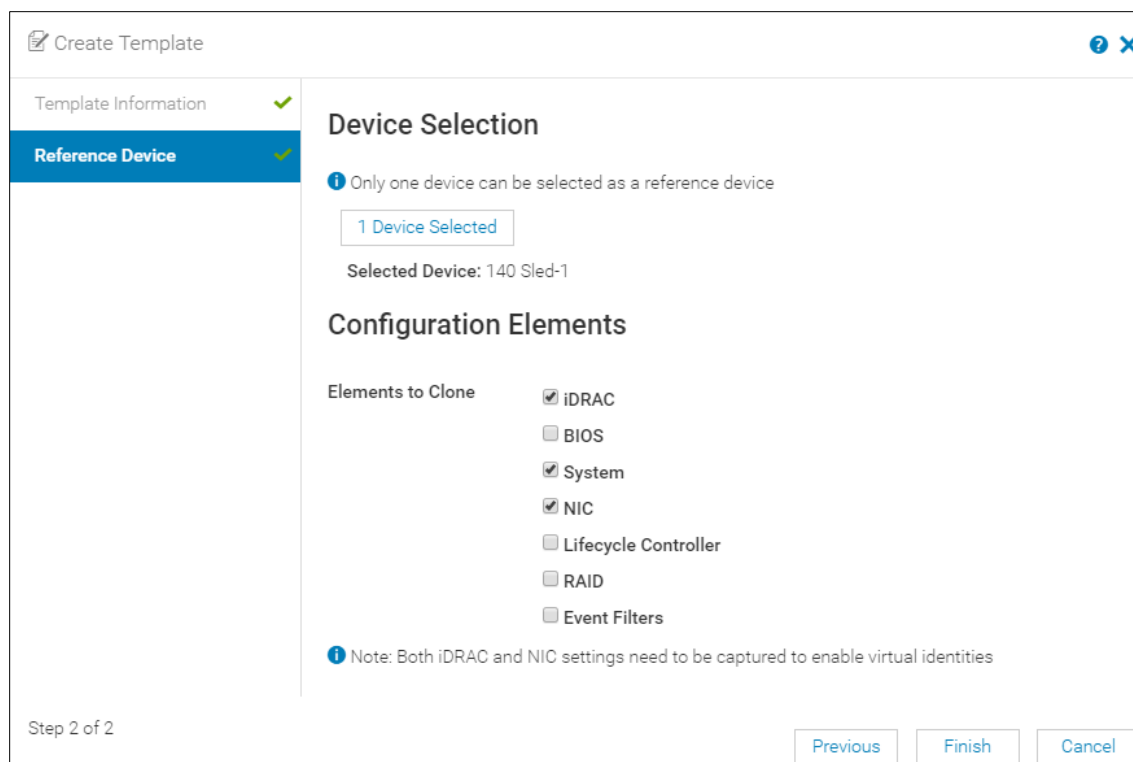


Figure 22 Select Devices and elements to clone

---

**Note:** iDRAC, System, and NIC are the minimum elements required for this example to function correctly. Other elements are optional.

A job starts, and the new server template displays on the list. When complete, the **Completed successfully** status displays.

### 5.6.1 Associate server template with a VLAN and Identity Pool

After successfully creating a new template, associate the template with a network:

1. From the **Deploy** panel, select the **MX740c with FCoE CNA** server template.
2. Click **Edit Network**.
3. From the **Identity Pool** list, choose **Ethernet CNA**.
4. For the NIC in **Mezzanine 1A Port 1**, locate the **Tagged Network** list and choose **FC A1**.
5. For the NIC in **Mezzanine 1A Port 2**, locate the **Tagged Network** list and choose **FC A2**, then click **Finish**.

---

**Note:** Ports running FCoE must have their untagged VLAN undefined, as SmartFabric mode auto-assigns VLAN 1 for FIP packets. Failure to do so will throw an error in creating or updating the Server Interface Profile.

### 5.6.2 Deploy a server template

To deploy the server template, complete the following steps:

1. From the **Deploy** panel, select the **MX740c with FCoE CNA** server template.

2. Click **Deploy Template** to open the **Deploy Template** window.
3. Click the **Select** button to choose slots for template deployment.
4. Select the **Do not forcefully reboot the host OS** option.
5. Click **Next, Run Now**, and then click **Finish**.

Interfaces on the MX9116n FSE are updated automatically. SmartFabric configures each interface with the untagged and tagged VLANs configured above. Additionally, SmartFabric deploys associated QoS settings, as shown in Appendix A.5. On the server reboot or when turning the power on, the server may run a configuration job to change settings to match the reference server.

---

**Note:** Only Fabric A was used in this section.

---

Your system is now ready to connect to Fibre Channel storage. See Appendix B for setting up storage logical unit numbers (LUNs).

## 6 Scenario 2: Connect MX5108n to Fibre Channel storage

This chapter provides instructions for connecting the MX5108n to Fibre Channel storage. Two options are included for attaching using FCoE FSB. Each option is described below and includes steps for deployment.

---

**Note:** The FCoE (FSB) examples in this chapter use the Dell EMC Networking MX5108n. The same instructions may also be applied and used with the MX9116n.

---

**Note:** The deployment steps below assume the chassis are connected to a management network with assigned management IP addresses, and joined in a chassis group. Refer to the [Dell EMC PowerEdge MX Network Architecture Guide](#) for detailed steps for configuring these management features.

---

### 6.1 Option 1: Connect to S4148U-ON in F\_port mode

The first FCoE (FSB) topology uses the S4148U-ON in F\_port mode to connect the MX5108n to an existing Fibre Channel network. The S4148U-ON switches are in F\_port mode and do not have a VLT or LAG. The MX5108ns have two FCoE VLANs that do not traverse the VLTi link.

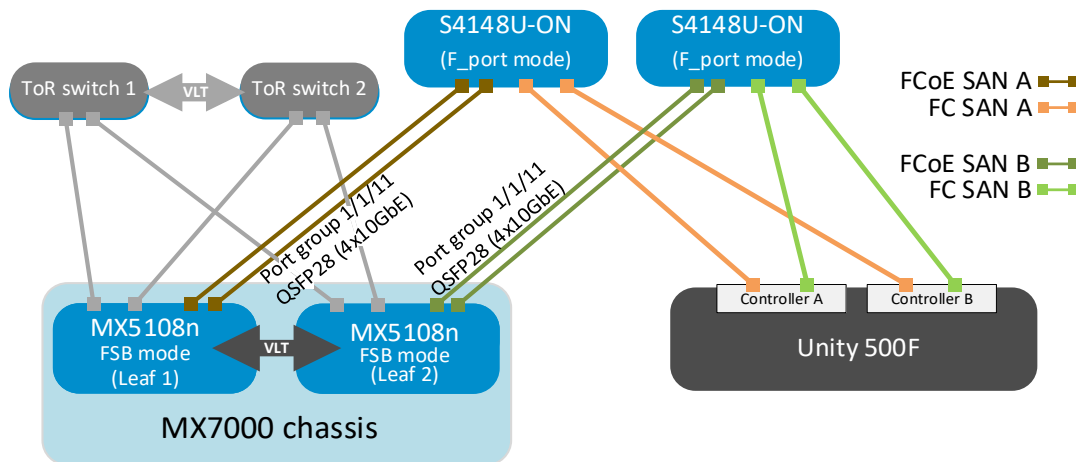


Figure 23 FCoE (FSB) Network to Dell EMC Unity through F\_port mode switch

---

**Note:** Use the steps described in Section 6.3 to configure the MX chassis for this option. See the [FCoE-to-Fibre Channel Deployment with S4148U-ON](#) document for instructions on configuring S4148U-ON in F\_port mode.

---

### 6.2 Option 2: Connect to S4148U-ON in NPG mode

Figure 24 shows how the MX5108n can add FCoE functionality to an existing LAN-only topology as an FCoE FSB.

The Dell EMC Networking MX5108n is configured to operate in FSB mode. The MX5108n snoops FIP packets on FCoE-enabled VLANs and discovers the following information:

1. End nodes (ENodes)
2. Fibre Channel forwarder (FCF)
3. Connections between ENodes and FCFs

#### 4. Sessions between ENodes and FCFs

Using the discovered information, the switch installs ACL entries that provide security and point-to-point link emulation.

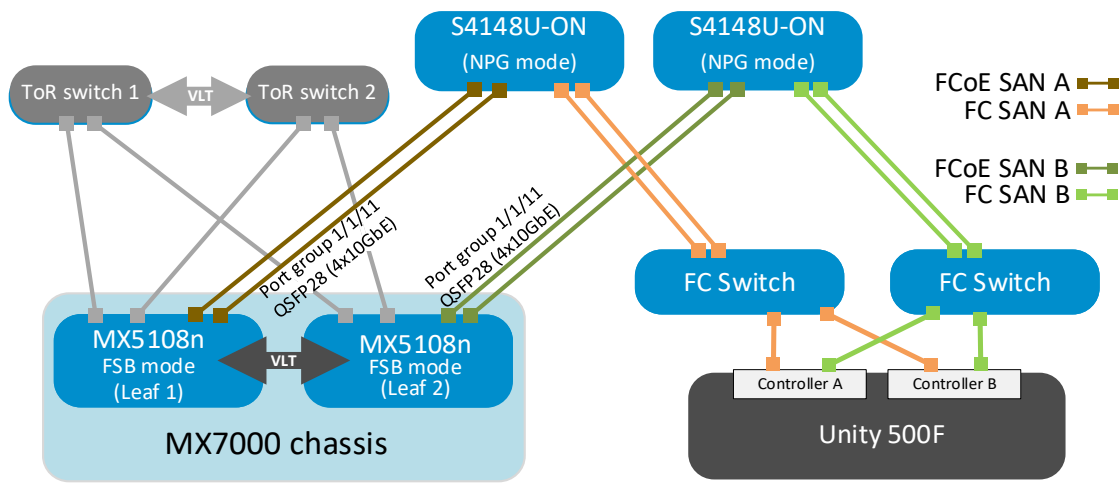


Figure 24 FCoE (FSB) Network to Dell EMC Unity through NPG mode switch

S4148U-ON switches in Figure 24 do not have a VLT or LAG and are NPG mode devices only. The figure has two FCoE VLANs that do not traverse the VLTi link.

---

**Note:** See the [OS10 Enterprise Edition User Guide](#) for configuring NPG mode globally on the S4148U-ON switches.

---

## 6.3 Deployment steps

SmartFabric deployment for both topologies above consists of six steps completed using the OME-M console:

1. Define VLANs to be used in the fabric.
2. Create Identity Pools.
3. Select switches and create the SmartFabric based on the physical topology desired.
4. Breakout uplink ports.
5. Define uplinks from the fabric to the existing network and assign VLANs to those uplinks.
6. Create and deploy the appropriate server templates to the compute sleds.

Follow the instructions in the remainder of this chapter to configure the MX7000 to attach to the FC array.

### 6.3.1 Define VLANs

To define VLANs using the OME-M console, perform the following steps:

1. Open the OME-M console.
2. From the navigation menu, click **Configuration**, then click **Networks**.
3. In the **Networks** panel, click **Define**.
4. In the **Name** box, enter **Default**.
5. Optionally, enter a description in the **Description** box.

6. Enter **1** in the **VLAN ID** box.
7. From the **Network Type** list, select **General Purpose (Bronze)**, then click **Finish**.
8. Using the information in Table 5, repeat the steps in this section to create the remaining VLANs.

Table 5 VLAN attributes

Name	Description	Network Type	VLAN ID
Default	Default VLAN	General Purpose (Bronze)	1
FC A1	FCOE A1	Storage - FCoE	30
FC A2	FCOE A2	Storage - FCoE	40

**Note:** For information on Network Types, see Appendix A.5.

In a SmartFabric deployment, there is no default VLAN, such as VLAN 1. The default VLAN must be created for any untagged traffic to cross the fabric. Figure 25 shows the three VLANs after being created.

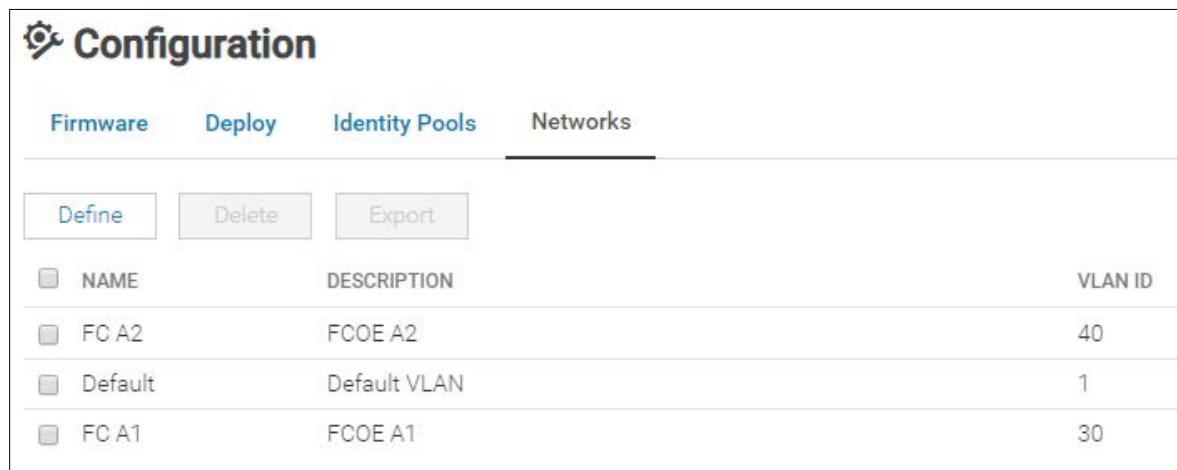


Figure 25 Defined VLAN list

### 6.3.2 Create identity pools

Identity Pools are used in conjunction with server templates to automate network onboarding of compute sleds. To create an ID pool, follow these steps:

1. Open the OME-M console.
2. From the navigation menu, click **Configure**, then click **Identity Pools**.
3. In the **Network** panel, click **Create**.
4. Enter Ethernet CNA in the **Pool Name** box.
5. Optionally, enter a description in the **Description** box.
6. Click **Next**.
7. Select the **Include ethernet virtual MAC Addresses** option.
8. In the **Starting MAC Address** box, type a unique address (e.g. 06:3C:F9:A4:CC:00).
9. Enter 255 in the **Number of Virtual MAC Identities** box.
10. Click **Next**, then click **Next** again.
11. Select the **Include FCoE Identity** option.

12. In the **Starting MAC Address** box, type a unique address (e.g. 06:3C:F9:A4:CD:00).
13. Enter 255 in the **Number of Virtual MAC Identities** box.
14. Click **Finish**, then click **Finish** again.

---

**Note:** The starting MAC address must be a locally administered unicast MAC address. These addresses are user-defined.

---

### 6.3.3 Create the SmartFabric

To create a SmartFabric, use the OME-M console and perform the following steps:

1. Open the OME-M console.
2. From the navigation menu click **Devices**, then click **Fabric**.
3. In the **Fabric** panel, click **Add Fabric** to open the **Create Fabric** window.
4. Enter **SmartFabric** in the **Name** box.
5. Optionally, enter a description in the **Description** box.
6. Click **Next**.
7. From the **Design Type** list, select **2x M5108n Fabric Switching Engine in same chassis**.
8. From the **Chassis-X** list, select the first MX7000 chassis.
9. From the **Switch-A** list, select **Slot-IOM-A1**.
10. From the **Chassis-Y** list, select the second MX7000 chassis to join the fabric.
11. From the **Switch-B** list, select **Slot-IOM-A2** then click **Next**.

The screenshot shows the 'Create Fabric' configuration window. The 'Design Type' is set to '2xMX5108n Ethernet Switches in same chassis'. A diagram illustrates two server chassis connected by a fabric. The 'Chassis-X' is 'Chassis SKY002Z', 'Switch-A' is 'Slot-IOM-A1: BZTQPK2', and 'Switch-B' is 'Slot-IOM-A2: 6L59XM2'. The window includes 'Previous', 'Next', and 'Cancel' buttons and indicates 'Step 2 of 3'.

Figure 26 Select Design Type

12. On the **Summary** page, verify the proposed configuration and click **Finish**.

---

**Note:** A list of the physical cabling requirements can be printed from the Summary window.

---

## 6.3.4 Breakout ports

Perform the following steps on each MX5108n switch.

**Note:** In this example, we are going to break out port group 1/1/11 into 4 x 10GbE interfaces. To do so, the port group is put into default configuration, and then put into 4x10GbE mode.

1. From the navigation menu click **Devices**, then click **I/O Modules**.
2. In the **Devices** panel, click to select the IOM to configure.
3. In the **IOM** panel click **Hardware**, then click **Port Information**.

The screenshot shows the OpenManage Enterprise Modular interface for IOM-B1. The top navigation bar includes Home, Devices, Configuration, Alerts, Monitor, and Application Settings. The main content area is titled 'IOM-B1' and shows Health: Ok, State: On, IP: 100.67.163.224, and Service Tag: CBJP9N2. Below this are tabs for Overview, Hardware, Firmware, Alerts, and Settings. The Hardware tab is active, and the Port Information sub-tab is selected. A table of port information is displayed below the sub-tab.

PORT NUMBER	PORT NAME	PORT DESCRIPTION	OPERATIONAL STATUS	ADMIN STATE	CURRENT SPEED
<input type="checkbox"/>	ethernet1/1/1		Down	Enabled	0.00 Kb/s
<input type="checkbox"/>	ethernet1/1/2		Down	Enabled	0.00 Kb/s
<input type="checkbox"/>	ethernet1/1/3		Up	Enabled	25.00 Gb/s
<input type="checkbox"/>	ethernet1/1/4		Down	Enabled	0.00 Kb/s
<input type="checkbox"/>	ethernet1/1/5		Down	Enabled	0.00 Kb/s
<input type="checkbox"/>	ethernet1/1/6		Down	Enabled	0.00 Kb/s
<input type="checkbox"/>	ethernet1/1/7		Down	Enabled	0.00 Kb/s

Figure 27 Port information tab for the IOM

4. Set the port-group to **hardware default**.
5. Click the **port-group 1/1/11** check box, then click **Configure breakout**.

The screenshot shows the 'Configure Breakout' dialog box. It contains the following fields:

- I/O Module Name: IOM-B1
- Selected Ports: phy-port1/1/11
- Breakout Type: Select Breakout Type (dropdown menu)

A 'Cancel' button is located at the bottom right of the dialog box.

Figure 28 Configure Breakout dialog box



6. In **Configure breakout** panel, select **HardwareDefault** as the breakout type.



Configure Breakout

I/O Module Name IOM-B1

Selected Ports phy-port1/1/11

Breakout Type HardwareDefault

Finish Cancel

Figure 29 Set Breakout type to HardwareDefault

7. Click **Finish**.
8. To set port group 1/1/11 to **4X10GbE**, select the **port-group 1/1/11** check box, and then click **Configure breakout**.
9. In **Configure breakout** panel, select **4X10GbE** as the breakout type.



Configure Breakout

I/O Module Name IOM-B1

Selected Ports phy-port1/1/11

Breakout Type 4X10GE

Finish Cancel

Figure 30 Set desired Breakout Type

10. Click **Finish**. Be sure to repeat these steps for both MX5108n switches.

### 6.3.5 Define uplinks

After initial deployment, the new fabric shows **Uplink Count** as 'zero' and shows a warning (⚠️) icon. The lack of a fabric uplink results in a failed health check and displays the ❌ icon. To create uplinks, follow these steps:

1. Open the OME-M console.
2. From the navigation menu click **Devices**, then click **Fabric**.
3. Click on the **SmartFabric** fabric name.
4. In the **Fabric Details** panel, click **Uplinks**, then click the **Add Uplinks** button.
5. Using the information from Table 4, enter the **Uplink Name** in the **Name** box
6. Optionally, enter a description in the **Description** box.
7. From the **Uplink Type** list, select **FCoE** (as defined in Table 6)
8. Click **Next**
9. From the **Switch Ports** list, select the FCoE ports defined in Table 6.

- From the **Tagged Networks** list, select the VLAN defined in Table 6, then click **Finish**. SmartFabric creates the uplink object and the status for the fabric changes to **OK** (✓).

For the examples in this section, create the two uplinks using the information provided in Table 6.

Table 6 Uplink attributes

Uplink Name	Description	Type	Ports	VLAN (Tagged)
FCoE A1	FC Uplink for switch in Slot A1	FCoE	1/1/11:1 and 1/1/11:2	30
FCoE A2	FC Uplink for switch in Slot A2	FCoE	1/1/11:1 and 1/1/11:2	40

---

**Note:** Use the guidelines provided in chapter 8 to verify the SmartFabric configuration.

---

### 6.3.6 Create an FCoE server template

A server template contains the parameters extracted from a server and allows these parameters to be quickly applies to multiple compute sleds. The templates contain settings for the following categories:

- Local access configuration
- Location configuration
- Power configuration
- Chassis network configuration
- Slot configuration
- Setup configuration

Additionally, server templates also allow an administrator to associate VLANs to compute sleds.

To create a server template, follow these steps:

1. Open the OME-M console.
2. From the navigation menu, click **Configuration**, then click **Deploy**.
3. From the center pane, click **Create Template**, then click **From Reference Device**.
4. Enter **MX740c with FCOE CNA** in the **Template Name** box.

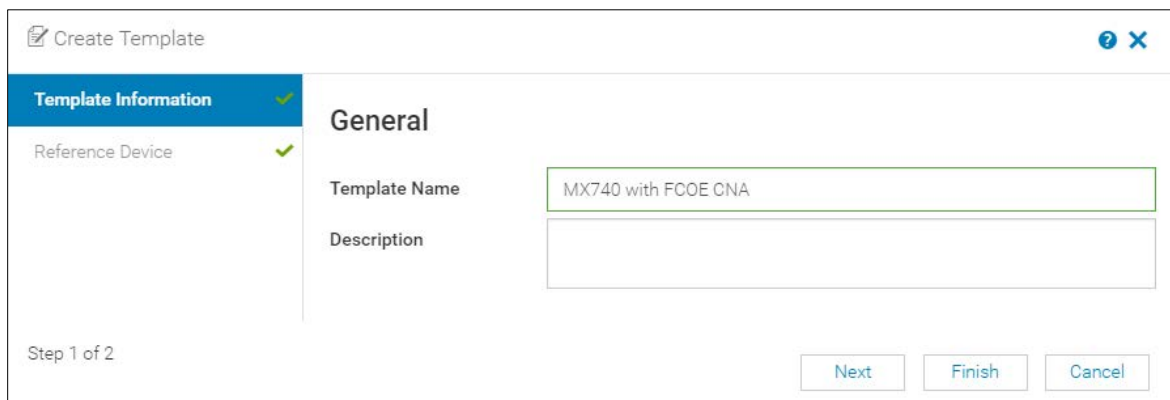


Figure 31 Create Template dialog box

5. Optionally, enter a description in the **Description** box.

6. Click **Next**.
7. In the **Device Selection** panel, click **Select Device**.
8. In the **Select Devices** window, choose the server previously configured.
9. In the **Select Devices** window, click **Finish**.
10. From the **Elements to Clone** list, select the **iDRAC**, **System**, and **NIC** options, then click **Finish**.

**Note:** iDRAC, System, and NIC are the minimum elements required for this example to function correctly. Other elements are optional.

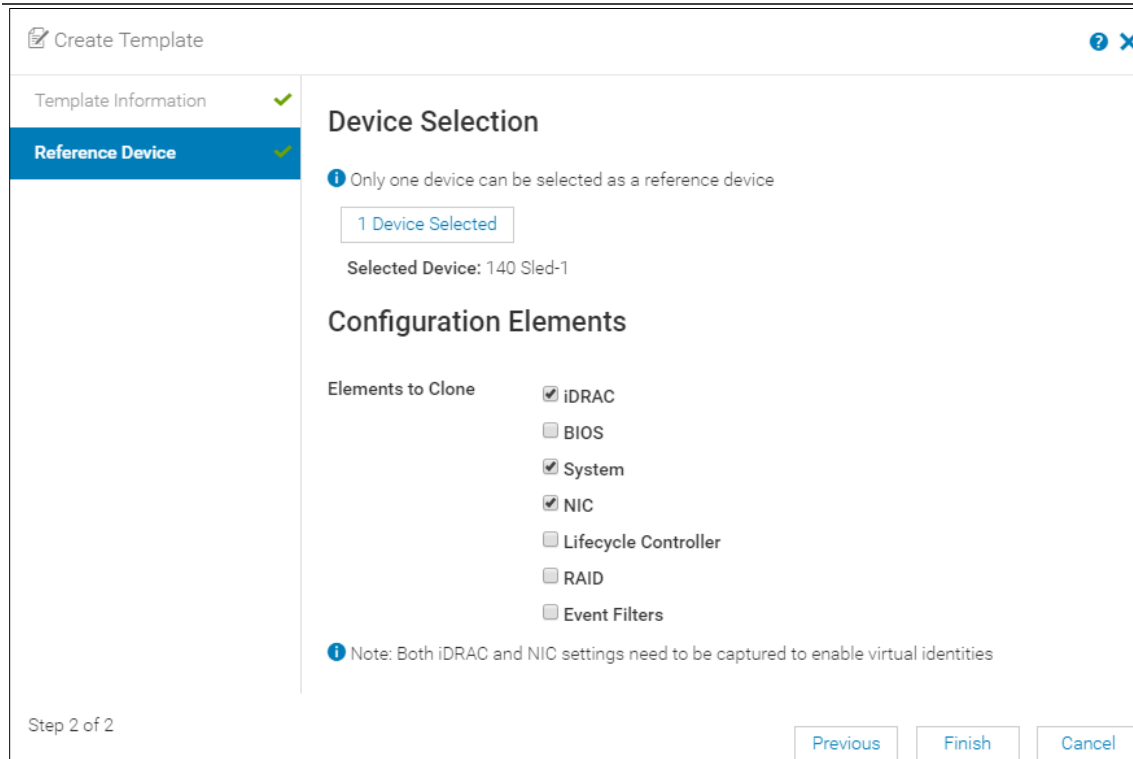


Figure 32 Select Device and elements to clone

A job starts and the new server template displays on the list. When complete, the status displays as **Completed successfully**. Next, associate the VLANs created previously with the template.

### 6.3.6.1 Associate server template with a VLAN and Identity Pool

After successfully creating a new template, associate the template with a network:

1. From the **Deploy** panel, select the **MX740c with FCOE CNA server** template.
2. From the **Deploy** panel, click **Edit Network** to open the **Edit Network** window.
3. From the **Identity Pool** list, choose **Ethernet CNA**.
4. For NIC in Mezzanine 1A Port 1, from the **Tagged Network** list, choose **FC A1**.
5. For NIC in Mezzanine 1A Port 2, from the **Tagged Network** list, choose **FC A2**.
6. Click **Finish**.

### 6.3.6.2 Deploy a server template

To deploy the server template, complete the following steps:

1. From the **Deploy** pane, select the **MX740c with FCOE CNA server** template.
2. From the **Deploy** pane, click **Deploy Template** to open the **Deploy Template** window.
3. Click the **Select** button to select the slots to deploy to the template.
4. Select the Do not forcefully reboot the host OS option.
5. Click **Next, Run Now**, then click **Finish**.

The interfaces on the MX9116n FSE are updated automatically. SmartFabric configures each interface with an untagged VLAN and any tagged VLANs. Additionally, SmartFabric deploys associated QoS settings (see Appendix A.5). On the server reboot or power on the server may run a configuration job to change settings to match the reference server.

---

**Note:** In this section, only Fabric A is used.

---

Your system is now ready to connect to Fibre Channel storage. See Appendix B to set up LUNs.

## 7 Scenario 3: Boot from SAN

An MX Server's host OS can boot from a remote FC storage array utilizing the IOMs. Booting to an OS through FC direct attach (F\_port), FC (NPG), and FCoE (FSB) scenarios is supported.

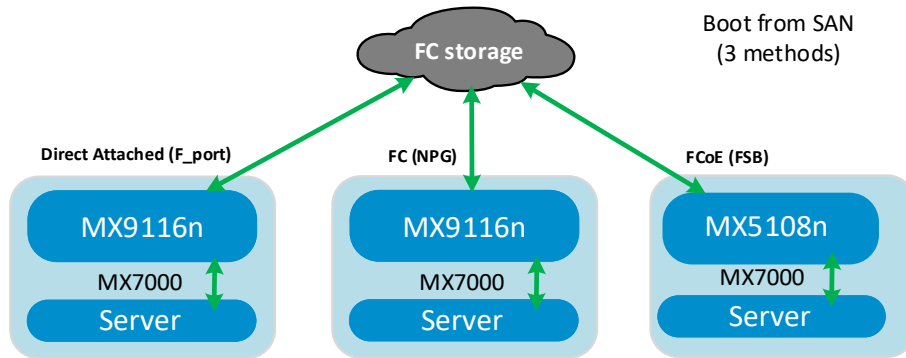


Figure 33 Boot from SAN

Figure 34 shows the example topology used in this chapter to demonstrate Boot from SAN. The required steps are provided to configure NIC partitioning, system BIOS, an FCoE LUN, and an OS install media device required for Boot from SAN.

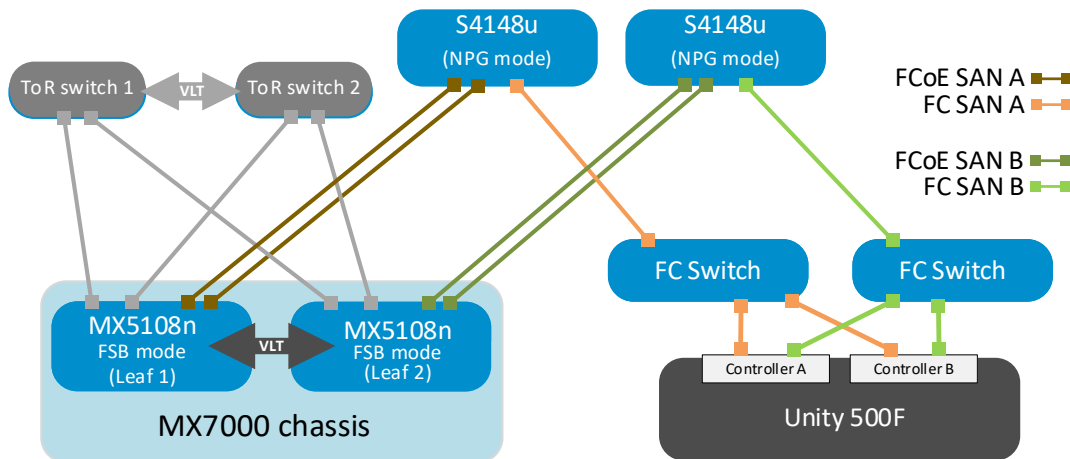


Figure 34 FCoE Boot from SAN

**Note:** See the [OS10 Enterprise Edition User Guide](#) document for configuring NPG mode globally on the S4148U-ON switches.

### 7.1 Configure NIC Boot Device

In this section, each QLogic CNA port is partitioned into one Ethernet and one FCoE partition.

**Note:** This is only done on CNA ports that carry converged traffic. In this example, these are the two 25GbE QLogic CNA ports on each server that attach to the switches internally through an orthogonal connection.

1. Prepare the server per chapter 4.

2. Connect to the server's iDRAC in a web browser and launch the virtual console.
3. In the virtual console, select **BIOS Setup** from the **Next Boot** menu.
4. Reboot the server.
5. On the System Setup Main Menu, select Device Settings.
6. Select the first CNA port.
7. Select FCoE Configuration.
8. Set the **Virtual LAN ID** (30 is used in this example).
9. Set Connect 1 to Enabled.
10. Set the **World Wide Port Name Target 1** to the connected port on **Unity** (see Appendix B).

Main Configuration Page • FCoE Configuration

FCoE General Parameters

Virtual LAN ID .....	30
Connect 1 .....	<input checked="" type="radio"/> Enabled <input type="radio"/> Disabled
World Wide Port Name Target 1 .....	50:06:01:61:47:E4:1B:19
Boot LUN 1 .....	0

Figure 35 FCoE configuration

11. Click **Back** and then click **Finish**.
12. When prompted, answer **Yes** to save changes and click **OK** in the **Success** window
13. Select the second CNA port and repeat the steps in this section for port 2.
14. Click **Finish** to exit to the System Setup Main Menu.

## 7.2 Configure BIOS settings

To allow Boot from SAN, perform the following steps in the system BIOS settings to disable the PXE BIOS.

1. Select System BIOS from the System Setup Main Menu.
2. Select Network Settings.
3. Click Disable for all PXE Devices.
4. Click **Back**.
5. Click **Finish**, click **Finish** again, then select **Yes** to exit and reboot.

---

**Note:** As discussed in previous chapters, this server configuration may be used to generate a template to deploy to other servers with identical hardware. When a template is not used, steps in this chapter should be repeated for each MX server sled requiring access to the FC storage.

---

## 7.3 Connect FCoE LUN

The server should be provisioned to connect to an FCoE boot LUN before moving on to Section 7.4. Follow the procedures in Chapter 5 or Chapter 6 to configure and connect to an FCoE LUN. Once connected, continue to the steps below to complete the Boot from SAN configuration.

## 7.4 Set up install media connection

**Note:** The iDRAC's Java Virtual console was used to complete the steps below.

1. Connect to the server's iDRAC in a web browser and launch the virtual console.
2. In the virtual console, from the **Virtual Media** menu, select **Virtual Media**.
3. In the virtual console, from the **Virtual Media** menu, select **Map CD/DVD**.
4. Click **Browse** to find the location of the OS install media then click **Map Device**.
5. In the virtual console, from the **Next Boot** menu, select **Lifecycle Controller**.
6. Reboot the server.

## 7.5 Set up OS driver install media using Lifecycle Controller

Some Operating System's install media do not contain the necessary FCoE drivers to boot from a FCoE LUN. Use this procedure to create an internal OS install media device. For VMware ESXi, refer to the the Dell customized media instructions located on the [Dell EMC Support website](#).

1. In Lifecycle Controller, select OS Deployment, then select Deploy OS.
2. From the **Select an Operating System** screen, verify that **Boot mode** is set to **UEFI**.
3. Select an OS to be installed to the boot LUN.

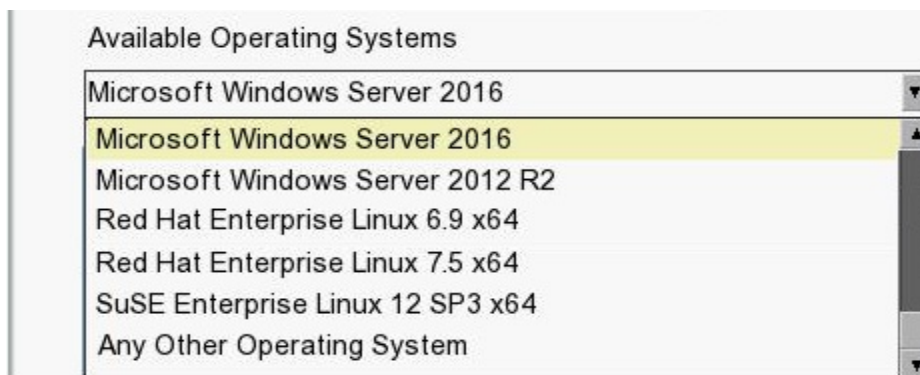


Figure 36 Lifecycle Controller OS deployment menu

4. Click Next.
5. Click the **Manual Install** check box, then click **Next**.
6. Click Next on the Insert OS Media Screen.
7. Click **Finish** when prompted on the **Reboot System** screen.
8. System reboots to Virtual Media. Press any key to boot install media when prompted.
9. Follow the OS prompts to install the OS to the FCoE storage LUN.

## 8 Verify configuration

This chapter covers the validation of the SmartFabric configuration.

### 8.1 Validate MX7000 setup using OME-M

This section covers validation specific to the Dell EMC PowerEdge MX7000 using the OME-M console.

#### 8.1.1 Show the MX chassis group topology

The OME-M console can be used to show the physical cabling of the SmartFabric.

1. Open the OME-M console.
2. In the left pane click **View Topology**.
3. Click the lead chassis and then click **Show Wiring**.
4. The blue check icons (✓) can be clicked to show cabling.

Figure 37 shows the current wiring of the SmartFabric.

The screenshot shows the OME-M console interface. The main area displays the 'Group Topology: SKY003Z' with two chassis: MX-SKY003Z and MX-SKY002Z. Each chassis has a port configuration table with columns for port numbers (17, 19, 21, 23, 25, 27, 29, 31, 33, 35, 37, 39, 41, 42, 43, 44) and a '2' column. Blue checkmarks are visible on the port configuration tables. The right pane shows 'Validation Errors (0)' and 'Shared Chassis (2)'. The shared chassis are MX-SKY003Z and MX-SKY002Z. For MX-SKY003Z, the IOM-A1 module (CBJXLN2 - Dell EMC MX9116n Fabric Engine) has three port-to-destination mappings: 1/1/39, 1/1/40 to SKY002Z, Slot IOM-A2: 1/1/39, 1/1/40; 1/1/37, 1/1/38 to SKY002Z, Slot IOM-A2: 1/1/37, 1/1/38; and 1/1/17, 1/1/18 to SKY002Z, Slot IOM-A1: UPLINK-1. The IOM-A2 module (110DXC2 - MX7116n Fabric Expander Module) has one mapping: UPLINK-1 to SKY002Z, Slot IOM-A2: 1/1/17, 1/1/18. For MX-SKY002Z, the IOM-A1 module (D10DXC2 - MX7116n Fabric Expander Module) has one mapping: UPLINK-1 to SKY003Z, Slot IOM-A1: 1/1/17, 1/1/18. The IOM-A2 module (F13RPK2 - Dell EMC MX9116n Fabric Engine) has three mappings: 1/1/17, 1/1/18 to SKY003Z, Slot IOM-A2: UPLINK-1; 1/1/39, 1/1/40 to SKY003Z, Slot IOM-A1: 1/1/39, 1/1/40; and 1/1/37, 1/1/38 to SKY003Z, Slot IOM-A1: 1/1/37, 1/1/38.

Figure 37 SmartFabric cabling



## 8.1.2 Show the SmartFabric status

The OME-M console can be used to show the overall health of the SmartFabric.

1. Open the OME-M console.
2. From the navigation menu, click **Devices**, then click **Fabric**.
3. Select **SmartFabric1** to expand the details of the fabric.

Figure 38 shows the details of the fabric.



Figure 38 Fabric status details

The **Overview** tab shows the current inventory, including switches, servers, and interconnects between the switches in the fabric (e.g. MX9116n FSEs). Figure 39 shows the SmartFabric switch in a healthy state.

The screenshot shows the 'Overview' tab of the SmartFabric switch inventory. The table has columns for 'HEALTH', 'POWER STATE', 'SERVICE TAG', 'CHASSIS', 'SLOT', and 'MODEL'. There are two rows of switch data.

	HEALTH	POWER STATE	SERVICE TAG	CHASSIS	SLOT	MODEL
Switches	Ok	On	CBJXLN2	MX-SKY003Z	IOM-A1	Dell EMC MX9116n Fabric Engine
Servers	Ok	On	F13RPK2	MX-SKY002Z	IOM-A2	Dell EMC MX9116n Fabric Engine

Figure 39 SmartFabric switch inventory

Figure 40 shows the participating servers in a healthy state.

Overview		Topology						
Uplinks		<b>Servers</b>						
Switches		HEALTH	POWER STATE	NAME	SERVICE TAG	CHASSIS	SLOT	MODEL
Servers		✓ Ok	On	Sled-1	CF52XM2	MX-SKY002Z	Sled-1	PowerEdge MX740c
ISL Links		✓ Ok	On	Sled-2	1S35MN2	MX-SKY003Z	Sled-2	PowerEdge MX740c
		✓ Ok	On	Sled-1	CBMP9N2	MX-SKY003Z	Sled-1	PowerEdge MX740c
		✓ Ok	On	Sled-2	1S34MN2	MX-SKY002Z	Sled-2	PowerEdge MX740c

Figure 40 SmartFabric server inventory

Figure 41 shows the **Topology** tab and the VLTi between two MX9116n FSEs automatically created by SmartFabric mode. If using MX5108n Ethernet switches, the picture will be similar, and depict the VLTi connection between those switches.

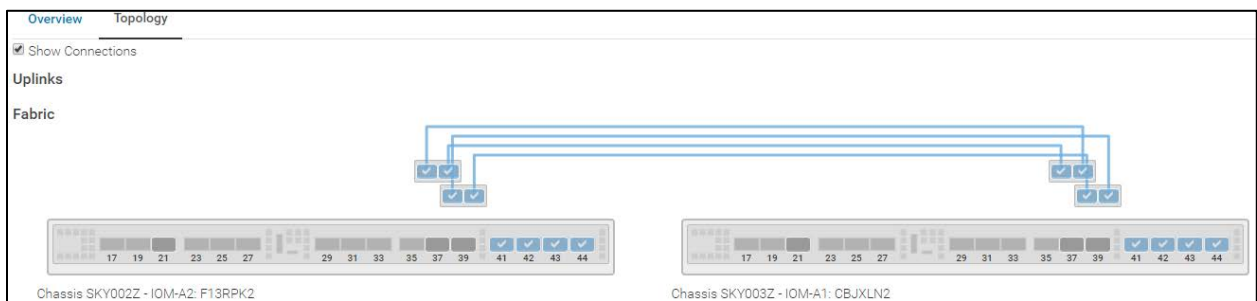


Figure 41 SmartFabric overview fabric diagram

Figure 42 displays the wiring diagram table from the **Topology** tab.

Wiring Diagram						
^ Chassis SKY002Z						
^ IOM-A2: F13RPK2 - Dell EMC MX9116n Fabric Engine						
PORT NUMBER	OPERATIONAL STATUS	PORT CONFIGURATION	PORT ROLE	UPLINK NAME	DESTINATION	
ethernet1/1/37	Up	NoBreakout	ISL		SKY003Z, Slot IOM-A1: ethernet1/1/37	
ethernet1/1/39	Up	NoBreakout	ISL		SKY003Z, Slot IOM-A1: ethernet1/1/39	
ethernet1/1/40	Up	NoBreakout	ISL		SKY003Z, Slot IOM-A1: ethernet1/1/40	
ethernet1/1/38	Up	NoBreakout	ISL		SKY003Z, Slot IOM-A1: ethernet1/1/38	
^ Chassis SKY003Z						
v IOM-A1: CBJXLN2 - Dell EMC MX9116n Fabric Engine						

Figure 42 SmartFabric topology wiring diagram table

## 8.1.3 Show port status

The OME-M console can be used to show MX9116n FSE port status.

1. Open the OME-M console.
2. From the navigation menu, click **Devices**, then click **I/O Modules**.
3. Select an IOM and click the **View Details** button to the right of the inventory screen. The **IOM overview** for that device is displayed.
4. From the IOM Overview, click **Hardware**.
5. Click the Port Information tab.

Figure 43 shows ethernet 1/1/1, 1/1/3, 1/71/1, and 1/71/2 in the correct operational status (Up). These interfaces correspond to the MX740c compute sleds in slots 1 and 2 in both chassis. The figure also shows the VLTi connection (port-channel1000) and the uplinks (port-channel1) to the leaf switches.

PORT NUMBER	P...	P...	OPERATIONAL STATUS	ADMIN STATE	CURRENT SPEED	PORT CONFIGURATION	OPTICS TYPE	MTU SIZE	AUTO NEGOTI...	PORT ROLE
ethernet1/1/1			Up	Enabled	25.00 Gb/s		Fixed	1532	Enabled	EndHost
ethernet1/1/3			Up	Enabled	25.00 Gb/s		Fixed	1532	Enabled	EndHost
port-channel1			Up	Enabled	200.00 Gb/s			1532	Disabled	Uplink
ethernet1/1/42			Up	Enabled	100.00 Gb/s		QSFP28	1532	Disabled	Uplink
ethernet1/1/41			Up	Enabled	100.00 Gb/s		QSFP28	1532	Disabled	Uplink
port-channel1000			Up	Enabled	400.00 Gb/s			9216	Disabled	ISL
ethernet1/1/37			Up	Enabled	100.00 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/39			Up	Enabled	100.00 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/40			Up	Enabled	100.00 Gb/s		QSFP28-DD	9216	Enabled	ISL
ethernet1/1/38			Up	Enabled	100.00 Gb/s		QSFP28-DD	9216	Enabled	ISL
port-group1/1/1					0.00 Kb/s	FabricExpander				
ethernet1/71/2			Up	Enabled	25.00 Gb/s			1532	Enabled	EndHost
ethernet1/71/1			Up	Enabled	25.00 Gb/s			1532	Enabled	EndHost

Figure 43 IOM Port Information

## 8.2 Switch CLI validation

After connected devices are configured, CLI commands are available to validate the network configuration. Use the IP address of each IOM (see Figure 43) to establish an SSH connection for entering CLI commands. This section provides a list of the most common commands and their output for this topology.

### 8.2.1 show switch-operating-mode

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

```
C140A1# show switch-operating-mode
```

Switch-Operating-Mode : Smart Fabric Mode

## 8.2.2 show discovered-expanders

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

```
C140A1# show discovered-expanders
Service   Model   Type   Chassis   Chassis-slot   Port-group   Virtual
tag                               service-tag                               Slot-Id
-----
D10DXC2   MX7116n  1      SKY002Z   A1              1/1/1       71
          FEM
```

## 8.2.3 show unit-provision

The `show unit-provision` command is only available on the MX9116n FSE and displays the unit ID and server tag of the MX7116n FEM attached to the MX9116n FSE.

```
C140A1# show unit-provision
Node ID | Unit ID | Provision Name | Discovered Name | State |
-----+-----+-----+-----+-----+
1       | 71     | D10DXC2       | D10DXC2         | up   |
```

## 8.2.4 show lldp neighbors

The `show lldp neighbors` command shows information about LLDP neighbors. `Ethernet1/1/1`, `ethernet 1/1/3`, and `ethernet 1/1/71-1/1/72` represent the two MX740c compute sleds. The first entry is the iDRAC for the compute sled. The iDRAC uses connectivity to the mezzanine card to advertise LLDP information. The second entry is the mezzanine card itself.

Ports `ethernet1/71/1` and `ethernet1/71/2` represent the MX7116n FEM connections to the MX740c. Ethernet range `ethernet1/1/37-1/1/40` are the VLTi interfaces for the SmartFabric.

```
C140A1# show lldp neighbors
Loc PortID      Rem Host Name      Rem Port Id      Rem Chassis Id
-----
ethernet1/1/1  iDRAC-CBMP9N2     CBMP9N2 NIC.Mezzanine.1A-1-1  d0:94:66:2a:07:2f
ethernet1/1/1  Not Advertised    24:6e:96:9c:e3:50  24:6e:96:9c:e3:50
ethernet1/1/3  iDRAC-1S35MN2     1S35MN2 NIC.Mezzanine.1A-1-1  d0:94:66:29:fa:f4
ethernet1/1/3  Not Advertised    24:6e:96:9c:e5:48  24:6e:96:9c:e5:48
ethernet1/1/37 C160A2            ethernet1/1/37    20:04:0f:00:a1:9e
ethernet1/1/38 C160A2            ethernet1/1/38    20:04:0f:00:a1:9e
ethernet1/1/39 C160A2            ethernet1/1/39    20:04:0f:00:a1:9e
ethernet1/1/40 C160A2            ethernet1/1/40    20:04:0f:00:a1:9e
ethernet1/71/1 Not Advertised    24:6e:96:9c:e5:d8  24:6e:96:9c:e5:d8
ethernet1/71/1 iDRAC-CF52XM2     CF52XM2 NIC.Mezzanine.1A-1-1  d0:94:66:29:fe:b4
ethernet1/71/2 Not Advertised    24:6e:96:9c:e5:da  24:6e:96:9c:e5:da
```

```
ethernet1/71/2 iDRAC-1S34MN2 1S34MN2 NIC.Mezzanine.1A-1-1 d0:94:66:29:ff:27
```

## 8.2.5 show vlan

Port VLAN assignment is verified with the `show vlan` command. Po1000 is the VLTi, and it is automatically created when the SmartFabric is created. VLAN 4094 is reserved and is the internal control VLAN for the VLT domain.

```
C140A1# show vlan
```

```
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs
```

```
Q: A - Access (Untagged), T - Tagged
```

	NUM	Status	Description	Q Ports
*	1	Active		T Po1000
				A Eth1/1/1,1/71/1-1/71/2
	30	Active		T Eth1/1/1,1/71/1-1/71/2
	40	Active		T Po1000
	4004	Active		T Mgmt1/1/1
M	4020	Active		A Mgmt1/1/1
	4094	Active		T Po1000

## 8.2.6 show vlt domain id

The `show vlt domain id` command validates the VLT configuration status when the VLTi Link Status is up. The role of one switch in the VLT pair is primary (shown) and its peer switch is assigned the secondary role.

```
C140A1# show vlt 255
```

```
Domain ID           : 1
Unit ID             : 1
Role                 : primary
Version             : 1.0
Local System MAC address : 4c:76:25:e8:f2:c0
VLT MAC address     : 4c:76:25:e8:f2:c0
IP address           : fda5:74c8:b79e:1::1
Delay-Restore timer  : 90 seconds
Peer-Routing         : Disabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
  port-channel1000   : up
```

VLT Peer Unit ID	System MAC Address	Status	IP Address	Version
2	4c:76:25:e8:e8:40	up	fda5:74c8:b79e:1::2	1.0

## 8.2.7 show spanning-tree brief

The `show spanning-tree brief` command validates that STP is enabled on the leaf switches. All interfaces are forwarding (FWD), as shown in the **Sts** column.

```

C140A1# show spanning-tree brief
Spanning tree enabled protocol rapid-pvst
VLAN 1
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32769, Address 2004.0f00.a19e
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID   Priority 32769, Address 2004.0f00.cd1e
Configured hello time 2, max age 20, forward delay 15
Flush Interval 200 centi-sec, Flush Invocations 3
Flush Indication threshold 0 (MAC flush optimization is disabled)
Interface
Name          PortID  Prio  Cost  Sts  Cost  Designated Bridge ID          PortID
-----
ethernet1/1/1 128.260 128   800   FWD  1     32769   2004.0f00.cd1e 128.260

Interface
Name          Role    PortID  Prio  Cost  Sts  Cost  Link-type  Edge
-----
ethernet1/1/1 Desg    128.260 128   800   FWD  1     AUTO   No
VLAN 30
Executing IEEE compatible Spanning Tree Protocol
Root ID    Priority 32798, Address 2004.0f00.cd1e
Root Bridge hello time 2, max age 20, forward delay 15
Bridge ID   Priority 32798, Address 2004.0f00.cd1e
We are the root of VLAN 30
Configured hello time 2, max age 20, forward delay 15
Flush Interval 200 centi-sec, Flush Invocations 1
Flush Indication threshold 0 (MAC flush optimization is disabled)
Interface
Name          PortID  Prio  Cost  Sts  Cost  Designated Bridge ID          PortID
-----
ethernet1/1/1 128.260 128   800   FWD  0     32798   2004.0f00.cd1e 128.260

Interface
Name          Role    PortID  Prio  Cost  Sts  Cost  Link-type  Edge
-----
ethernet1/1/1 Desg    128.260 128   800   FWD  1     AUTO   No
:
<output truncated>

```

## 8.2.8 show lldp dcbx interface ethernet *interface\_number*

The `show lldp dcbx interface ethernet interface_number` command is used to verify DCBX is operational and in sync with the connected CNA. The local DCBX TLVs Transmitted field is `ErPFFi`, confirming ETS, PFC, and FCoE TLVs are transmitted to the CNA. Refer to the table at the top of the command output.

---

**Note:** TLV stands for type, length, and value. Groups of type, length, and value fields, transmitted in Link Layer Discovery Protocol (LLDP) frames, are commonly referred to as TLVs.

---

The ETS TLV transmits the bandwidth allocation, the PFC TLV enables priority flow control, and the FCoE TLV assigns the dot1p priority value to FCoE traffic.

```
S4148U-Leaf1# show lldp dcbx interface ethernet 1/1/1
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled         r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled         p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled  f-Application Priority for FCOE
disabled
I-Application priority for iSCSI enabled  i-Application Priority for iSCSI
disabled
```

---

```
Interface ethernet1/1/1
  Port Role is Manual
  DCBX Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBX Compatibility mode is IEEEv2.5
  Local DCBX Configured mode is AUTO
  Peer Operating version is IEEEv2.5
  Local DCBX TLVs Transmitted: ERPfi
  21209 Input PFC TLV pkts, 126 Output PFC TLV pkts, 0 Error PFC pkts
  21201 Input ETS Conf TLV Pkts, 173 Output ETS Conf TLV Pkts, 7 Error ETS
Conf TLV Pkts
  0 Input ETS Reco TLV pkts, 173 Output ETS Reco TLV pkts, 0 Error ETS Reco
TLV Pkts
  21208 Input Appln Priority TLV pkts, 126 Output Appln Priority TLV pkts, 0
Error Appln Priority TLV Pkts

Total DCBX Frames transmitted 598
Total DCBX Frames received 63618
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0
```

### 8.2.9 `show lldp dcbx interface ethernet interface_number pfc detail`

The `show lldp dcbx interface ethernet interface_number pfc detail` command is used to verify PFC is enabled on dot1p priority 3 traffic, and its status is operational. It shows the FCoE TLV is enabled and the FCoE priority map is set to 0x08, which maps to dot1p priority 3. (Hex 08 is binary 1000. Counting bits from right to left and starting at 0, 1000 represents priority 3.)

```
C140A1# show lldp dcbx interface ethernet 1/1/1 pfc detail
Interface ethernet1/1/1
  Admin mode is on
  Admin is enabled, Priority list is 3
  Remote is enabled, Priority list is 3
  Remote Willing Status is enabled
```

```

Local is enabled, Priority list is 3
Oper status is init
PFC DCBX Oper status is Up
State Machine Type is Symmetric
PFC TLV Tx Status is enabled
Application Priority TLV Parameters :
-----
ISCSI TLV Tx Status is disabled
Remote FCOE PriorityMap is 0x08

21215 Input TLV pkts, 126 Output TLV pkts, 0 Error pkts
21214 Input Appln Priority TLV pkts, 126 Output Appln Priority TLV pkts, 0
Error Appln Priority TLV Pkts

```

### 8.2.10 show lldp dcbx interface ethernet *interface\_number* ets detail

The `show lldp dcbx interface ethernet interface_number ets detail` command is used to verify configured ETS bandwidth settings are applied in the correct percentages to the correct priority numbers, and ETS status is operational.

```
C140A1# show lldp dcbx interface ethernet 1/1/1 ets detail
```

```

Interface ethernet1/1/1
Max Supported PG is 8
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
-----
Admin is enabled

PG-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,5,6,7    98%            ETS
1           1               0%            SP
2           2               0%            SP
3           3               1%            ETS
4           4               1%            ETS
5           5               0%            SP
6           6               0%            SP
7           7               0%            SP

```

```

Remote Parameters :
-----
Remote is enabled

PG-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,5,6,7    98%            ETS
1           1               0%            SP
2           2               0%            SP
3           3               1%            ETS

```



```

4          4                      1%          ETS
5          4                      0%          SP
6          4                      0%          SP
7          4                      0%          SP

```

Remote Willing Status is enabled

Local Parameters :

-----

Local is enabled

PG-grp	Priority#	Bandwidth	TSA
0	0,1,2,5,6,7	98%	ETS
1		0%	SP
2		0%	SP
3	3	1%	ETS
4	4	1%	ETS
5		0%	SP
6		0%	SP
7		0%	SP

Oper status is init

ETS DCBX Oper status is Up

State Machine Type is Asymmetric

Conf TLV Tx Status is enabled

Reco TLV Tx Status is enabled

21210 Input Conf TLV Pkts, 173 Output Conf TLV Pkts, 7 Error Conf TLV Pkts

0 Input Reco TLV Pkts, 173 Output Reco TLV Pkts, 0 Error Reco TLV Pkts

## 8.2.11 show qos system

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

```

C140A1# show qos system
Service-policy (input): PM_VLAN
ETS Mode : off

```

## 8.2.12 show policy-map

Using the service policy from `show qos system`, the `show policy-map type qos PM_VLAN` command displays QoS policy details including associated class maps, for example, `CM1`, and QoS queue settings, `qos-group 2`.

```

C140A1# show policy-map
Service-policy (qos) input: PM_VLAN
Class-map (qos): CM1
set qos-group 2

```

### 8.2.13 show qos interface ethernet *interface\_number*

The `show qos interface ethernet interface_number` command is used to confirm ETS and PFC are both enabled on the interface.

```
C140A1# show qos interface ethernet 1/1/1
Interface
unknown-unicast-storm-control : Disabled
multicast-storm-control : Disabled
broadcast-storm-control : Disabled
flow-control-rx : Disabled
flow-control-tx : Disabled
ets mode : Enabled
pfc mode : Enabled
```

### 8.2.14 show class-map

The command `show class-map` displays details for all the configured class-maps. For example, the association between CM1 and VLAN 1 is shown.

```
C140A1# show class-map
Class-map (application): class-iscsi
Class-map (qos): class-trust
Class-map (qos): CM1(match-any)
Match: mac vlan 1
```

### 8.2.15 show fc ns switch

The `show fc ns switch` command shows all device ports logged into the fabric. In this deployment, four ports are logged in to each switch: two storage ports and two CNA ports.

```
C140A1# show fc ns switch

Total number of devices = 3
Switch Name                10:00:20:04:0f:00:cd:1e
Domain Id                  1
Switch Port                fibrechannel1/1/44:1
FC-Id                      01:00:00
Port Name                  50:06:01:61:47:e4:1b:19
Node Name                  50:06:01:60:c7:e0:1b:19
Class of Service           8
Symbolic Port Name        UNITY::::SPA13::FC::::::
Symbolic Node Name        UNITY::::SPA::FC::::::
Port Type                  N_PORT
Registered with NameServer Yes
Registered for SCN        Yes

Switch Name                10:00:20:04:0f:00:cd:1e
Domain Id                  1
Switch Port                ethernet1/71/1
FC-Id                      01:01:00
```

```

Port Name                20:01:06:c3:f9:a4:cd:03
Node Name                20:00:06:c3:f9:a4:cd:03
Class of Service        8
Symbolic Port Name
Symbolic Node Name
Port Type                N_PORT
Registered with NameServer Yes
Registered for SCN      Yes

Switch Name             10:00:20:04:0f:00:cd:1e
Domain Id               1
Switch Port            ethernet1/1/1
FC-Id                  01:02:00
Port Name              20:01:f4:e9:d4:73:d0:0c
Node Name              20:00:f4:e9:d4:73:d0:0c
Class of Service        8
Symbolic Port Name      QLogic qedf v8.24.8.0
Symbolic Node Name      QLogic qedf v8.24.8.0
Port Type                N_PORT
Registered with NameServer Yes
Registered for SCN      Yes

```

## 8.2.16 show fcoe sessions

The `show fcoe sessions` command shows active FCoE sessions. The output includes MAC addresses, Ethernet interfaces, the FCoE VLAN ID, FC IDs, and WWPNs of logged-in CNAs.

---

**Note:** Due to the width of the command output, each line of output is shown on two lines below.

---

```

C140A1# show fcoe sessions
Enode MAC           Enode Interface   FCF MAC           FCF interface     VLAN
FCoE MAC           FC-ID    PORT WWPN           PORT WWNN
-----
06:c3:f9:a4:cd:03   Eth 1/71/1        20:04:0f:00:ce:1d ~                   30
0e:fc:00:01:01:00   01:01:00  20:01:06:c3:f9:a4:cd:03  20:00:06:c3:f9:a4:cd:03
f4:e9:d4:73:d0:0c   Eth 1/1/1        20:04:0f:00:ce:1d ~                   30
0e:fc:00:01:02:00   01:02:00  20:01:f4:e9:d4:73:d0:0c  20:00:f4:e9:d4:73:d0:0c

```

## 8.2.17 show vfabric

The `show vfabric` command output provides a variety of information including the default zone mode, the active zone set, and interfaces that are members of the vfabric.

```

C140A1# show vfabric
Fabric Name          New vfabric
Fabric Type          FPORT
Fabric Id            1

```

```

Vlan Id          30
FC-MAP          0xEFC00
Vlan priority    3
FCF Priority     128
FKA-Adv-Period  Enabled,8
Config-State    ACTIVE
Oper-State      UP
=====
Switch Config Parameters
=====
Domain ID       1
=====
Switch Zoning Parameters
=====
Default Zone Mode:  Allow
Active ZoneSet:    None
=====
Members
fibrechannell1/1/44:1
ethernet1/1/1
ethernet1/71/1
ethernet1/71/2
=====

```

## 8.2.18 show fc switch

The `show fc switch` command verifies the switch mode (e.g. F\_Port) for FC traffic.

```

C140A1# show fc switch
Switch Mode : FPORT
Switch WWN  : 10:00:e4:f0:04:6b:04:42

```

## A Additional information

This section contains additional details to clarify the concepts covered in the main body of the document.

### A.1 Factory default PowerEdge MX7000

This section contains details on factory defaulting a PowerEdge MX7000 using SmartFabric mode IOMs.

#### A.1.1 Remove the SmartFabric

To remove the SmartFabric using the OME-M console, perform the following steps:

1. Open the OME-M console
2. From the navigation menu click **Devices** → **Fabric**
3. Select SmartFabric
4. In the **Fabric** pane, click **Delete Fabric**
5. In the **Delete Fabric** dialog box click **Yes**

All participating switches reboot to FullSwitch mode.

---

**Note:** Any configuration not completed by the OME-M console is lost when switching between IOM operating modes.

---

#### A.1.2 Use RACADM to reset each chassis

To reset the chassis to factory default:

1. Connect to the MM IP address using SSH.
2. In the RACADM shell type the `racadm racresetcfg` command. The Factory reset process is initiated and a status message displays.

---

**Note:** The process takes several minutes to complete. Once complete, use the LCD screen or other method to assign a static IP. See the [Dell EMC PowerEdge MX SmartFabric Network Deployment Guide](#) for information on assigning an IP address.

---

## A.2 Factory default OS10EE

To reset OS10EE switches back to the factory default configuration, enter the following commands:

```
C140A1#delete startup-configuration
  Proceed to delete startup-configuration [yes/no(default)]:yes
C140A1#reload

System configuration has been modified. Save? [yes/no]:no
Proceed to reboot the system? [confirm yes/no]:yes
```

The switch reboots with default configuration settings.

## A.3 Spanning Tree Protocol recommendations

By default, OS10EE uses Rapid per-VLAN Spanning Tree Plus (RPVST+) across all switching platforms including MX-series networking IOMs. OS10EE also supports RSTP and Multiple Spanning Tree (MST). However, MST is not currently supported with VLT.

RPVST+ works fine when the number of VLANs is small, usually up to 48 VLANs with the below configuration:

- Spanning-tree hello interval = 10 Seconds
- MAC Flush threshold = 5
- MAC Flush timer = 500 centi-seconds

In cases where users need more than 48 VLANs, Dell EMC recommends using RSTP. This can be configured in either SmartFabric or Full Switch mode.

Caution should be taken when connecting an RPVST+ to an existing RSTP environment. RPVST+ creates a single topology per VLAN with the default VLAN, typically VLAN 1, for the Common Spanning Tree (CST) with RSTP.

For non-native VLANs, all Bridge Protocol Data Unit (BPDU) traffic is tagged and forwarded by the upstream, RSTP-enable switch, with the associated VLAN. These BPDUs use a protocol-specific multicast address. Any other RPVST+ tree attached to the RSTP tree might processes these packets accordingly leading to the potential of unexpected trees.

---

**Note:** Dell EMC Networking recommends, that when connecting to an existing environment that is not using RPVST+, to change to the existing STP protocol before connecting an OS10EE switch.

---

In the example below, RSTP is enabled globally on an MX9116n FSE.

```
C140A1(config)# spanning-tree mode rstp
C140A1(config)# end
C140A1#show spanning-tree brief
```

```
Spanning tree enabled protocol rstp with force-version rstp
Executing IEEE compatible Spanning Tree Protocol
Root ID      Priority 0, Address 4c76.25e8.f2c0
Root Bridge hello time 2, max age 20, forward delay 15
```

```
Bridge ID      Priority 32768, Address 2004.0f00.cd1e
Configured hello time 2, max age 20, forward delay 15
Flush Interval 200 centi-sec, Flush Invocations 95
Flush Indication threshold 0 (MAC flush optimization is disabled)
```

## A.4 QSFP28 double density connectors

Quad Small Form-Factor Pluggable 28 Double Density, or QSFP28-DD connectors, expand on the QSFP28 pluggable form factor. By doubling the number of available lanes from four to eight, with each lane operating at 25 Gbps, the result is 200 Gbps for each connection.

---

**Note:** A QSFP28-DD transceiver is not compatible with a QSFP28 port due to the specifications required to lengthen the PCB connector to allow for the additional four lanes. However, a QSFP28 transceiver can be inserted into a QSFP28-DD port.

---




Figure 44 QSFP28-DD and QSFP28 physical interfaces

The MX9116n FSE supports direct attach cables (DAC), active optic cables (AOC), as well as multi-mode fiber (MMF) cables with supported Dell EMC optics.

## A.5 VLAN management and automated QoS

In addition to being able to assign VLANs to server profiles, SmartFabric automates QoS settings based on the input. Figure 45 shows that when defining a VLAN, one of 11 options are pre-defined. Each of these options represents a queue.

Figure 45 QoS options available in SmartFabric mode

Table 7 lists the traffic types and related settings. The QoS group is the numerical value for the queues available in SmartFabric mode. Available queues include 2 through 5 with queues 1, 6, and 7 in reserve.

**Note:** In SmartFabric mode, an administrator cannot change the default weights for the queues.

Table 7 Traffic types and default QoS settings

Traffic Type	Description	QoS Group
General Purpose (Bronze)	Used for low priority data traffic	2
General Purpose (Silver)	Used for standard/default priority data traffic	3
General Purpose (Gold)	Used for high priority data traffic	4
General Purpose (Platinum)	Used for extremely high priority data traffic	5
Cluster Interconnect	Used for cluster heartbeat VLANs	5
Hypervisor Management	Used for hypervisor management connections such as the ESXi management VLAN	5
Storage - iSCSI	Used for iSCSI VLANs	5
Storage - FCoE	Used for FCoE VLANs	5



Storage - Data Replication	Used for VLANs supporting storage data replication such as for VMware VSAN	5
VM Migration	Used for VLANs supporting vMotion and similar technologies	5
VMware FT Logging	Used for VLANs supporting VMware Fault Tolerance	5

When a VLAN-capable server template deploys, SmartFabric creates a class map. For example, class map CM10, matching all traffic associated with VLAN 10. Then a policy map, for example, PM\_VLAN, sets this class map to the appropriate queue, as in `qos-group 2`.

## A.6 Rate limited 32 Gb Fibre Channel

When using 32 Gb FC, the actual data rate is 28 Gbps due to 64b/66b encoding. Figure 46 shows the unified port group 15. The port group is set to 4x 32 Gb FC mode. However, each of the four lanes is 25 Gbps, not 28 Gbps. When these lanes are mapped from the Network Processing Unit (NPU) to the FC ASIC, for conversion to FC signaling, the 32GFC interfaces are mapped to four 25 Gbps lanes. With each lane operating at 25 Gbps, not 28 Gbps, the result is rate limited 32 Gb FC.

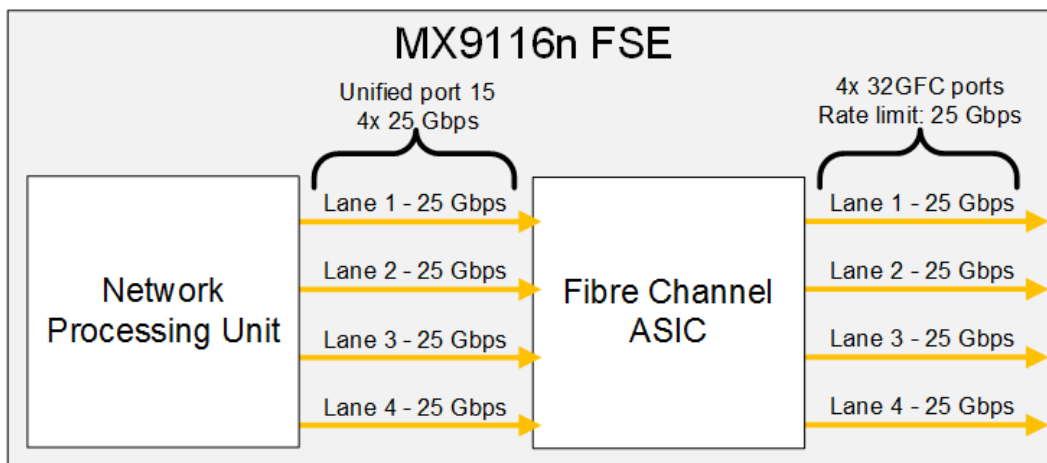


Figure 46 4x 32 Gb FC breakout mode, rate limit of 25 Gbps

While each 32 Gb FC connection provides 25 Gbps, the overall FC bandwidth available is 100 Gbps per unified port group, or 200 Gbps for both ports. However, if an application requires the maximum 28 Gbps throughput per port, use the 2x 32 Gb breakout mode. This mode configures the connections between the NPU and the FC ASIC as shown Figure 47.

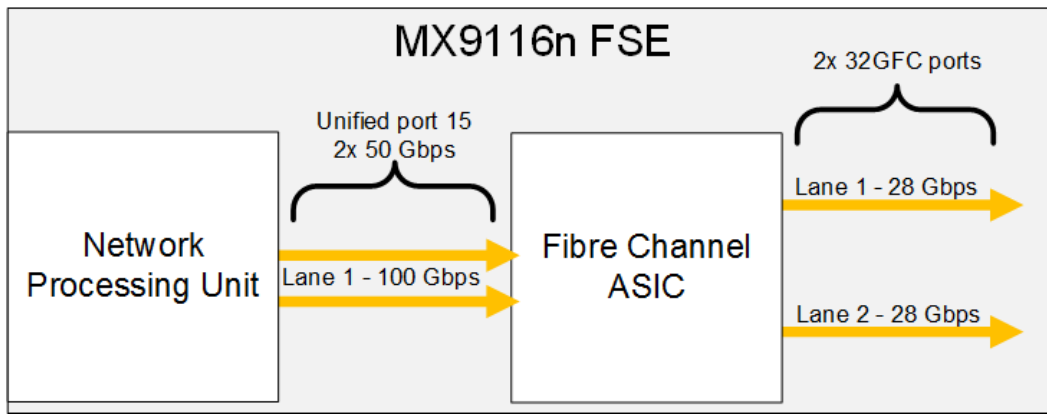


Figure 47 2x 32 Gb FC breakout mode

In 2x 32 Gb FC breakout mode, the MX9116n FSE binds two 50 Gbps links together to provide a total of 100 Gbps bandwidth per lane to the FC ASIC. This results in the two FC ports operating at 28 Gbps. The overall FC bandwidth available is 56 Gbps per unified port, or 112 Gbps for both. Compared to the 200 Gbps using 4x 32 Gb FC.

**Note:** Rate limited ports are not oversubscribed ports. There is no FC frame drop on these ports and buffer to buffer credit exchanges ensure flow consistency.

## A.7 Matching FC uplink port speed

SmartFabric does not support setting the Fibre Channel port speed of IOM uplinks. The upstream FC device must be able to match the speed of the FC optics used in the MX9116n FSE. For example, if the MX9116n FSE is using 64G optics, broken out as 4x16G FC, the upstream device must also support 16G.

## A.8 Fibre Channel optics for MX9116n

The MX9116n FSE supports the use of two Fibre Channel QSFP transceivers on unified ports 43 and 44. Table 8 shows the two transceiver options that are available: 4x16GFC and 4x32GFC.

Table 8 Supported GFC transceivers for MX9116n

Transceiver	Description
4x16GFC	Dell Networking Transceiver, 16G QSFP+ SWL Fibre Channel QSFP (4x16GFC, Supports 8/16 GFC)
4x32GFC	Dell Networking Transceiver, 32G QSFP28 SWL Fibre Channel QSFP (4x32GFC, Supports 8/16/32 GFC)


**Note:** The above listed optics in Table 8 are not supported in the MXG610s FC IOM.

## B Storage connection information

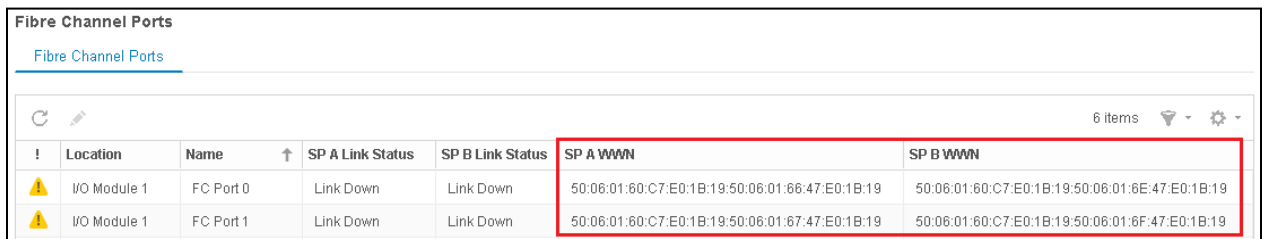
This section shows how an administrator can determine the WWPNs of CNAs and storage targets. The WWPNs are used to connect FC storage targets to specific servers for file storage or OS boot.

### B.1 Determine Unity 500F storage array FC WWPNs

The WWPNs of FC adapters in storage arrays are also used for FC configuration. WWPNs on Unity storage arrays are determined as follows:

1. Connect to the Unisphere GUI in a web browser and log in.
2. Click the **Settings** icon  near the top-right corner of the page.
3. In the left pane of the **Settings** window, select **Access**, then **Fibre Channel**.

The Fibre Channel Ports page is displayed as shown in Figure 48. A zoomed-in view of the area inside the red box is shown in Figure 49.





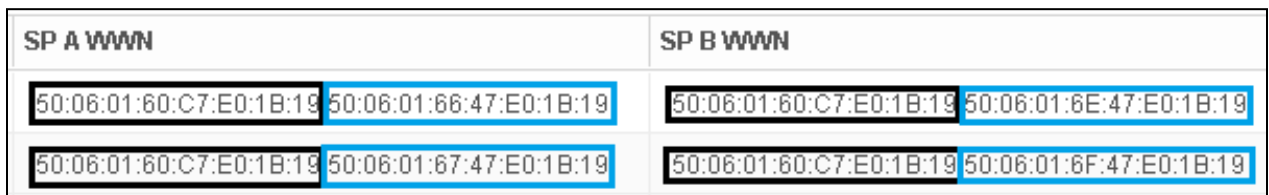
!	Location	Name	↑	SP A Link Status	SP B Link Status	SP A WWN	SP B WWN
	I/O Module 1	FC Port 0		Link Down	Link Down	50:06:01:60:C7:E0:1B:19 50:06:01:66:47:E0:1B:19	50:06:01:60:C7:E0:1B:19 50:06:01:6E:47:E0:1B:19
	I/O Module 1	FC Port 1		Link Down	Link Down	50:06:01:60:C7:E0:1B:19 50:06:01:67:47:E0:1B:19	50:06:01:60:C7:E0:1B:19 50:06:01:6F:47:E0:1B:19

Figure 48 Unisphere Fibre Channel Ports page



SP A WWN	SP B WWN
50:06:01:60:C7:E0:1B:19 50:06:01:66:47:E0:1B:19	50:06:01:60:C7:E0:1B:19 50:06:01:6E:47:E0:1B:19
50:06:01:60:C7:E0:1B:19 50:06:01:67:47:E0:1B:19	50:06:01:60:C7:E0:1B:19 50:06:01:6F:47:E0:1B:19

Figure 49 Zoomed-in view of SP A and SP B WWNs

Two WWNs are listed for each port. The World Wide Node Name (WWNN), outlined in black, identifies this Unity storage array (the node). The WWPNs, outlined in blue, identify the individual ports. Record the WWPNs as shown in Table 8.

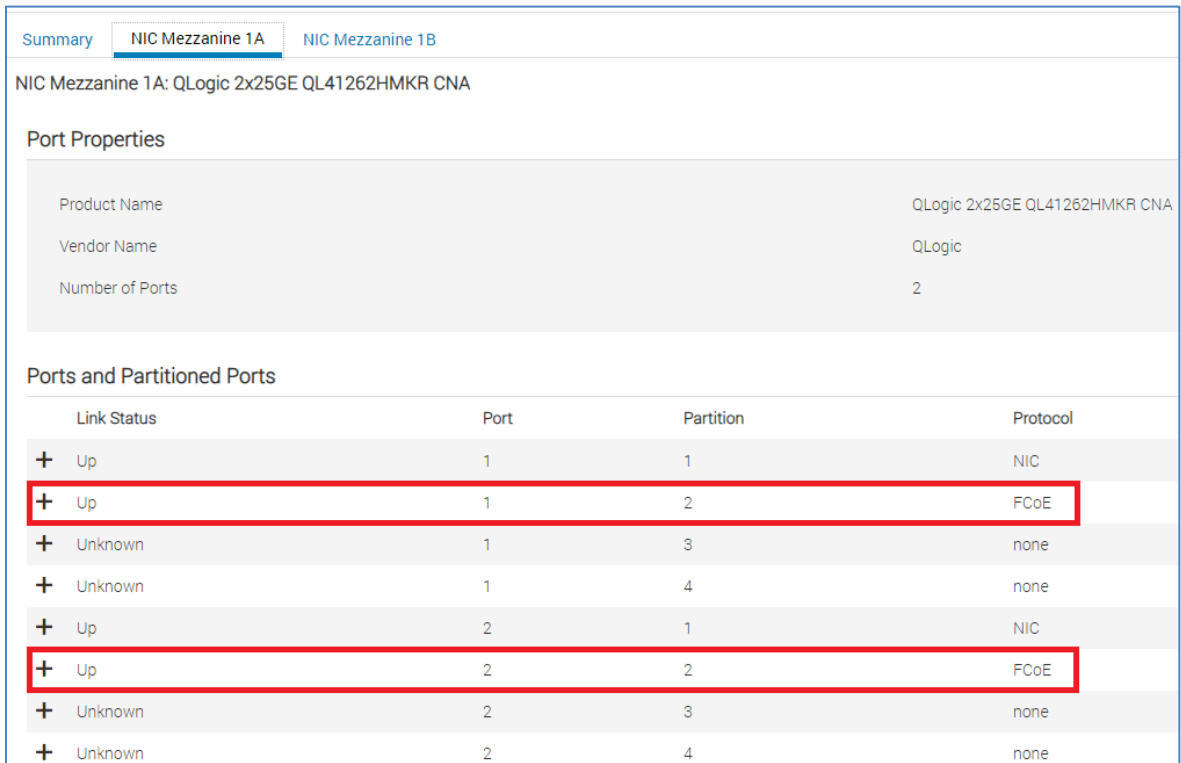
Table 9 Storage array FC adapter WWPNs

Service processor	Port	WWPN
SP A	0	50:06:01:66:47:E0:1B:19
SP A	1	50:06:01:67:47:E0:1B:19
SP B	0	50:06:01:6E:47:E0:1B:19
SP B	1	50:06:01:6F:47:E0:1B:19

## B.2 CNA FCoE port WWPNs

The PowerEdge server's FCoE adapter World Wide Port Names (WWPNs) are used for FC connection configuration. Adapter WWPNs are determined as follows:

1. Connect to the first server's iDRAC in a web browser and log in.
2. Select **System**, then click **Network Devices**.
3. Click the CNA. In this example, it is **NIC Mezzanine 1A**. Under **Ports and Partitioned Ports**, the FCoE partition for each port is displayed as shown in Figure 50:



Summary **NIC Mezzanine 1A** NIC Mezzanine 1B

NIC Mezzanine 1A: QLogic 2x25GE QL41262HMKR CNA

**Port Properties**

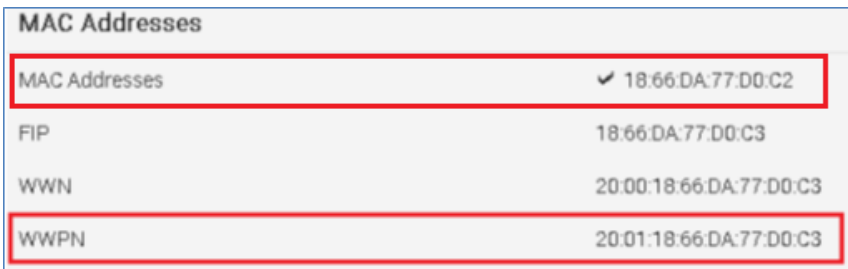
Product Name	QLogic 2x25GE QL41262HMKR CNA
Vendor Name	QLogic
Number of Ports	2

**Ports and Partitioned Ports**

Link Status	Port	Partition	Protocol
+ Up	1	1	NIC
+ Up	1	2	FCoE
+ Unknown	1	3	none
+ Unknown	1	4	none
+ Up	2	1	NIC
+ Up	2	2	FCoE
+ Unknown	2	3	none
+ Unknown	2	4	none

Figure 50 FCoE partitions in iDRAC

4. The first FCoE partition is **Port 1, Partition 2**. Click the **+** icon to view the MAC Addresses as shown in Figure 51:



**MAC Addresses**

MAC Addresses	✓ 18:66:DA:77:D0:C2
FIP	18:66:DA:77:D0:C3
WWN	20:00:18:66:DA:77:D0:C3
WWPN	20:01:18:66:DA:77:D0:C3

Figure 51 MAC address and FCoE WWPN for CNA port 1

5. Record the **MAC Address** and **WWPN**, outlined in red above.

---

**Note:** A convenient method is to copy and paste these values into a text file.

---

6. Repeat steps 4 and 5 for the FCoE partition on port 2.
7. Repeat in this section for the remaining servers.

The FCoE WWPNs and MAC addresses used in this deployment example are shown in Table 9.

Table 10 Server CNA FCoE port WWPNs and MACs

Server	Port	WWPN	MAC
MX740c-1	1	20:01:18:66:DA:71:50:AD	18:66:DA:71:50:AC
MX740c-1	2	20:01:18:66:DA:71:50:AF	18:66:DA:71:50:AE
MX740c-2	1	20:01:18:66:DA:77:D0:C3	18:66:DA:77:D0:C2
MX740c-2	2	20:01:18:66:DA:77:D0:C5	18:66:DA:77:D0:C4

## B.3 Configure Unity FC storage

This section covers configuration of a Dell EMC Unity 500F storage array. Refer to the storage system documentation for other FC storage devices.

### B.3.1 Create a storage pool

1. Connect to the Unisphere GUI in a web browser and log in using the required credentials.
2. In the left pane under **STORAGE**, select **Pools**.
3. Click the (+) icon.
4. In the **Create Pool** dialog box, provide a name in the required field, then click **Next**.
5. Select appropriate storage tiers and RAID configuration for the pool, then click **Next**.
6. Under **Select Amount of Storage**, select the desired number of drives. The total number of drives and the total capacity will be displayed next to **Totals**. Click **Next**.
7. The **Capability Profile Name** section is optional. Click **Next**.
8. Review selections on the **Summary** page and click **Finish** to create the pool. Once the Overall status shows 100%, click **Close**.
9. The pool is displayed on the **STORAGE > Pools** page as shown in Figure 52.




Figure 52 Storage pool created

### B.3.2 Add hosts

1. In the **Unisphere** left pane under **ACCESS**, select **Hosts**.
2. On the **Hosts** tab, click the (+) icon, then click **Host**.
3. Enter the **Name** of the server in the field provided, then click **Next**.
4. Select the Initiator IQN/WWN checkbox in the Discovered Initiators panel.
5. Click **Next**.
6. Click **Finish** in the Review the host configuration page.

---

**Note:** Additional hosts may be added as needed by clicking the + icon from the Hosts tab.

---

### B.3.3 Create LUNs and configure host access

1. In the **Unisphere** left pane under **STORAGE**, select **Block**
2. On the LUNs tab, click the (+) icon to open the **Create LUNs** dialog box.
3. On the Configure LUN(s) page, select the Number of LUNs.
4. Provide a name in the field provided, then select the **Storage Pool**.
5. From the Size section, make the required modifications, then click **Next**.
6. On the **Access** page, click the (+) icon and the select host(s) to be granted access to the LUN
7. Click **OK**, then click **Next**.
8. On the **Snapshot** page, leave settings at their defaults and click **Next**.

9. On the **Replication** page, leave settings at their defaults and click **Next**.
10. On the **Summary** page, review the details and click **Finish** to create the LUN.
11. On the Results page, click Close when Overall status shows 100% Completed.


The newly created LUN is now visible on the **LUNs** tab as shown in Figure 53. In this example, a LUN named FC-80GB that is 80GB in size has been created.

LUNs						
Consistency Groups						
iSCSI Interfaces						
+ [trash] [refresh] [edit] More Actions ▾ 1 item [signal] [gear] [download]						
<input type="checkbox"/>	!	Name ↓	Size (GB)	Allocated (%)	Pool	Thin Clone Base
<input type="checkbox"/>	✓	FC-80GB	80.0	<input type="text"/>	Pool_1	--

Figure 53 LUN Created

Create additional LUNs and grant access (map) to hosts as needed.

---

**Note:** To modify host access at any time, check the box next to the LUN to select it. Click the pencil  icon, and select the Host Access tab.

---

## C Validated components

The tables in this section include the hardware, software, and firmware used to configure and validate the examples in this document.

Table 11 Switches and OS versions

Qty	Item	Software Version
1	Dell EMC Networking S3048-ON management switch	10.4.0E(R3P2)
2	Dell EMC Networking MX9116n FSE	10.4.0E(R3S)
2	Dell EMC Networking MX5108	10.4.0E(R3S)
2	Dell EMC Networking S4148U-ON	10.4.0E(R3)
2	Dell EMC Networking MX7116n FEM	--

Table 12 MX-series components

Qty	Item	Version
4	Dell EMC PowerEdge M9002m modules	1.00.01
4	Dell EMC PowerEdge MX740c compute sleds	A01

Table 13 MX740c compute sled details

Qty per sled	Item	Firmware Version
1	QLogic QL41262HMKR (25G) mezzanine CNA	14.07.07
2	Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz	--
12	16GB DDR4 DIMMs (192GB total)	--
1	Boot Optimized Storage Solution (BOSS) Controller w/ 2x240GB SATA SSDs	2.6.13.3011
1	PERC H730P MX	25.5.5.0005
3	600GB SAS HDD	--
-	Dell OS Driver Pack	18.07.13
-	BIOS	1.0.2
-	iDRAC with Lifecycle Controller	3.20.20.20



## D FC attachment methods defined

The purpose of this document is to instruct the reader on how to connect MX server sleds to FC storage in three different ways: F\_port, NPG, and FSB. This appendix briefly discusses why you should use one over the other along with a basic explanation of how each method works. Which method to implement depends greatly on the existing infrastructure, as well as present and future scaling requirements that will fit your needs. Consult your Dell EMC representative for more information.

### D.1 F\_port

The term “F\_port” describes a port on the switch that connects to a node point-to-point (i.e. connects to an N\_port). It is also known as fabric port. A switch supporting F\_Port functionality allows FC storage to directly connect to the switch. The switch supports some of the services such as name server and zoning that are typical of fabric switches. The switch does not support creation of a multi-switch fabric because of lack of support for features such as FSPF. Two examples are provided in this guide that discuss F\_port mode. The topologies are shown in Figure 16 and Figure 23.

---

**Note:** N\_port is a port on the node (e.g. host or storage device). It is also known as the node port.

---

### D.2 NPIV Proxy Gateway (NPG)

NPIV Proxy Gateway (NPG) is a type of switch that significantly simplifies the Fibre Channel (FC) storage area network by reducing the number of domain IDs in the network as well as reducing the number of switches that need to be managed in the network. Dell NPG switches use VF\_Ports and N\_Ports to connect devices to the FC switch and/or FC storage. The connected FC switch connects to the NPG switch N\_Ports through F\_Ports and presents a variety of ports for connection to FC fabric devices.

NPG switches manage the following items:

- DCB (PFC, ETS and DCBx)
- FIP discovery and initialization
- FLOGI and FDISC conversion process
- FIP keep alives (FKA)

### D.3 FCoE (FSB)

FSB is a type of Ethernet switch that can exist between an FCoE server node and an FCF / NPG switch. FSB switches snoop FIP (FCoE Initialization Protocol) packets during the discovery and login phases. The switch then implements dynamic data integrity mechanisms using ACLs that permit valid FCoE traffic between the FCoE server nodes and FCF / NPG switch. Implementing such security mechanisms ensures that only valid FCoE traffic is allowed

## E Configure SRI-OV on blade servers

As data centers move into the realm of the Software Defined Data Center (SDDC), enhancements are needed to bring virtualized environment to its maximum potential. The significant efficiency that the hardware virtualizations achieved was not realized in the I/O devices. To support the virtualized environment with efficient I/O utilization, Single Root I/O Virtualization (SR-IOV) was introduced.

Sharing of a single PCIe device between various guest operating systems is available using SR-IOV. SR-IOV comprises of two main functions, namely, Physical Functions (PF) and Virtual Functions (VF). Each of the virtual function is associated with the physical functions of a device. A VF shares resources such as memory of the physical function. The VF can be associated directly associated with the guest operating system.

Follow the below steps to enable SR-IOV on the compute sleds with Intel NIC:

1. Using a web browser, connect to the iDRAC and launch the **Virtual Console**.
2. From the **Virtual Console**, click **Next Boot** menu then select **BIOS Setup**.
3. Select the option to reboot the server.
4. On the System Setup Main Menu, select **System Bios**.
5. Select **Processor Settings** on the System Bios Settings menu.
6. Verify if **Virtualization Technology** is set to **Enabled**. If not enabled, change the **Virtualization Technology** option to **Enabled**.

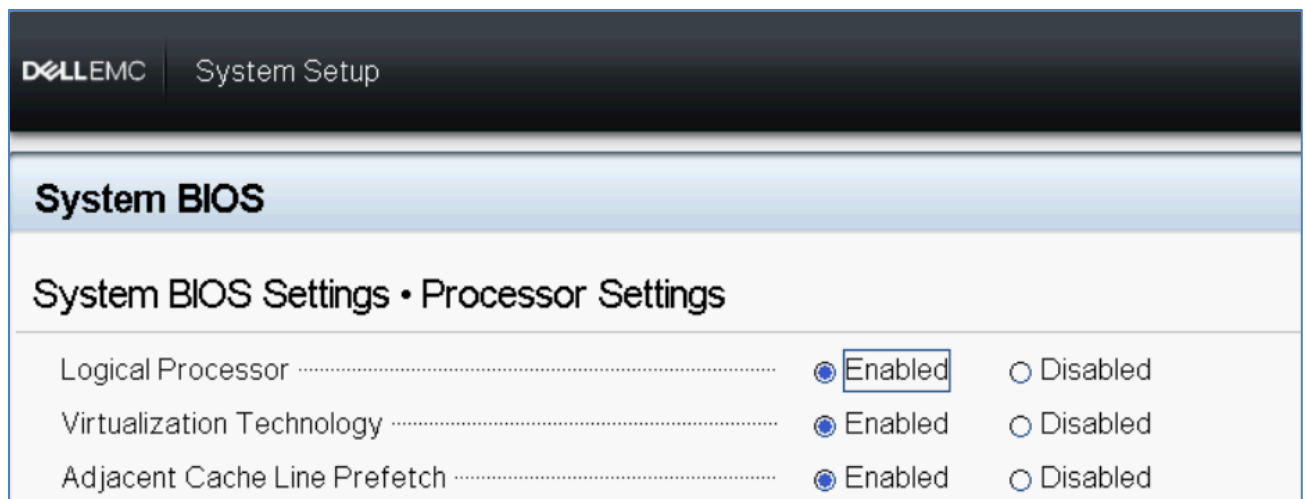


Figure 54 CNA Processor Settings configuration

7. Click **Back**.
8. In the **System Bios Settings** menu, choose **Integrated Devices**.
9. Select **SR-IOV Global Enable** option to enable.

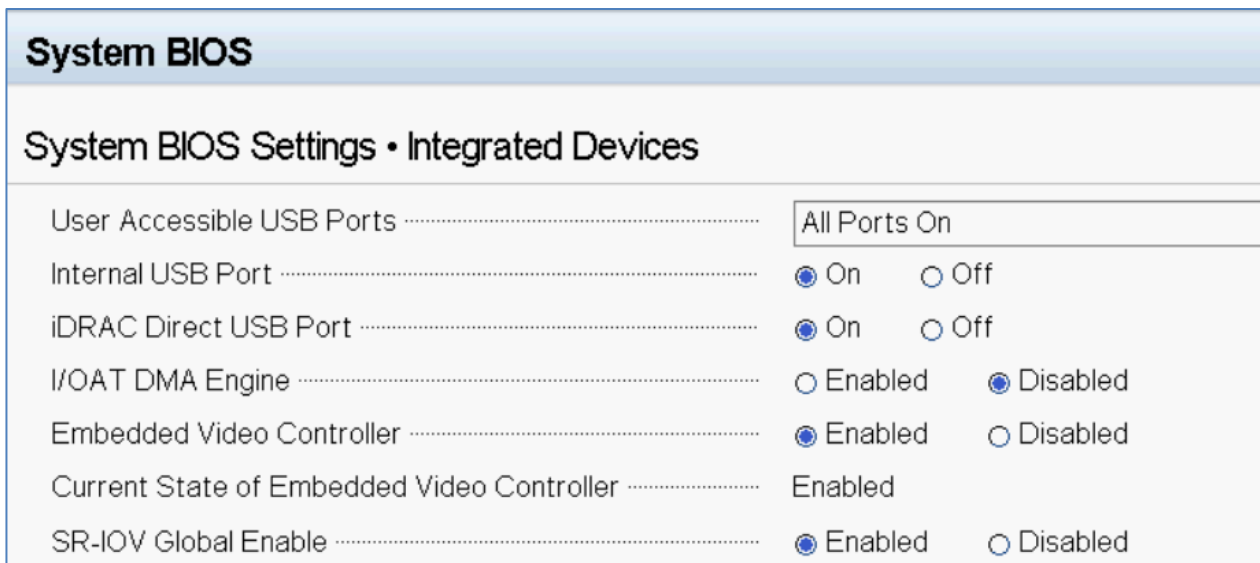


Figure 55 Integrated Devices Settings configuration

10. Click **Back**.
11. Click **Finish**.
12. Choose **Yes** to save the changes.
13. Click **OK**.
14. In the **System Setup Main Menu**, choose **Device Settings**.
15. In the **Device Settings** menu, choose the desired NIC for SR-IOV.

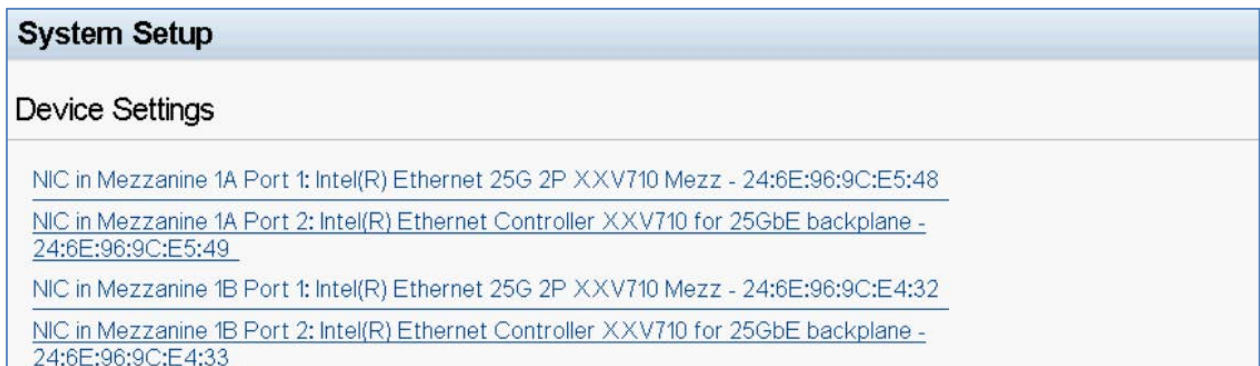


Figure 56 Device Settings configuration

**Note:** In this example, NIC in Mezzanine 1A Port 1: Intel(R) Ethernet 25G 2P XXV710 Mezz – 24:6E:96:9C:E5:48 is selected.

16. In the **Main Configuration** page that displays, select **Device Level Configuration**.
17. In the **Device Level Configuration** page, select the **Virtualization mode** to SR-IOV.

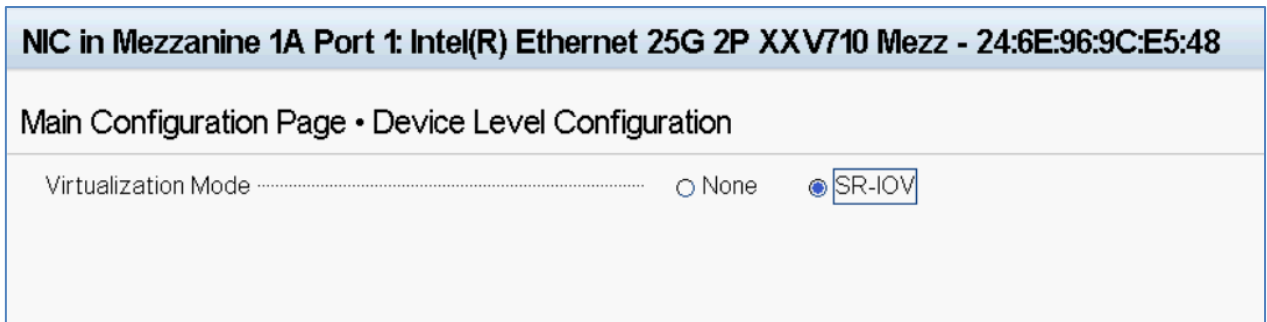


Figure 57 Device Level configuration

18. Click **Back**.
19. Click **Finish**. Choose **Yes** to save the changes. Click **OK**.
20. Click **Finish** to go to the **System Setup Main Menu**.
21. Click **Finish**. Click **Yes** to confirm the exit and reboot. The server will now reboot.

---

**Note:** Further steps needed to complete enabling SR-IOV are done using the operating system installed on the blade server and steps vary based on the OS. See your OS guide for further details on enabling SR-IOV on the OS.

---

## F Technical resources

[Dell Community](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware, and services. Other helpful links related to the concepts in this guide are listed below.

[Dell EMC Networking Guides](#)

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

[Dell EMC PowerEdge MX SmartFabric Network Deployment Guide](#)

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

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

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

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

[FCoE-to-Fibre Channel Deployment with S4148U-ON in F\\_port Mode](#)

[Fibre Channel Deployment with S4148U-ON in F\\_port Mode](#)

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

[Manuals and documentation for Dell EMC Networking S4148U-ON](#)

[Manuals and documentation for Dell EMC Networking S3048-ON](#)

[Manuals and documentation for Dell EMC Networking Z9264F-ON](#)

[Manuals and documentation for Dell EMC Networking Z9100-ON](#)

## G Support and feedback

### Contacting Technical Support

Support Contact Information

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

Telephone: USA: 1-800-945-3355

### Feedback for this document

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