FX2 Storage Networking with FC (NPG)

Connecting PowerEdge FX2 blade servers to Fibre Channel storage using FN2210S I/O Modules and FC switches

Dell EMC Networking Solutions Engineering
December 2017

A Dell EMC Deployment and Configuration Guide

Revisions

Date	Rev.	Description	Authors
December 2017	3.1	Revised SAN topology diagrams. Revised 6510 zoning and validation. Updated VMware software versions and VMware documentation links. Updated Leaf-Spine Deployment Guide references, updated to DNOS 9.11.2.6.	Jim Slaughter Shree Rathinasamy
May 2017	3.0	Updated to include DNOS 9.11, Brocade 6510 and Unity FC storage. Added sections including configuration of S4048-ON, 6510, virtual switches and FC Storage. Revised existing sections.	Jim Slaughter, Jordan Wilson
July 2016	2.1	EMC Integration changes	Ed Blazek
December 2015	2.0	Added DNOS 9.9 configuration information	Curtis Bunch, Jim Slaughter, Mike Matthews
July 2015	1.1	Added additional SMUX configuration and quad-port configuration examples. Fixed formatting issues.	Curtis Bunch, Kevin Locklear, Mike Matthews
December 2014	1.0	Initial release	Curtis Bunch, Kevin Locklear, Mike Matthews

Table of contents

1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company in the data center - for not only today's applications, but for the cloud-native world we are entering. To attain that vision, the Dell EMC portfolio focuses not only on making every component of data center infrastructure (servers, storage, and networking) compelling to our customers, but to also make the value in the integration of those components greater than the sum of the parts.

This document focuses on three specific elements of the Dell EMC portfolio:

Dell EMC PowerEdge FX2. FX2 is an innovative design with modular IT building blocks to precisely address evolving workloads. The flexibility of the FX2 chassis and its variety of components enables a wide range of target customers; data centers building private clouds, web service providers and dedicated hosting organizations and enterprises that want easily scalable platforms to address growth. Specific targets are configuration dependent, so customers can tailor infrastructure precisely, with the right power, storage and connectivity to meet specific workload needs.

Dell EMC Networking. Running internal to the FX2 chassis, Dell EMC Networking developed the FN-series switches. An integrated networking solution, the FN-series provides Ethernet as well as LAN/SAN convergence with iSCSI and FCoE support. Scaling beyond multiple chassis, S-series switches are multi-layer Ethernet switches that run on Dell EMC's own operating system, or the operating system from one of our Open Networking ecosystem partners.

Dell EMC Unity. The Unity family delivers high-performance hybrid or All-Flash midrange unified storage, with NAS and SAN connectivity. Unity simplifies and modernizes today's data center with a powerful combination of enterprise capabilities and cloud-like simplicity.

This guide provides assistance for a step-by-step deployment of Fibre Channel using Dell PowerEdge FX2s and Dell EMC Unity 500F. It includes configuration of physical switches, ESXi hosts, a Virtual Distributed Switch and Unity 500F storage. The goal of this guide is to enable a network administrator or engineer with traditional networking and VMware ESXi experience to deploy Fibre Channel using the Dell EMC hardware and software outlined in this guide.

1.1 Typographical Conventions

The command line examples in this document use the following conventions:

`Monospace Text`	CLI examples
`Italic Monospace Text`	Variables in CLI examples
`Bold Monospace Text`	Commands entered at the CLI prompt

2 Hardware overview

This section briefly describes the primary hardware used to validate this deployment. A complete listing of hardware validated for this guide is provided in Appendix C.

2.1 Dell EMC PowerEdge FX2s enclosure and supported modules

The PowerEdge FX2s enclosure is a 2-rack unit (RU) computing platform. It has capacity for two FC830 full-width servers, four FC630 half-width servers or eight FC430 quarter-width servers. The enclosure is also available with a combination of servers and FD332 storage sleds. The FX2s enclosure used in this guide contains four FC630 servers as shown in Figure 1.

Figure 1 Dell EMC PowerEdge FX2s (front) with four PowerEdge FC630 servers

The back of the FX2s enclosure includes a Chassis Management Controller (CMC), two FN IO Modules (FN IOMs), eight Peripheral Component Interconnect Express (PCIe) expansion slots and redundant power supplies.

Figure 2 Dell EMC PowerEdge FX2s (back)

Note: The Dell EMC PowerEdge FX enclosure is currently available in two models: FX2 and FX2s. The FX2s adds eight PCIe slots, shown in Figure 2, and provides the option to add PowerEdge FD332 storage sleds in the front slots. In this deployment, four of the PCIe slots are populated with 1GbE Intel I350-T network adapters (for ESXi host management) and four PCIe slots are unused.

2.1.1 PowerEdge FC630 server

The PowerEdge FC630 server is a half-width, 2-socket server. Four FC630 servers in the FX2s enclosure are used in the deployment example in this guide.

Figure 3 PowerEdge FC630

2.1.2 PowerEdge FN I/O Modules

PowerEdge FN I/O Modules (IOMs) are network switches housed in the back of the FX2s enclosure. Dell EMC offers three FN IOM options providing Ethernet as well as LAN/SAN convergence with Internet Small Computer System Interface (iSCSI) and Fibre Channel over Ethernet (FCoE) support. In addition to these features, the FN2210S supports native FC traffic using N_Port ID Virtualization (NPIV) Proxy Gateway (NPG) mode for connections to an intermediate FC switch or F_Port mode for direct connections to FC storage arrays.

All three FN IOM options, FN410S, FN410T and FN2210S, provide eight 10GbE internal, server-facing ports and four external ports.

FN410S
SFP+ I/O Module
Provides four external SFP+ 10GbE ports. Supports optical and direct attach copper (DAC) cables.

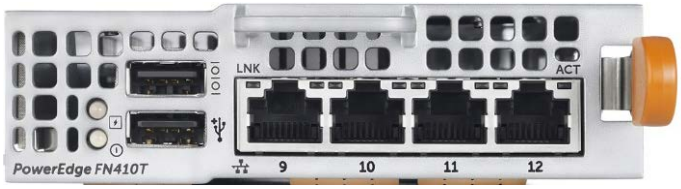

FN410T
10GBASE-T I/O Module
Provides four external 10GbE Base-T ports. Supports cost-effective twisted-pair copper cables up to 100 meters.

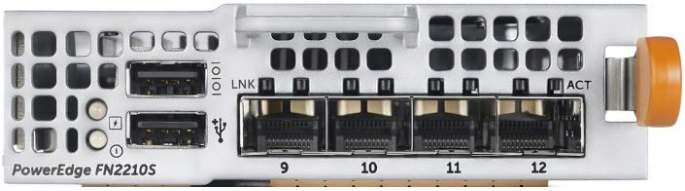

FN2210S
Combo FC/Ethernet I/O Module
Provides four external SFP+ ports configurable in any combination of 10GbE or 2/4/8 Gbps FC. Supports optical and DAC cables.

Figure 4 FN IOMs

Two PowerEdge FN2210S IOMs in NPG mode are used in the FX2s enclosure for the deployment example in this guide. Two external ports are configured for Ethernet traffic and two are configured for FC traffic.

2.2 Dell EMC Networking S3048-ON

The S3048-ON is a 1-RU switch with forty-eight 1GbE Base-T ports and four 10GbE SFP+ ports. In this guide, one S3048-ON switch supports management traffic in each rack.

Figure 5 Dell EMC Networking S3048-ON

2.3 Dell EMC Networking S4048-ON

The S4048-ON is a 1-RU, multilayer switch with forty-eight 10GbE SFP+ ports and six 40GbE QSFP+ ports. Two S4048-ONs are used as leaf switches in each rack as part of the data center's leaf-spine network.

Figure 6 Dell EMC Networking S4048-ON

2.4 Dell EMC Networking S6010-ON

The S6010-ON is a 1-RU, multilayer switch with thirty two 40GbE QSFP+ ports. In this guide, one S6010-ON is used as a spine switch.

Figure 7 Dell EMC Networking S6010-ON

2.5 Brocade 6510 / Connectrix DS-6510B

Brocade 6510 and Connectrix DS-6510B switches are 1-RU, 2/4/8/16 Gbps FC switches with up to 48 ports. In this deployment guide, the term "6510" is used to refer to either model. Two 6510 switches provide connections from FN2210S IOMs to FC storage.

Figure 8 6510 Fibre Channel switch

3 NPG and the example environment

Two FN2210S IOMs are installed in the FX2s chassis for redundancy and performance. Each FN2210S provides four external ports that may be configured as FC or Ethernet ports. The mix of FC and Ethernet ports allow it to simultaneously connect to an FC SAN for storage traffic and a LAN for TCP/IP traffic on the production network.

In this deployment example, two ports on each FN2210S are configured for FC and are connected to an FC storage area network (SAN).

FC SAN traffic is sent between an FC630 server's Converged Network Adapter (CNA) and the FN2210S using FCoE. Native FC is used between the FN2210S and FC switches. In this configuration, the FN2210S acts as an NPG.

LAN traffic sent from the CNA is forwarded to a pair of leaf switches in each rack that are part of the data center's leaf-spine network.

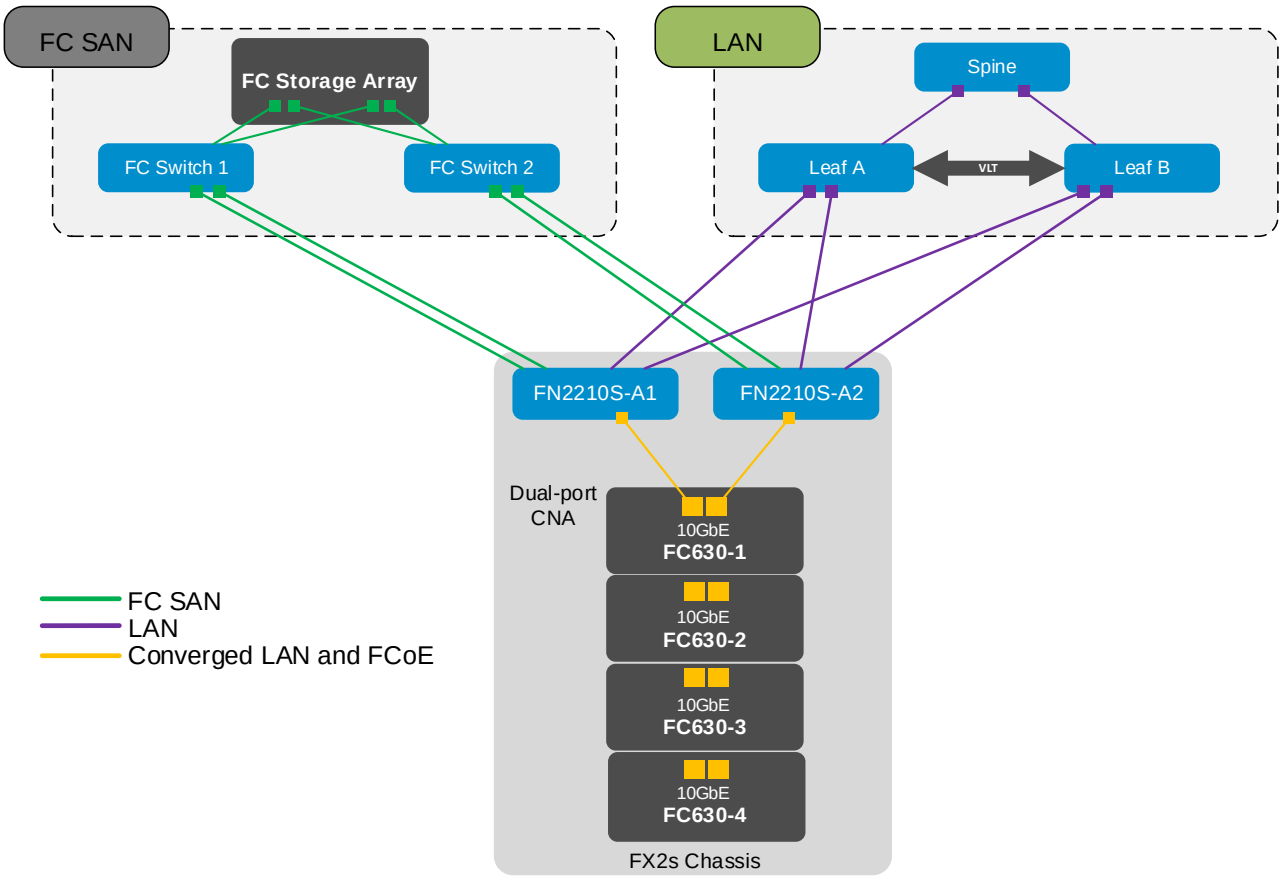

Figure 9 Combined LAN and SAN topologies

Note: In this deployment guide, "LAN" is used to refer to the data center's TCP/IP network and "SAN" refers to the FCoE and FC network.

3.1 SAN topology

The storage network used throughout this guide is shown in Figure 10. It provides connections from the FC630 servers in the FX2s chassis to the FC storage array.

Note: The FC storage array used in all examples in this guide is a Dell EMC Unity 500F with redundant storage processors, SP A and SP B. Other validated FC storage arrays are listed in Appendix C.3.

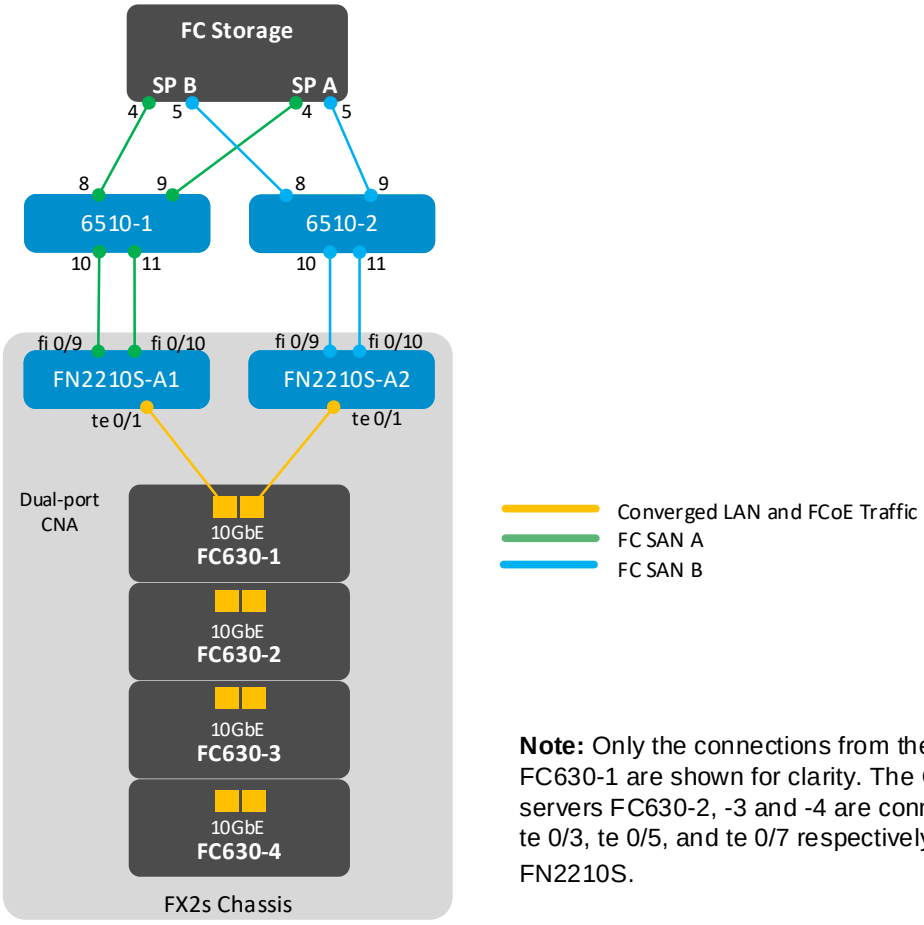

Note: Only the connections from the CNA in FC630-1 are shown for clarity. The CNAs in servers FC630-2, -3 and -4 are connected to ports te 0/3, te 0/5, and te 0/7 respectively on each FN2210S.

Figure 10 SAN topology

Inside the FX2s chassis, each FC630 server has one dual-port CNA. One port is connected to FN2210S-A1 and the other connected to FN2210S-A2. These connections carry converged LAN and FCoE SAN traffic.

Each FN2210S is connected to a 6510 switch with 8 Gbps FC connections using external Fibre Channel ports fi 0/9 and fi 0/10. Each 6510 switch has 16Gbps connections to FC storage.

Storage traffic is sent from the server CNAs using FCoE. The FN2210S decapsulates the FC payload from the FCoE frames and forwards it through the 6510 switch to the FC storage array. Incoming FC frames received by the FN2210S are encapsulated using FCoE and forwarded to the appropriate CNA.

3.2 LAN topology - production network

The LAN topology for the production network includes the converged connections and leaf-spine network connections for TCP/IP traffic as shown in Figure 11.

S4048-ON leaf switches are configured as a Virtual Link Trunking (VLT) pair, with one pair at the top of each rack. The two external Ethernet ports on each FN2210S, te 0/11 and te 0/12, are configured in LACP port channel 128, with one connection going to each leaf switch. The corresponding leaf switch port channel numbers used in this example are 101 (to FN2210S-A1) and 102 (to FN2210S-A2).

Note: Using a leaf-spine network in the datacenter as shown in Figure 11 is considered a best practice. Configuration of the S6010-ON spine switch is outside the scope of this document. See the Dell EMC Leaf-Spine Deployment and Best Practices Guide for a detailed description and configuration instructions for a leaf-spine network.

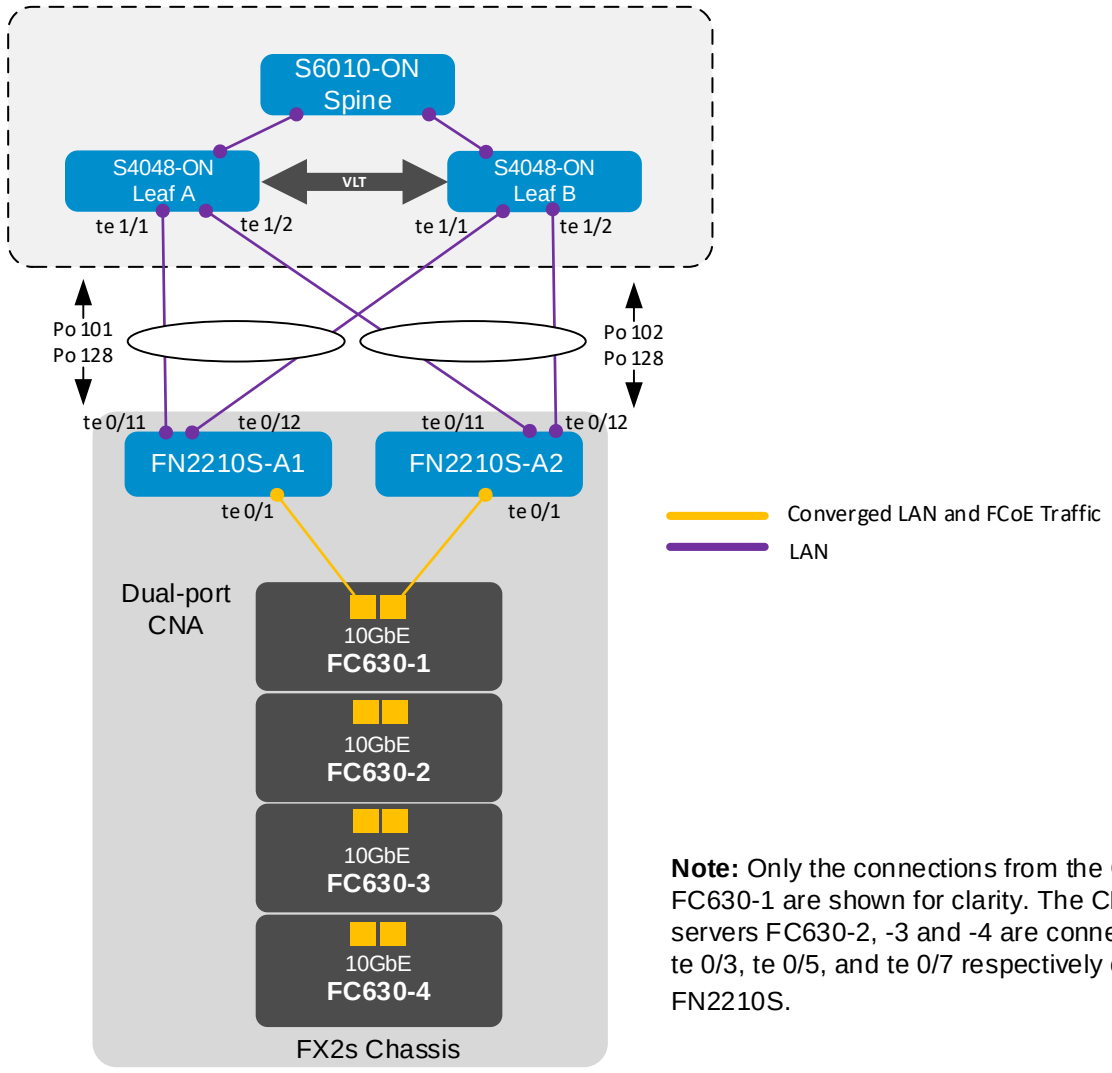

Note: Only the connections from the CNA in FC630-1 are shown for clarity. The CNAs in servers FC630-2, -3 and -4 are connected to ports te 0/3, te 0/5, and te 0/7 respectively on each FN2210S.

Figure 11 LAN topology – production network

3.3 LAN topology - management network

This guide uses an isolated network for management traffic. One S3048-ON switch installed in each rack provides connectivity to the management network.

The PowerEdge FX2s has four 1GbE Intel I350-T network adapters installed in the PCIe slots in the back of the chassis. Each is connected internally to an FC630 server and used for ESXi host management. The FX2s Chassis Management Controller (CMC) is also connected to the S3048-ON for out-of-band (OOB) management. The CMC provides management access to the FN IOMs and to each FC630 server's integrated Dell Remote Access Controller (iDRAC).

The management ports of all switches and storage devices in the rack are also connected to the S3048-ON as shown in Figure 12.

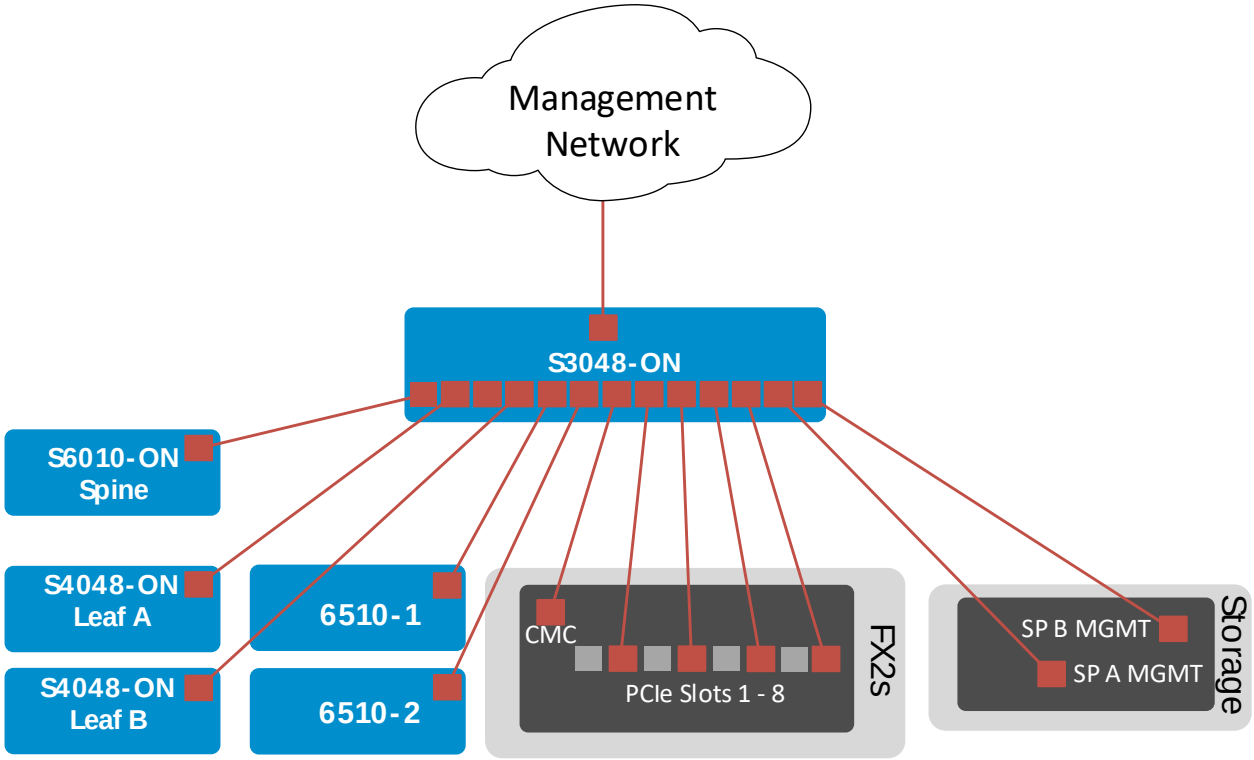

Figure 12 LAN topology - management network

3.4 ESXi Hosts

The four ESXi hosts used in this deployment example appear in screenshots by their management IP addresses. The host names and management IP addresses used are shown in Table 1.

Table 1 ESXi host addresses

Host name	Management IP address
FC630-1	100.67.169.51
FC630-2	100.67.169.52
FC630-3	100.67.169.53
FC630-4	100.67.169.54

3.5 VLANs

VLAN 1002 is the factory default FCoE VLAN on FN2210S IOMs. The FCoE VLAN ID may be changed in Full-switch mode but this guide uses 1002 in all examples for consistency. The FCoE VLAN must only carry FCoE traffic.

Configuration of VLANs for non-storage traffic is outside the scope of this document.

4 Configuration prerequisites

Note: Exact iDRAC and CMC steps in this section may vary slightly depending on hardware, software and browser versions used. See your PowerEdge server documentation for steps to connect to the iDRAC and CMC.

4.1 Reset server CNAs to factory Defaults

Note: Resetting to defaults is only necessary if installed CNAs have been modified from their factory default settings.

1. Connect to the server's iDRAC in a web browser and launch the virtual console.
2. In the virtual console, from the **Next Boot** menu, select **BIOS Setup**.
3. Reboot the server.
4. From the **System Setup Main Menu**, select **Device Settings**.
5. From the **Device Settings** page, select the first port of the first adapter in the list.
6. From the **Main Configuration Page**, click the **Default** button followed by **Yes** to load the default settings. Click **OK**.
7. To save the settings, click **Finish** then **Yes** to save changes. Click **OK**.
8. Repeat for each adapter and port listed on the **Device Settings** page.

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 the QLogic 57810 CNA Network Daughter Card (NDC) in each server. It is not done on the Intel I350T PCIe adapters used for ESXi host management.

If you are already in System Setup from the prior section, skip to step 4.

1. Connect to the server's iDRAC in a web browser and launch the virtual console.
2. In the virtual console, from the **Next Boot** menu, select **BIOS Setup**.
3. Reboot the server.
4. On the **System Setup Main Menu**, select **Device Settings**.
5. Select **Integrated NIC 1 Port 1**.
6. Select **Device Level Configuration**. Set **Virtualization Mode** to **NPar** and click **Back**.
7. Select **NIC Partitioning Configuration.**
8. Select **Partition 1 Configuration**.
   a. Set **NIC Mode** to **Enabled**.
   b. Set **iSCSI Offload Mode** to **Disabled**.
   c. Set **FCoE Mode** to **Disabled** and click **Back**.
9. Select **Partition 2 Configuration**
   a. Set **NIC Mode** to **Disabled**.
   b. Set i**SCSI Offload Mode** to **Disabled**.
   c. Set **FCoE Mode** to **Enabled** (screen shown in Figure 13) and click **Back**.

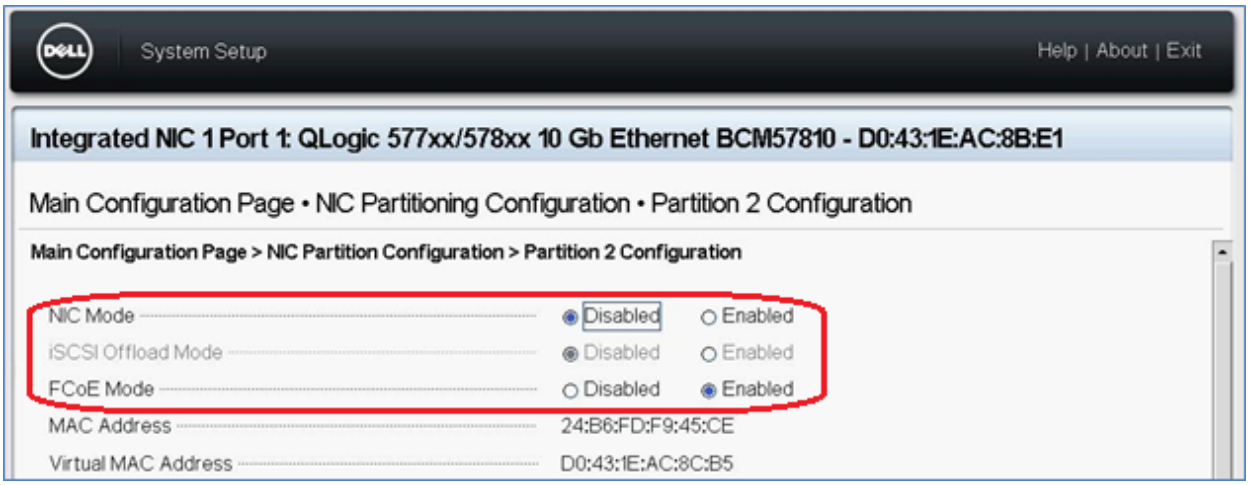

Figure 13 System Setup – Partition 2 configured for FCoE

10. Select **Partition 3 Configuration**. Set NIC, iSCSI and FCoE modes to **Disabled** and click **Back**.
11. Select **Partition 4 Configuration**. Set NIC, iSCSI and FCoE modes to **Disabled** and click **Back**.
12. Click **Back > Finish**.
13. When prompted, answer **Yes** to save changes and click **OK** in the **Success** window.
14. Select **Integrated NIC 1 Port 2** and repeat steps 6-13 above for port 2.
15. Click **Finish > Finish** and answer **Yes** to exit and reboot.

Repeat on remaining servers in the FX2s chassis.

4.3 Determine CNA FCoE port WWPNs

The FC630 server's FCoE adapter World Wide Port Names (WWPNs) are required for zone configuration on the FC switches. Server adapter WWPNs are determined as follows:

1. Connect to the FX2s chassis' CMC in a web browser and log in.
2. In the left pane of the GUI, select **Server Overview.** After the page loads, make sure the **Properties** tab is open and click **WWN/MAC**.
3. Under **WWN/MAC Addresses**, click **Expand/Collapse All** to expand. The screen will show the WWNs and MAC addresses for all servers in the chassis.
4. Filter the results as follows:
   a. Under **Protocol**, select **4xFC FlexIO Module** (or **Fibre Channel** if present).
   b. Under **Partition Status**, select **Enabled**.

The FCoE WWNs (also known as WWPNs) of all configured CNAs in the FX2s chassis are displayed as shown in Figure 14.

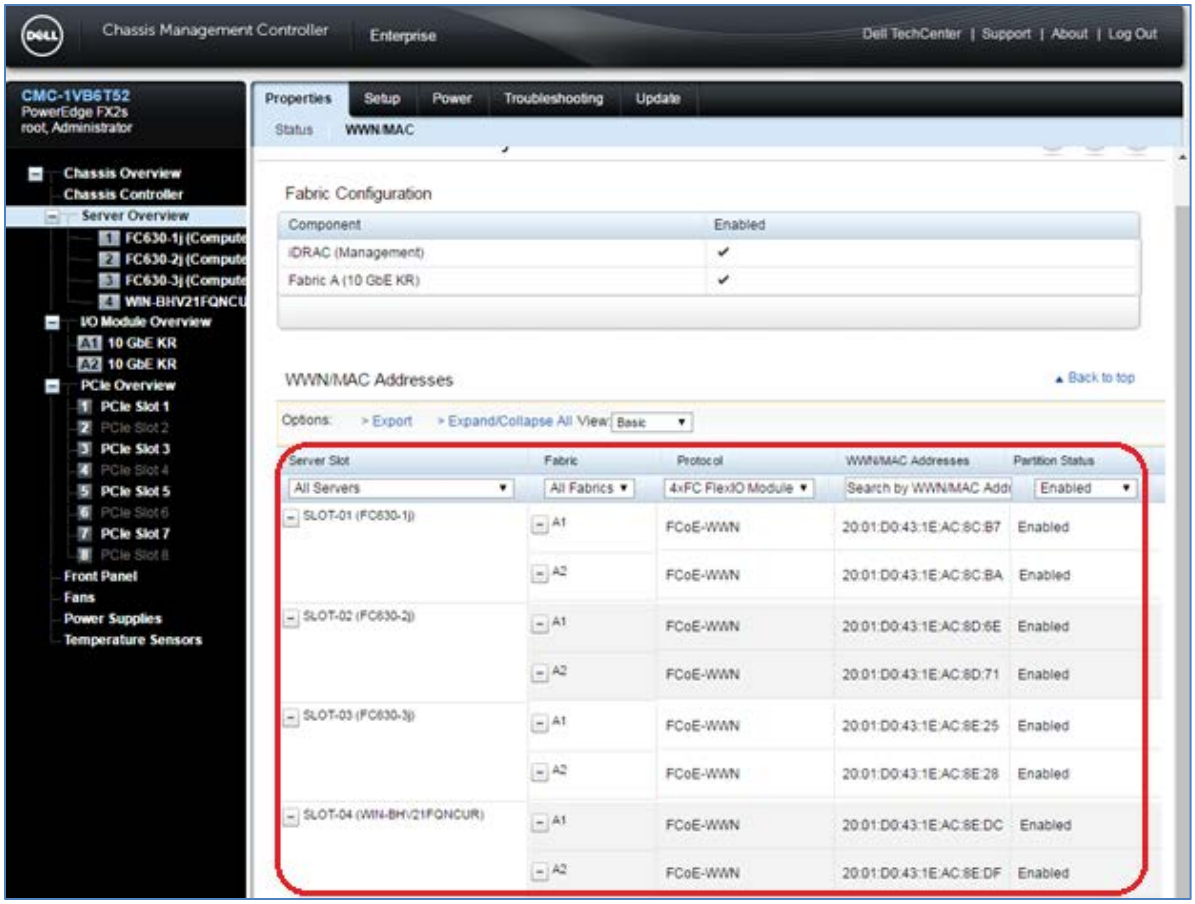

Figure 14 FCoE-WWNs of converged network adapters

Record the WWPNs as shown in Table 2.

Note: The **Export** feature on the WWN/MAC page may be used to export all addresses to a .csv file.

Table 2 Server CNA FCoE port WWPNs

Server	Fabric	WWPN
FC630-1	FN2210S-A1	20:01:D0:43:1E:AC:8C:B7
FC630-1	FN2210S-A2	20:01:D0:43:1E:AC:8C:BA
FC630-2	FN2210S-A1	20:01:D0:43:1E:AC:8D:6E
FC630-2	FN2210S-A2	20:01:D0:43:1E:AC:8D:71
FC630-3	FN2210S-A1	20:01:D0:43:1E:AC:8E:25
FC630-3	FN2210S-A2	20:01:D0:43:1E:AC:8E:28
FC630-4	FN2210S-A1	20:01:D0:43:1E:AC:8E:DC
FC630-4	FN2210S-A2	20:01:D0:43:1E:AC:8E:DF

4.4 Determine storage array FC WWPNs

Note: This section shows the procedure for a Dell EMC Unity 500F storage array. Refer to your storage system documentation to determine WWPNs on other FC storage devices.

The WWPNs of FC adapters in storage arrays are used for configuring zoning on the FC switches. WWPNs on Unity storage arrays are determined as follows:

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

The Fibre Channel ports page is displayed as shown in Figure 15. A zoomed-in view of the area inside the red box is shown in Figure 16.

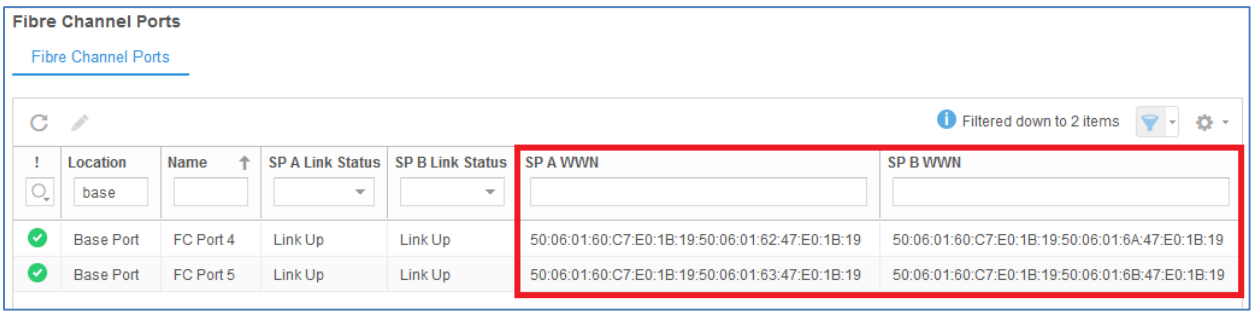

Figure 15 Fibre Channel Ports page

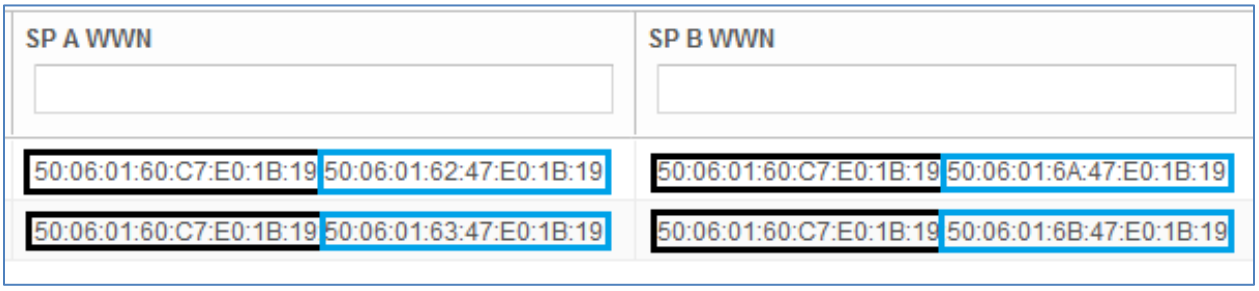

Figure 16 Zoomed-in view of SP A and SP B WWNs on Fibre Channel Ports page

3. Two WWNs are listed for each port. The World Wide Node Name (WWNN) is outlined in black and is not used for zoning. The WWPNs are outlined in blue. Record the four WWPNs as shown in Table 3.

Table 3 Storage array FC adapter WWPNs

Port	WWPN
SP A Port 4	50:06:01:62:47:E0:1B:19
SP B Port 4	50:06:01:6A:47:E0:1B:19
SP A Port 5	50:06:01:63:47:E0:1B:19
SP B Port 5	50:06:01:6B:47:E0:1B:19

4.5 VMware preparation

4.5.1 VMware ESXi download and installation

Install VMware ESXi 6.5 U1 on each FC630 server in the FX2s chassis.

Dell EMC recommends using the latest Dell EMC customized ESXi .iso image available on support.dell.com. The correct drivers for your PowerEdge hardware are built into this image. This image can be used to install ESXi via CD/DVD, a USB flash drive, or by mounting the .iso image through the PowerEdge server's iDRAC interface.

4.5.2 Install and configure VMware vCenter Server 6.5 U1

The installation and initial configuration of VMware vCenter Server is not covered in this guide. For information on the installation and configuration of vCenter Server, refer to the *vSphere Installation and Setup* guide.

4.5.3 Add ESXi hosts to vCenter

The vSphere Web Client is a service running on vCenter Server. In the vSphere Web Client, a data center object named **Datacenter** is created for this deployment and the ESXi hosts are added to it.

A cluster named **FX2-FC630-FC** is created and added to the datacenter object. The four FC630 hosts are added to the FX2-FC630-FC cluster.

When complete, the vSphere Web Client **Hosts and Clusters** tab in the **Navigator** pane appears as shown in Figure 17.

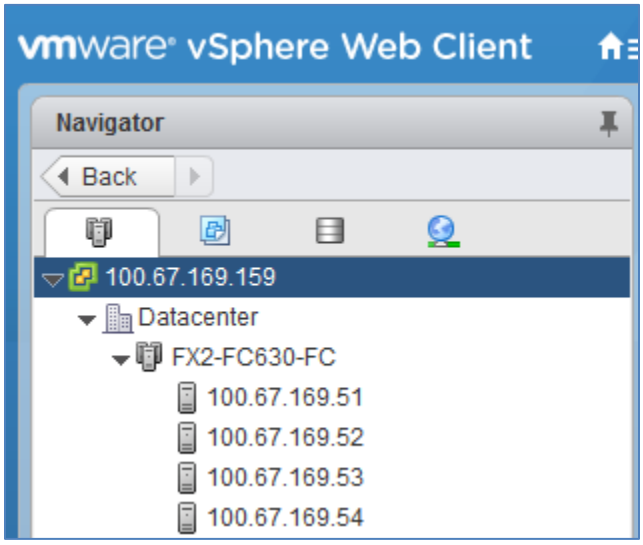

Figure 17 Datacenter and cluster created with four ESXi hosts

4.5.4 Add VMkernel adapters for FCoE

Note: Before starting this section, be sure you know the vmnic-to-physical adapter mapping for each host. This can be determined by going to **Home > Hosts and Clusters** and selecting the host in the **Navigator** pane. In the center pane go to **Configure > Networking > Physical adapters**. Vmnics and their MAC addresses are shown. In this example, the two FCoE adapters on each host are vmnic2 and vmnic3. Your vmnic numbering may vary.

To add VMkernel adapters for FCoE traffic:

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. In the left pane, select the first ESXi host. In the center pane, select **Configure > Networking > Virtual switches**.
3. For the first adapter, click the  icon and select **VMkernel Network Adapter**. Click **Next**.
4. Select **New standard switch** and click **Next.**
5. Under **Assigned adapters**, click the  icon. Select the first vmnic enabled for FCoE, e.g. vmnic2, and click **OK**. Make sure it is listed under **Active Adapters** and click **Next.**
6. Next to **Network label**, provide a name, e.g. FCoE1. Set the **VLAN ID** to 1002. (VLAN 1002 is used by default for FCoE traffic on FN2210S switches). Leave the remaining settings at their defaults and click **Next.**
7. On the **IPv4 settings** page, keep the default selection and click **Next > Finish.**
8. For the second adapter, click the  icon and select **VMkernel Network Adapter**. Click **Next**.
9. Select **New standard switch** and click **Next.**
10. Under **Assigned adapters**, click the  icon. Select the second vmnic enabled for FCoE, e.g. vmnic3. Click **OK > Next.**
11. Specify a **Network label**, e.g. FCoE2. Set the **VLAN ID** to 1002. Leave the remaining settings at their defaults and click **Next.**
12. On the **IPv4 settings** page, keep the default selection and click **Next > Finish.**

Repeat steps 2-12 for remaining ESXi hosts.

When complete, the **Configure > Networking > VMkernel adapters** page for each host appear should appear similar to Figure 18. (The IP addresses shown next to the two FCoE adapters are automatically assigned private addresses and are not used for FCoE).

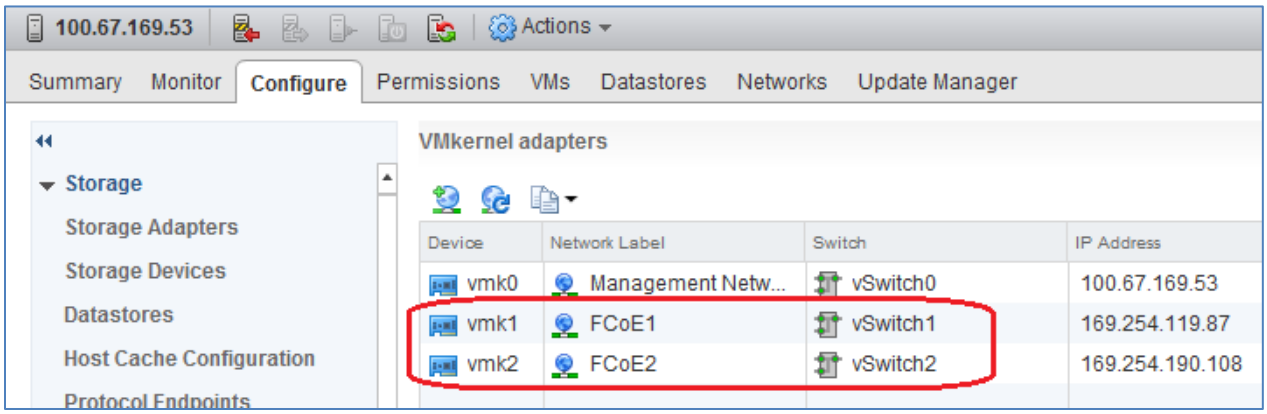

Figure 18 Host VMkernel adapters configured

4.5.5 Increase MTU size for FCoE

FCoE frames may be up to 2180 bytes in size. By default, VMware vSwitches and VMkernel adapters have the Maximum Transmission Unit (MTU) size set to 1500 bytes.

The following steps increase the MTU size to 2500 bytes on the vSwitches created for FCoE:

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. In the left pane, select the first ESXi host. In the center pane, select **Configure > Virtual switches**.
3. Click on the first FCoE vSwitch, e.g. **vSwitch1**, and click the upper icon as shown in Figure 19.

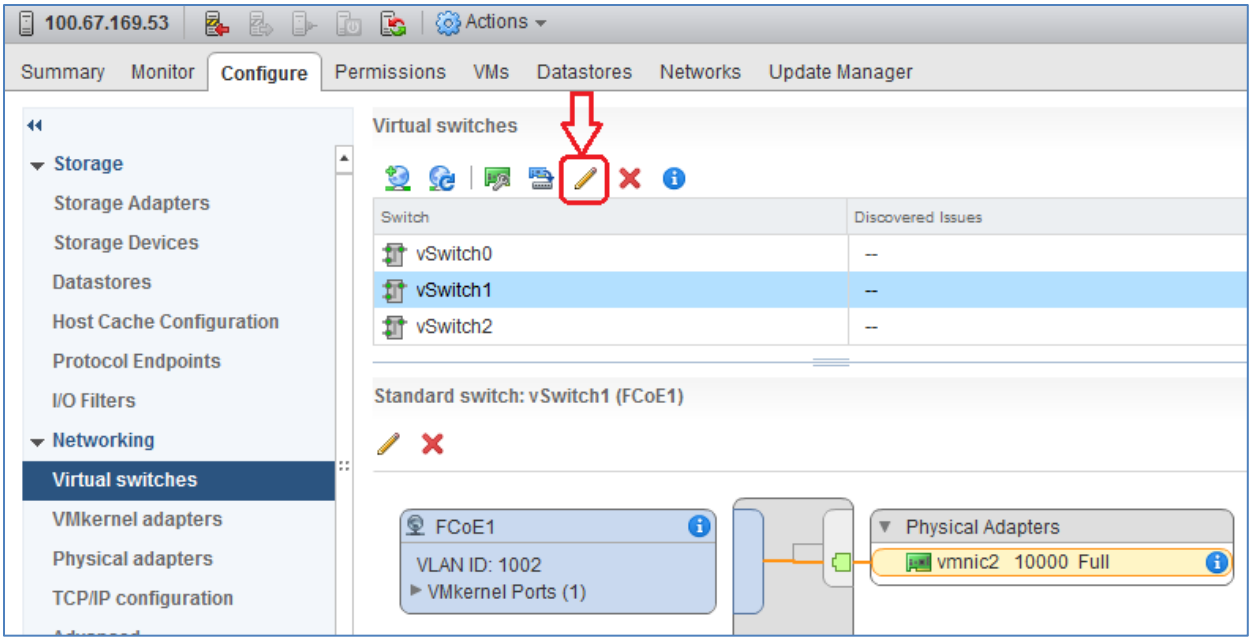

Figure 19 Selecting the correct vSwitch edit icon

4. In the **Edit Settings** dialog box, change the value of **MTU (Bytes)** to 2500 as shown in Figure 20.

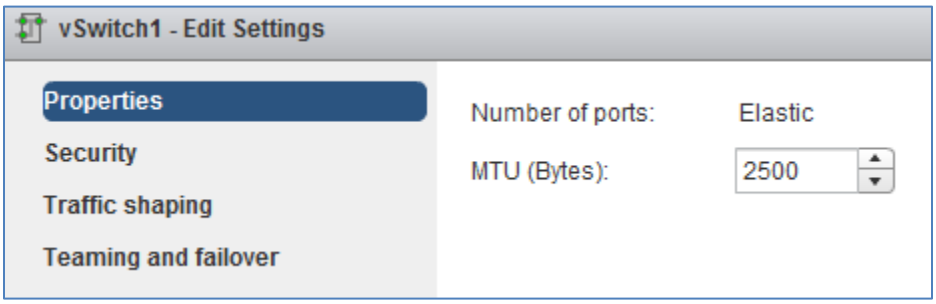

Figure 20 vSwitch MTU set to 2500 bytes

5. Click **OK** to apply the setting and close the box.
6. Repeat for the host's second FCoE vSwitch, e.g. vSwitch2.
7. Repeat the steps above for the remaining ESXi hosts.

The following steps increase the MTU size to 2500 bytes on the VMkernel adapters created for FCoE:

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. In the left pane, select the first ESXi host. In the center pane, select **Configure > VMkernel adapters**.
3. Click on the first FCoE VMkernel adapter, e.g. FCoE1.
4. Click the icon to open the **Edit Settings** dialog box and select **NIC settings**.
5. Change the **MTU** value to **2500** bytes.
6. Click **OK** to apply the setting and close the box.
7. Repeat for the host's second VMkernel adapter, e.g. FCoE2.

8. Repeat the steps above for the remaining ESXi hosts.

4.5.6 Add software FCoE adapters to ESXi Hosts

To add FCoE storage adapters:

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. In the left pane, select the first ESXi host. In the center pane, select **Configure > Storage > Storage Adapters**.
3. Under **Storage Adapters**, click the + icon and select **Software FCoE adapter** to open the **Add Software FCoE Adapter** dialog box. The first vmnic configured for FCoE is automatically selected.
4. Click **OK** to create the FCoE adapter.
5. Click the refresh icon to refresh the storage adapter view. The FCoE adapter appears under **QLogic 57810 10 Gigabit Ethernet Multi Function Adapter.**
6. Repeat steps 3-5 to add the second FCoE adapter.

When complete, two FCoE storage adapters, named vmhba33 and vmhba34 in this example, are listed as shown in Figure 21.

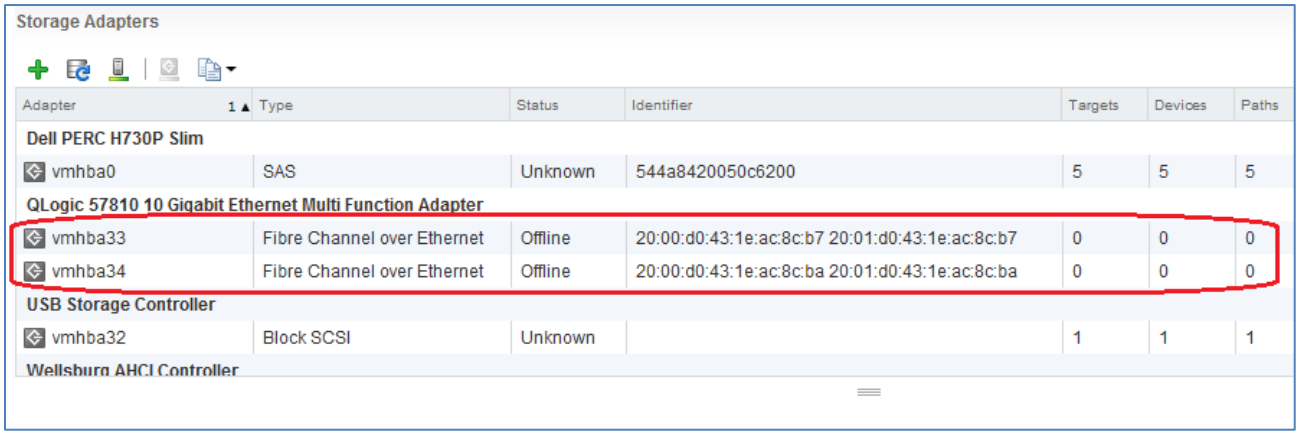

Figure 21 FCoE adapter created on ESXi host

Repeat the steps above for the remaining ESXi hosts.

5 Configure Dell EMC Networking switches

This section contains Ethernet switch configuration details with explanations for one switch in each role. Configuration files for both FN2210S switches in Full-switch mode and both S4048-ON switches are provided as attachments.

5.1 Factory default settings

The configuration commands in the sections below assume switches start at their factory default settings. All Dell EMC Networking switches in this guide can be reset to factory defaults as follows:

```
switch#restore factory-defaults stack-unit unit# clear-all
Proceed with factory settings? Confirm [yes/no]:yes
```

Factory settings are restored and the switch reloads. After reload, enter **A** at the [A/C/L/S] prompt as shown below to exit Bare Metal Provisioning mode.

```
This device is in Bare Metal Provisioning (BMP) mode.
To continue with the standard manual interactive mode, it is necessary to
abort BMP.

Press A to abort BMP now.
Press C to continue with BMP.
Press L to toggle BMP syslog and console messages.
Press S to display the BMP status.
[A/C/L/S]:A

% Warning: The bmp process will stop ...

Dell>
```

The switch is now ready for configuration.

Note: FN IOMs can also be reset to factory defaults using the Dell Blade I/O Manager GUI. This is covered in the attached document, **Dell Blade IO Manager v1.0.pdf**.

5.2 FN2210S Configuration

This section details FN2210S configuration in Standalone and Full-switch modes. Either mode may be used for this deployment.

The first section, Standalone mode configuration, covers the switch's default mode and highlights the low-touch benefits of the FN2210S.

The second section, Full-switch mode configuration, shows how to manually configure an FN2210S for this deployment. This mode allows for more granular control over the FN2210S and allows the IOM to behave more like a traditional switch.

5.2.1 Standalone mode configuration

The benefit of Standalone mode is ease of configuration.

Note: Standalone mode configuration can be done from the CLI or the Dell Blade I/O Manager GUI. CLI steps are covered in this section. See the attached document, **Dell Blade IO Manager v1.0.pdf**, to use the GUI.

If the FN2210S switches have been restored to their factory default settings, they will be in **Standalone mode**. This can be verified from the CLI with the command:

```
#show system stack-unit 0 iom-mode

Unit    Boot-Mode               Next-Boot
------------------------------------------------
0        standalone             standalone
```

No configuration is necessary on the FN2210S in Standalone mode (when at its factory default settings) to use the solution in this guide. All required features have either been preconfigured or are automatically configured. This includes:

- FCoE and Data Center Bridging (DCB) are auto-configured.
- The switch is in NPG mode and Fibre Channel ports fi 0/9 and fi 0/10 are preconfigured.
- External Ethernet ports te 0/11 and te 0/12 are configured in LACP port-channel 128.
- All Ethernet ports are in all VLANs except VLAN 1002 which is reserved for FCoE traffic. FCoE ports are automatically added to VLAN 1002.

Note: Uplink Failure Detection (UFD) is enabled on the FN2210S switches by default. Therefore, all server-facing ports (te 0/1 – te 0/8) will be down until the upstream LACP port channel 128 (consisting of ports te 0/11 and te 0/12) is up. Port channel 128 will come up when connected to a properly configured LACP port channel on an upstream switch. This is covered for the S4048-ON switches in the next section. You may also reference the attachment, **FN IOM Easy Deployment Guide v2.0.pdf**, for 3rd party upstream switch configuration examples and for more information on UFD.

Even though the FN2210S switches are fully operational in their factory default state, the following administrative commands may be run as needed:

Set the hostname:

```
Dell>enable
Dell#configure
Dell(conf)#hostname FN2210S-A1
```

Configure the serial console enable password, change the SSH/Telnet password for root and disable Telnet:

```
FN2210S-A1(conf)#enable password enable_password
FN2210S-A1(conf)#username root password ssh_password
FN2210S-A1(conf)#no ip telnet server enable
```

Note: SSH and Telnet are both enabled by default. It is a best practice to use SSH instead of Telnet for security. SSH can also be disabled with the command `(conf)#no ip ssh server enable`.

If the management IP address needs to be changed from its default Dynamic Host Configuration Protocol (DHCP) setting to a static address:

```
FN2210S-A1(conf)#interface managementethernet 0/0
FN2210S-A1(conf-if-ma-0/0)#ip address 100.67.169.91/24

Proceed with Static IP [confirm yes/no]: y
FN2210S-A1(conf-if-ma-0/0)#exit
```

Set the default gateway for the management IP address:

```
FN2210S-A1(conf)#management route 0.0.0.0/0 100.67.169.254
```

Save the configuration

```
FN2210S-A1(conf)#exit
FN2210S-A1#write
```

5.2.2 Full-switch mode configuration

The benefit of Full-switch mode is the ability to customize your configuration and the availability of extra features.

Note: Most features must be configured from the CLI on systems in Full-switch mode.

After restoring factory default settings per Section 5.1, change both FN2210S switches from Standalone mode to Full-switch mode as follows:

```
Dell>enable
Dell#configure
Dell(conf)#stack-unit 0 iom-mode full-switch
Dell(conf)#exit
Dell#write
Dell#reload
```

Answer **yes** when prompted to reload the first FN2210S. Repeat on the second FN2210S.

After the FN2210S switches have reloaded, verify both are in Full-switch mode with the command:

```
Dell#show system stack-unit 0 iom-mode
```

The command output is as follows:

```
Unit    Boot-Mode               Next-Boot
------------------------------------------------
0        full-switch            full-switch
```

Note: The following configuration details are specific to switch FN2210S-A1. The configuration for FN2210S-A2 is similar. See the **FN2210S-A1-fullswitchmode.txt** and **FN2210S-A2-fullswitchmode.txt** attachments for both switches. CLI prompts are not shown to allow for copying and pasting directly into a text editor or switch console. Commands should be run in the order shown.

Note: On FN IOMs, SSH and Telnet are both enabled by default. It is a best practice to use SSH instead of Telnet for security. SSH can optionally be disabled with the command `(conf)#`**`no ip ssh server enable`**.

Configure the serial console enable password and change the default SSH/Telnet password for root. Disable Telnet.

```
enable
configure

enable sha256-password enable_password
username root sha256-password ssh_password
no ip telnet server enable
```

Set the hostname, enable the FC feature and Link Layer Discovery Protocol (LLDP).

```
hostname FN2210S-A1
feature fc
protocol lldp
advertise management-tlv management-address system-name
```

By default, the management interface is configured for DHCP. If a static IP address is required, run the following commands to configure the address and default gateway. Substitute the addresses shown as needed for your network.

```
interface ManagementEthernet 0/0
ip address 100.67.169.91/24
yes
no shutdown

management route 0.0.0.0/0 100.67.169.254
```

Configure the DCB map. This example guarantees a minimum of 50% available bandwidth for FCoE traffic and 50% for all other traffic. The FCoE map is configured and the FCoE VLAN is set to 1002.

```
dcb-map SAN_DCB_MAP
priority-group 0 bandwidth 50 pfc off
priority-group 1 bandwidth 50 pfc on
priority-pgid 0 0 0 1 0 0 0 0
```

```
interface vlan 1002
description FCoE VLAN
no shut

fcoe-map SAN_FABRIC
fc-map 0efc00
fabric-id 1002 vlan 1002
exit
```

Configure the eight internal server-facing interfaces.

```
interface range tengigabitethernet 0/1-8
description To CNA
mtu 12000
no switchport
portmode hybrid
switchport
fcoe-map SAN_FABRIC
dcb-map SAN_DCB_MAP
no shutdown
```

Configure the external Fibre Channel and Ethernet interfaces and upstream port channel.

```
interface range fibrechannel 0/9-10
description To FC Storage
fabric SAN_FABRIC
no shutdown

interface range tengigabitethernet 0/11-12
description To S4048-ON
port-channel-protocol LACP
port-channel 128 mode active
no shutdown

interface port-channel 128
description To S4048-ON
no switchport
portmode hybrid
switchport
no shutdown
```

Configure UFD.

```
uplink-state-group 1
enable
downstream TenGigabitEthernet 0/1-8
upstream Port-channel 128
```

Save the configuration.

```
end
write
```

Note: Commands to validate the FN2210S configuration are covered in Section 7, after the S4048-ON and 6510 switches have been configured.

5.3 S4048-ON configuration and validation

The S4048-ON switches are part of the LAN topology for TCP/IP traffic on the production network. No FC or FCoE traffic is passed through these switches.

In this deployment example, two S4048-ON top-of-rack (ToR) switches are configured as a VLT pair. This provides redundancy and multiple active paths. They are leaf switches in the data center's leaf-spine network.

Note: Spine switch configuration is not covered in this guide. For a detailed configuration of a leaf-spine network, refer to the Dell EMC Leaf-Spine Deployment and Best Practices Guide.

This configuration example using redundant S4048-ON switches with VLT is a best practice. However, any ToR switch that supports port channel creation with LACP may be used for this solution. See the attachment, **FN IOM Easy Deployment Guide v2.0.pdf**, for switch configuration examples including Dell EMC Networking, Arista, Cisco, and Brocade.

5.3.1 S4048-ON configuration

The following section outlines the configuration commands issued to the S4048-ON switches. The switches start at their factory default settings per Section 5.1.

Note: The following configuration details are specific to S4048-ON-LeafA using the port numbering shown in Figure 11. The configuration for S4048-ON-LeafB is similar. See the **S4048-ON-LeafA.txt** and **S4048-ON-LeafB.txt** attachments for both switches. CLI prompts are not shown to allow for copying and pasting directly into a text editor or switch console. Commands should be run in the order shown.

Configure the serial console enable password and disable Telnet.

```
enable
configure
enable sha256-password enable_password
no ip telnet server enable
```

Note: On S4048-ON, Telnet is enabled and SSH is disabled by default. Both services require the creation of a non-root user account to login. If needed, it is a best practice to use SSH instead of Telnet for security. SSH can optionally be enabled with the command: `(conf)#ip ssh server enable.`
A user account can be created to access the switch via SSH with the command
`(conf)#username ssh_user sha256-password ssh_password`

Set the hostname, enable LLDP and configure the management interface and default gateway.

```
hostname S4048-ON-LeafA
protocol lldp
advertise management-tlv management-address system-description system-name
advertise interface-port-desc


interface manangementethernet 1/1
ip address 100.67.169.31/24
no shutdown

management route 0.0.0.0/0 100.67.169.254
```

Configure the VLT interconnect (VLTi) between Leaf A and Leaf B. The VLTi in this example uses interfaces fortyGigE 1/53-54 in static port-channel 127. The backup destination is the management IP address of the VLT peer switch, S4048-ON-Leaf B.

```
interface Port-channel 127
description VLTi
channel-member fortyGigE 1/53 - 1/54
no shutdown

interface range fortyGigE 1/53 - 1/54
description VLTi
no shutdown

vlt domain 127
peer-link port-channel 127
back-up destination 100.67.169.30
unit-id 0
exit
```

Configure the downstream interfaces to the FN2210S switches. Each interface is added to a port channel and the port channels are configured for VLT.

```
interface TenGigabitEthernet 1/1
description To FN2210S-A1
port-channel-protocol LACP
port-channel 101 mode active
no shutdown

interface TenGigabitEthernet 1/2
description To FN2210S-A2
port-channel-protocol LACP
port-channel 102 mode active
no shutdown

interface Port-channel 101
description To FN2210S-A1
portmode hybrid
switchport
vlt-peer-lag port-channel 101
no shutdown

interface Port-channel 102
description To FN2210S-A2
portmode hybrid
switchport
vlt-peer-lag port-channel 102
no shutdown
```

Save the configuration.

```
end
write
```

5.3.2 S4048-ON validation

After the S4048-ON and FN2210S switches have been configured, VLT is fully functional. It can be verified with the commands shown below.

Note: The commands and output shown below are for S4048-ON-LeafA. The output for S4048-ON-LeafB is similar.

5.3.2.1 show vlt brief

Run the `show vlt brief` command on each of the S4048-ON switches. The Inter-chassis link (ICL) Link Status, Heart Beat Status, and VLT Peer Status must be up. The role for one switch in the VLT pair will be primary and its peer switch (not shown) will be assigned the secondary role.

```
S4048-ON-LeafA#show vlt brief
 VLT Domain Brief
------------------
 Domain ID:                    127
 Role:                         Primary
 Role Priority:                32768
 ICL Link Status:              Up
 HeartBeat Status:             Up
 VLT Peer Status:              Up
 Local Unit Id:                1
 Version:                      6(7)
 Local System MAC address:     14:18:77:7c:c4:e8
 Remote System MAC address:    14:18:77:e0:69:31
 Remote system version:        6(7)
 Delay-Restore timer:          90 seconds
 Delay-Restore Abort Threshold: 60 seconds
 Peer-Routing :                Disabled
 Peer-Routing-Timeout timer:   0 seconds
 Multicast peer-routing timeout: 150 seconds
```

5.3.2.2 show vlt detail

Run the `show vlt detail` command on at least one S4048-ON switch in the VLT pair. The output indicates downstream port channels 101 and 102 connected to FN2210S-A1 and -A2 are up on the local switch (LeafA) and the VLT peer switch (LeafB).

```
S4048-ON-LeafA#show vlt detail
Local LAG Id  Peer LAG Id  Local Status  Peer Status  Active VLANs
------------  -----------  ------------  -----------  -------------
101           101          UP            UP            1
102           102          UP            UP            1
```

5.4 S3048-ON management switch configuration

For the S3048-ON management switches, all ports used are in layer 2 mode and are in the default VLAN. No additional configuration is required.

6 6510 configuration and validation

This section covers configuration of 6510 FC switches using the serial console CLI. 6510 switches also have a Java-based GUI available that may be used for zone configuration as an alternative. For more information on switch configuration and commands, refer to the 6510 documentation listed in Appendix E.

6.1 Reset to defaults

The commands in this guide are based on 6510 switches starting at their factory default settings. For example, the N_Port ID Virtualization (NPIV) feature must be enabled for this deployment. Resetting 6510 switches to defaults ensures NPIV is enabled and helps ensure other key settings are configured properly.

To restore system configuration parameters to default values, run the following commands on each 6510:

```
6510-1:admin> chassisdisable
Are you sure you want to disable all chassis ports now? (yes, y, no, n): [no] y
6510-1:admin> configdefault -all
WARNING:  This is a disruptive operation that requires a switch reboot.
Would you like to continue [Y/N]: y
```

Note: Some configured parameters are not affected by the `configdefault -all` command such as the switch name, IP address and zone configuration settings. For more information about this command, see the Brocade Fabric OS Command Reference Guide listed in Appendix E.3.

6.2 6510 configuration

The following section outlines the configuration commands run at the `admin>` prompt on the 6510 switches.

Note: Commands that wrap the page are `underlined` below. The following configuration details are specific to 6510-1. The configuration for 6510-2 is similar. See the **6510-1.txt** and **6510-2.txt** attachments for all configuration commands. Prompts are not shown to allow for copying and pasting directly into a text editor or switch console. Commands should be run in the order shown.

Configure the switch name, management IP address and default gateway.

```
switchname 6510-1
ipaddrset -ipv4 -add -ethip 100.67.169.10 -ethmask 255.255.255.0 -gwyip
100.67.169.254 -dhcp off
```

Enable the ports connected to FN IOMs and storage (shown in Figure 10.)

```
portcfgpersistentenable 8-11
```

Create aliases to use in zoning. Aliases are used to associate a user-friendly name to a WWPN. They are optional but can simplify working with WWPNs.

Create aliases for the four FC630 server FCoE ports connected to FN2210S-A1. This information is from Table 2 in Section 4.3.

```
alicreate "FC630_1_A1", "20:01:D0:43:1E:AC:8C:B7"
alicreate "FC630_2_A1", "20:01:D0:43:1E:AC:8D:6E"
alicreate "FC630_3_A1", "20:01:D0:43:1E:AC:8E:25"
alicreate "FC630_4_A1", "20:01:D0:43:1E:AC:8E:DC"
```

Create aliases for the two storage ports connected to 6510-1. This information is from Table 3 in Section 4.4.

```
alicreate "SPA_4", "50:06:01:62:47:E0:1B:19"
alicreate "SPB_4", "50:06:01:6A:47:E0:1B:19"
```

Next, the zones are created. Zones enable the SAN to be partitioned into groups of devices that can access each other. Dell EMC recommends creating a separate zone for each server CNA FCoE port. Using the aliases created above, a zone is created for each server port and both storage ports.

```
zonecreate "zone_FC630_1_A1", "FC630_1_A1;SPA_4;SPB_4"
zonecreate "zone_FC630_2_A1", "FC630_2_A1;SPA_4;SPB_4"
zonecreate "zone_FC630_3_A1", "FC630_3_A1;SPA_4;SPB_4"
zonecreate "zone_FC630_4_A1", "FC630_4_A1;SPA_4;SPB_4"
```

The four zones created above are added to a zone configuration named ZoneConfig1.

```
cfgcreate
"ZoneConfig1","zone_FC630_1_A1;zone_FC630_2_A1;zone_FC630_3_A1;zone_FC630_4_A1"
```

Note: Multiple zone configurations may be created on the switch, but only one zone configuration can be active at a time.

The configuration is not complete until the following two commands are run. Save the configuration with the `cfgsave` command and enable ZoneConfig1 with the `cfgenable` command. These commands are interactive and are answered manually as shown below.

```
6510-1:admin> cfgsave
Do you want to save the Defined zoning configuration only?  (yes, y, no, n):
[no] yes
sw0 Updating flash ...
2017/02/15-22:43:26, [ZONE-1024], 124, FID 128, INFO, 6510-1, cfgSave
completes successfully.

6510-1:admin> cfgenable ZoneConfig1
Do you want to enable 'ZoneConfig1' configuration  (yes, y, no, n): [no] yes
sw0 Updating flash ...
2017/02/15-22:47:58, [ZONE-1022], 126, FID 128, INFO, 6510-1, The effective
configuration has changed to ZoneConfig1.
zone config "ZoneConfig1" is in effect
```

6.3 6510 validation

The following commands may be run to verify ports are up and zoning is configured properly.

Note: The commands and output shown below are for 6510-1. The output for 6510-2 is similar.

6.3.1 switchshow

The `switchshow` command shows the port state and connected port information.

```
6510-1:admin> switchshow
switchName:     6510-1
switchType:     109.1
switchState:    Online
switchMode:     Native
switchRole:     Principal
switchDomain:   1
switchId:       fffc01
switchWwn:      10:00:c4:f5:7c:50:9d:0f
zoning:         ON (ZoneConfig1)
switchBeacon:   OFF
FC Router:      OFF
FC Router BB Fabric ID: 1
Address Mode:   0
HIF Mode:       OFF

Index Port Address  Media Speed   State       Proto
==================================================
   8   8   010800   id    N16     Online      FC  F-Port  50:06:01:6a:47:e0:1b:19
   9   9   010900   id    N16     Online      FC  F-Port  50:06:01:62:47:e0:1b:19
  10  10   010a00   id    N8      Online      FC  F-Port  1 N Port + 2 NPIV public
  11  11   010b00   id    N8      Online      FC  F-Port  1 N Port + 2 NPIV public
```

Note: Unused ports have been removed from the command output above for brevity.

On each 6510 switch, ports 8-9 are connected to the Unity storage array and the corresponding Unity WWPNs are shown in the far right column of the command output.

Note: With some FC storage arrays, storage ports may appear as NPIV ports.

6510-1 ports 10-11 are connected to FN2210S-A1. The output shows a total of four NPIV ports, one for each connected server adapter.

To view the WWPNs of devices on ports using NPIV (such as ports 10-11 in this example) use the `portshow` *`port#`* command.

6.3.2 zoneshow

The `zoneshow` command shows configured zone and alias information. Most importantly, the effective zone configuration is summarized at the end of the output.

```
6510-1:admin> zoneshow
Defined configuration:
 cfg:   ZoneConfig1
                zone_FC630_1_A1; zone_FC630_2_A1; zone_FC630_3_A1;
                zone_FC630_4_A1
 zone:  zone_FC630_1_A1
                FC630_1_A1; SPA_4; SPB_4
 zone:  zone_FC630_2_A1
                FC630_2_A1; SPA_4; SPB_4
 zone:  zone_FC630_3_A1
                FC630_3_A1; SPA_4; SPB_4
 zone:  zone_FC630_4_A1
                FC630_4_A1; SPA_4; SPB_4
 alias: FC630_1_A1
                20:01:d0:43:1e:ac:8c:b7
 alias: FC630_2_A1
                20:01:d0:43:1e:ac:8d:6e
 alias: FC630_3_A1
                20:01:d0:43:1e:ac:8e:25
 alias: FC630_4_A1
                20:01:d0:43:1e:ac:8e:dc
 alias: SPA_4   50:06:01:62:47:e0:1b:19
 alias: SPB_4   50:06:01:6a:47:e0:1b:19

Effective configuration:
 cfg:   ZoneConfig1
 zone:  zone_FC630_1_A1
                20:01:d0:43:1e:ac:8c:b7
                50:06:01:62:47:e0:1b:19
                50:06:01:6a:47:e0:1b:19
 zone:  zone_FC630_2_A1
                20:01:d0:43:1e:ac:8d:6e
                50:06:01:62:47:e0:1b:19
                50:06:01:6a:47:e0:1b:19
 zone:  zone_FC630_3_A1
                20:01:d0:43:1e:ac:8e:25
                50:06:01:62:47:e0:1b:19
                50:06:01:6a:47:e0:1b:19
 zone:  zone_FC630_4_A1
                20:01:d0:43:1e:ac:8e:dc
                50:06:01:62:47:e0:1b:19
                50:06:01:6a:47:e0:1b:19
```

7 FN2210S IOM validation

After connected devices are configured, several commands may be run on the FN2210S IOMs to validate the network configuration. These commands are the same on the FN2210S IOMs in Standalone or Full-switch modes.

Note: The commands and output shown below are for FN2210S-A1. The output for FN2210S-A2 is similar.

7.1 show interface port channel 128

Port channel 128 on FN2210S-A1 is connected to VLT port channel 101 on both S4048-ON leaf switches. On FN2210S-A2, port channel 128 is connected to VLT port channel 102. It must be up on both FN2210S switches with a line speed of 20000 Mbit. It consists of two members, Te 0/11 and Te 0/12.

```
FN2210S-A1#show interfaces port-channel 128
Port-channel 128 is up, line protocol is up
Created by LACP protocol
Hardware address is f8:b1:56:76:ba:db, Current address is f8:b1:56:76:ba:db
Interface index is 1258356736
Minimum number of links to bring Port-channel up is 1
Internet address is not set
Mode of IPv4 Address Assignment : NONE
DHCP Client-ID :f8b15676badb
MTU 12000 bytes, IP MTU 11982 bytes
LineSpeed 20000 Mbit
Members in this channel:  Te 0/11(U) Te 0/12(U)
```

(output truncated)

7.2 show interfaces status

When port channel 128 is up, UFD enables the server-facing interfaces, Te 0/1, 0/3, 0/5 and 0/7 in this example. Fibre Channel ports 9 and 10 are connected to 6510-1 and -2 and are up at 8000 Mbit (8 Gbps).

Note: Server-facing interfaces Te 0/2, 0/4, 0/6 and 0/8 are not used in this deployment example. They are reserved for quad-port CNAs or for quarter-width servers such as the FC430.

```
FN2210S-A1#show interfaces status
Port                Description  Status Speed        Duplex Vlan
Te 0/1                           Up     10000 Mbit   Full   1-4094
Te 0/2                           Down   Auto         Auto   1-1001,1003-4094
Te 0/3                           Up     10000 Mbit   Full   1-4094
Te 0/4                           Down   Auto         Auto   1-1001,1003-4094
Te 0/5                           Up     10000 Mbit   Full   1-4094
Te 0/6                           Down   Auto         Auto   1-1001,1003-4094
Te 0/7                           Up     10000 Mbit   Full   1-4094
Te 0/8                           Down   Auto         Auto   1-1001,1003-4094
```

```
Fc 0/9                              Up      8000 Mbit    Full   --
Fc 0/10                             Up      8000 Mbit    Full   --
Te 0/11                             Up      10000 Mbit   Full   --
Te 0/12                             Up      10000 Mbit   Full   --
```

7.3 show fc switch

The show fc switch command verifies the FN IOM is in NPG mode for FC traffic.

```
FN2210S-A1#show fc switch
Switch Mode : NPG
Switch WWN  : 10:00:f8:b1:56:76:ba:dd
```

7.4 show fip-snooping sessions

After VMware software FCoE adapters are added (Section 4.5.6) and 6510 switches are configured (Section 6.2), fip-snooping sessions start on each FN2210S. Run the `show fip-snooping sessions` command on both FN2210S IOMs to verify functionality.

Note: Due to the width of the `show fip-snooping sessions` command output, each line of output is shown on two lines below.

```
FN2210S-A1#show fip-snooping sessions
Enode MAC           Enode Intf     FCF MAC             FCF Intf       VLAN
FCoE MAC            FC-ID     Port WWPN                 Port WWNN

d0:43:1e:ac:8c:b7   Te 0/1         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:02   01:0a:02  20:01:d0:43:1e:ac:8c:b7   20:00:d0:43:1e:ac:8c:b7

d0:43:1e:ac:8d:6e   Te 0/3         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:01   01:0a:01  20:01:d0:43:1e:ac:8d:6e   20:00:d0:43:1e:ac:8d:6e

d0:43:1e:ac:8e:25   Te 0/5         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:03   01:0b:03  20:01:d0:43:1e:ac:8e:25   20:00:d0:43:1e:ac:8e:25

d0:43:1e:ac:8e:dc   Te 0/7         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:02   01:0b:02  20:01:d0:43:1e:ac:8e:dc   20:00:d0:43:1e:ac:8e:dc
```

There is one fip-snooping session for each server FCoE adapter port (four on FN2210S-A1 and four on FN2210S-A2). Server FCoE adapter WWPNs are shown in the `Port WWPN` column. The FCF interface used, Fc 0/9 or Fc 0/10, is whichever connects first.

In the above example, servers FC630-1 and FC630-2, connected to Te 0/1 and Te 0/3 respectively, use interface Fc 0/9. Servers FC630-3 and FC630-4, connected to Te 0/5 and Te 0/7 respectively, use interface Fc 0/10. The Fc interface shown in the command output is active for that adapter. The unused Fc interface is standby.

To demonstrate a simple failover, if interface Fc 0/9 in this example is shut down, fip-snooping sessions on ports Te 0/1 and Te 0/3 are cleared and new fip-snooping sessions are established using Fc 0/10. Fip-snooping sessions on ports Te 0/5 and Te 0/7 remain established using Fc 0/10. This is shown in the following two commands:

```
FN2210S-A1(conf-if-fc-0/9)#shutdown

FN2210S-A1#show fip-snooping sessions
Enode MAC           Enode Intf     FCF MAC             FCF Intf       VLAN
FCoE MAC            FC-ID     Port WWPN                  Port WWNN

d0:43:1e:ac:8c:b7   Te 0/1         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:05   01:0b:05  20:01:d0:43:1e:ac:8c:b7   20:00:d0:43:1e:ac:8c:b7

d0:43:1e:ac:8d:6e   Te 0/3         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:04   01:0b:04  20:01:d0:43:1e:ac:8d:6e   20:00:d0:43:1e:ac:8d:6e

d0:43:1e:ac:8e:25   Te 0/5         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:03   01:0b:03  20:01:d0:43:1e:ac:8e:25   20:00:d0:43:1e:ac:8e:25

d0:43:1e:ac:8e:dc   Te 0/7         f8:b1:56:76:ba:e1   Fc 0/10        1002
0e:fc:00:01:0b:02   01:0b:02  20:01:d0:43:1e:ac:8e:dc   20:00:d0:43:1e:ac:8e:dc
```

Interface Fc 0/9 is brought back up as follows:

```
FN2210S-A1(conf-if-fc-0/9)#no shutdown
```

Interface Fc 0/9 is now up, but existing fip-snooping sessions do not switch back to Fc 0/9, they remain connected to Fc 0/10 as long as connections are functional.

If interface Fc 0/10 is shut down, fip-snooping sessions are cleared and new fip-snooping sessions are established using Fc 0/9. This is shown in the following two commands:

```
FN2210S-A1(conf-if-fc-0/10)#shutdown

FN2210S-A1#show fip-snooping sessions
Enode MAC           Enode Intf     FCF MAC             FCF Intf       VLAN
FCoE MAC            FC-ID     Port WWPN                  Port WWNN


d0:43:1e:ac:8c:b7   Te 0/1         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:04   01:0a:04  20:01:d0:43:1e:ac:8c:b7   20:00:d0:43:1e:ac:8c:b7

d0:43:1e:ac:8d:6e   Te 0/3         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:03   01:0a:03  20:01:d0:43:1e:ac:8d:6e   20:00:d0:43:1e:ac:8d:6e

d0:43:1e:ac:8e:25   Te 0/5         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:01   01:0a:01  20:01:d0:43:1e:ac:8e:25   20:00:d0:43:1e:ac:8e:25

d0:43:1e:ac:8e:dc   Te 0/7         f8:b1:56:76:ba:e0   Fc 0/9         1002
0e:fc:00:01:0a:02   01:0a:02  20:01:d0:43:1e:ac:8e:dc   20:00:d0:43:1e:ac:8e:dc
```

8 Configure FC storage

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

8.1 Create a Storage Pool

1. Connect to the Unisphere GUI in a web browser and log in.
2. In the left pane under **STORAGE**, select **Pools**.
3. Click the + icon. In the **Create Pool** dialog box, provide a **Name** and click **Next**.
4. Select appropriate storage tiers for the pool. Change the RAID configuration for the chosen tiers if needed and click **Next**.
5. In the **Drives** section, select the **Amount of Storage.** The total number of disks and the total capacity will be displayed next to **Totals** label. Click **Next**.
6. Leave the **Capability Profile Name** section as is and click **Next**.
7. Review your selections in the **Summary** section and click **Finish**. Once the **Overall status** shows 100%, click **Close**.
8. The newly created Storage Pool is now visible under **Pools** section, as shown in Figure 22.

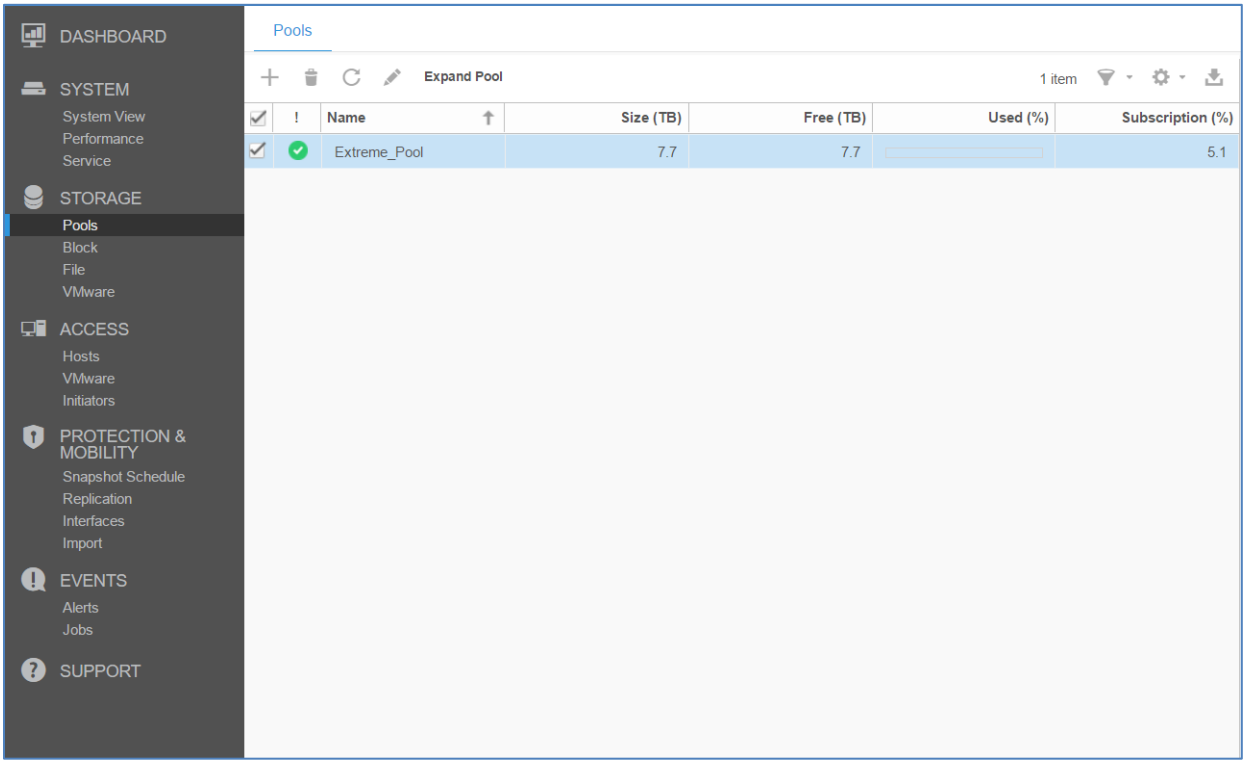

Figure 22 Storage pool created

8.2 Add ESXi Hosts

1. In the Unisphere left pane under **ACCESS**, select **VMware.**
2. On the **vCenters** tab, click the + icon to open the **Add vCenter** dialog box.
3. Enter the **Network Name or Address** of the vCenter server. Enter the vCenter **User name** and **Password** and click **Find**.
4. A list of **ESXi hosts** that can be added from vCenter is displayed. Select the applicable ESXi hosts and click **Next**.
5. On the **Summary** page, review the ESXi Hosts to be added. Click **Finish**.
6. When the **Overall status** shows **100% Completed**, click **Close**.
7. The vCenter server is displayed as shown in Figure 23.

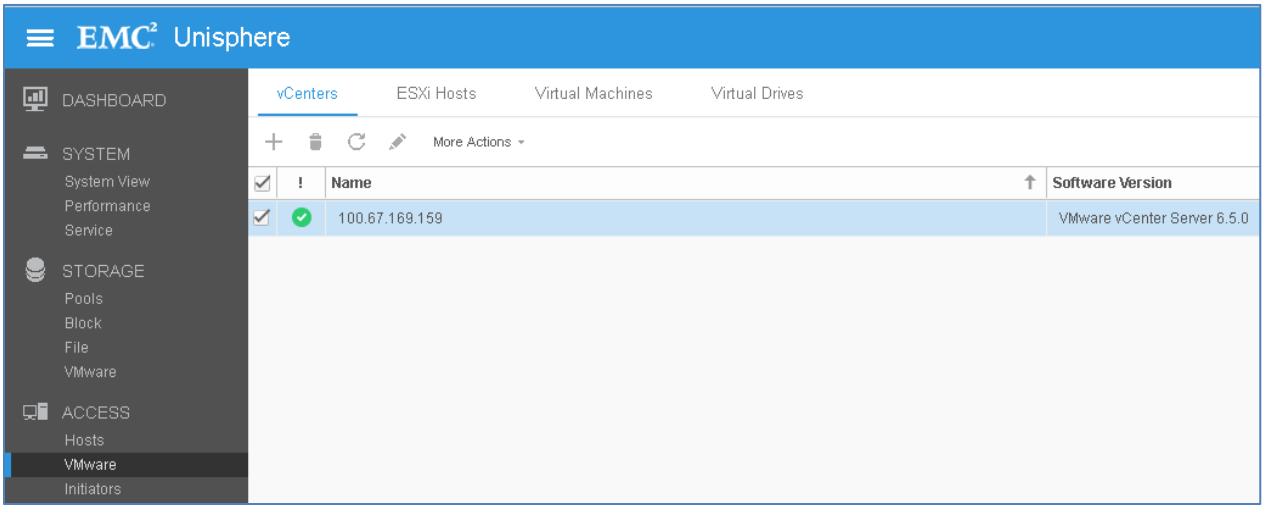

Figure 23 vCenter server added to Unisphere

8. The list of added ESXi hosts is displayed under the **ESXi Hosts** tab, as shown in Figure 24.

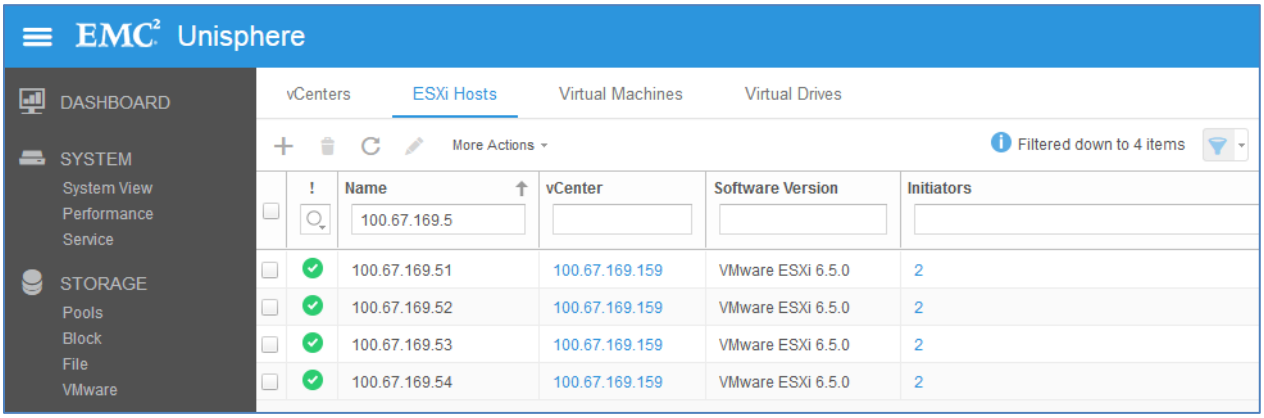

Figure 24 ESXi Hosts added to Unisphere

Note: Additional hosts may be added as needed on the **ESXi Hosts** tab. Click the + icon, enter the credentials for the vCenter Server or ESXi host, and follow the prompts displayed on the screen.

8.3 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**. Provide a **Name** and select the **Storage Pool.** Modify the **Size** as required and click **Next**.
4. On the **Access** page, click + icon and select host(s) to be granted access to the LUN. Click **OK > Next**.
5. On the **Snapshot** page, click **Next.**
6. On the **Replication** page, click **Next**.
7. On the **Summary** page, review the details and click **Finish**.
8. 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 25. In this example, a LUN named FC-51GB that is 51 GB in size has been created.

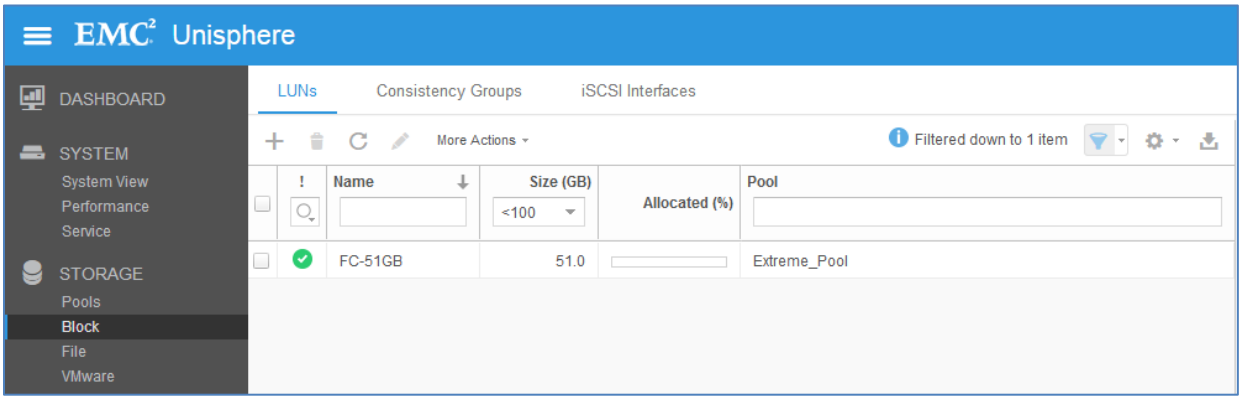

Figure 25 LUN Created

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 icon, and select the **Host Access** tab.

9 Configure storage on ESXi hosts

In this section, the example LUN created on the storage array is used to create a datastore on an ESXi host. The datastore is used to create a virtual disk on a virtual machine (VM) residing on the ESXi host. This process may be repeated as needed for additional LUNs, ESXi hosts and VMs.

9.1 Rescan storage

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. In the **Navigator** pane, select an ESXi host with LUN access configured on the FC storage array.
3. In the center pane, select **Configure > Storage Adapters**.
4. Select host's first FCoE adapter, vmhba33 in this example, and click the  icon to rescan storage.
5. The LUN on the storage array mapped to this host appears under **Adapter Details** on the **Devices** tab as shown in Figure 26.

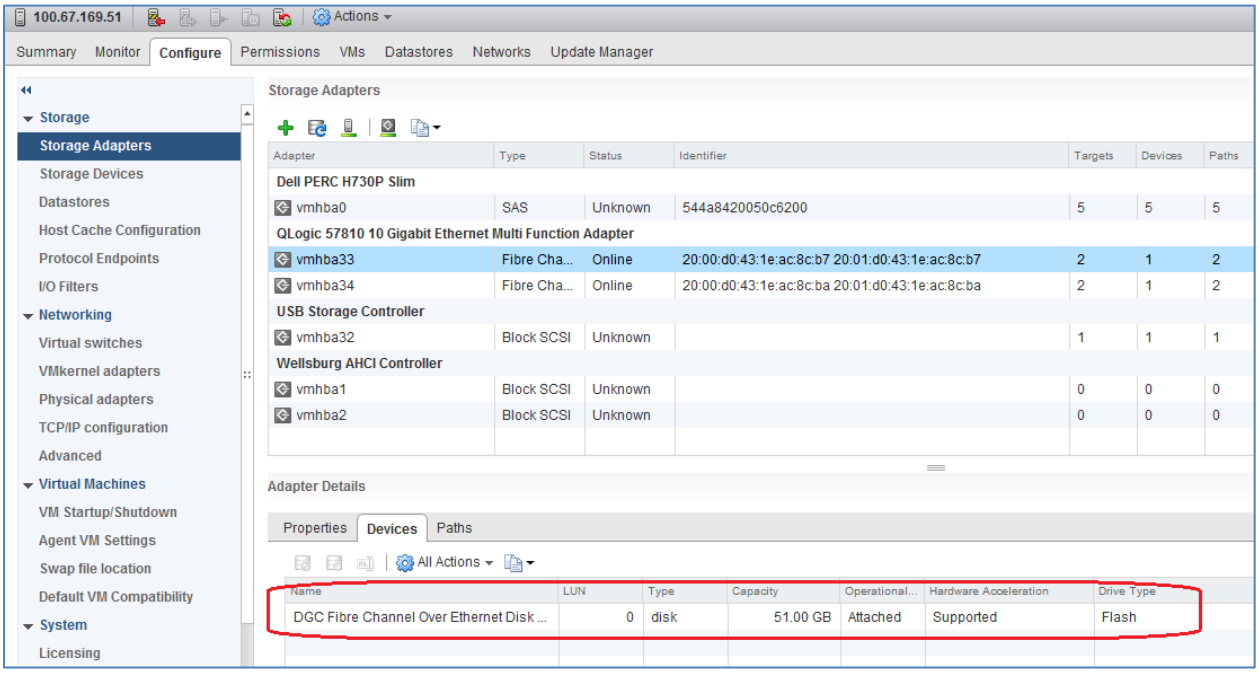

Figure 26 LUN visible to ESXi host

6. Repeat for host's second adapter, vmhba34 in this example. The LUN information on the **Devices** tab is identical to the first adapter.
7. Select the **Paths** tab as shown in Figure 27. The target, LUN number (e.g. LUN 0) and path status are shown. The target field includes the active storage WWPNs. The status field is marked either ◆ **Active** or ◆ **Active (I/O)** for each path.

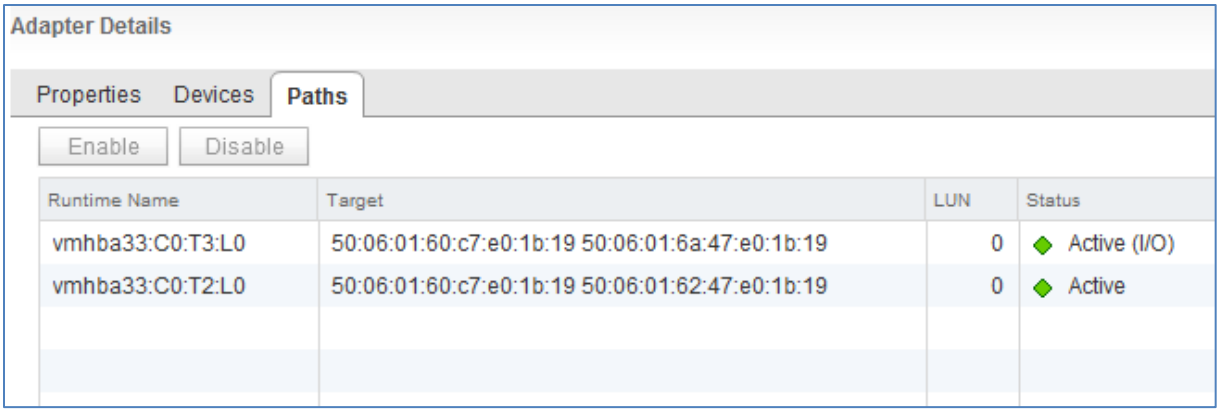

Figure 27 Adapter Details - Paths tab

9.2 Create a datastore

In this section, a datastore that uses the LUN is created on an ESXi host.

To create the datastore:

1. In the vSphere Web Client, go to **Home** > **Hosts and Clusters**.
2. In the **Navigator** pane, right click on an ESXi host and select **Storage > New Datastore**.
3. In the **New Datastore** window, leave the **Type** set to **VMFS** and click **Next**.
4. The **Name and device selection** page appears as shown in Figure 28. In this example, the 51 GB LUN mapped to this host appears in the list of devices.

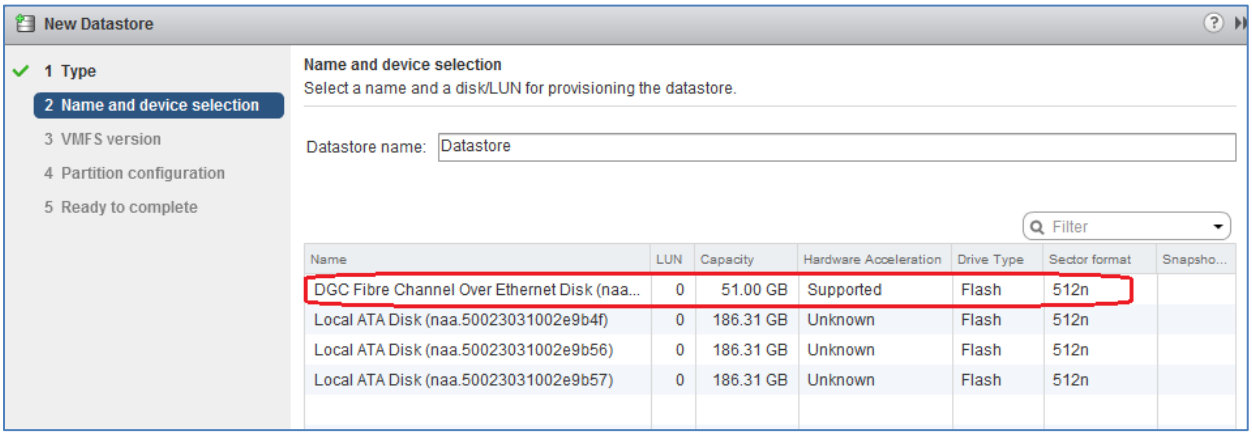

Figure 28 Name and device selection page

5. Provide a **Datastore name**, e.g. **Unity 51 GB LUN**, select the LUN in the list, and click **Next.**
6. Select the **VMFS version**. For this guide it is left at its default setting, **VMFS 5**. Click **Next.**
7. Leave the **Partition configuration** at its default settings and click **Next > Finish** to create the datastore.

The datastore can now be accessed by selecting the host in the **Navigator** pane. Select the **Configure** tab **> Datastores** as shown in Figure 29.

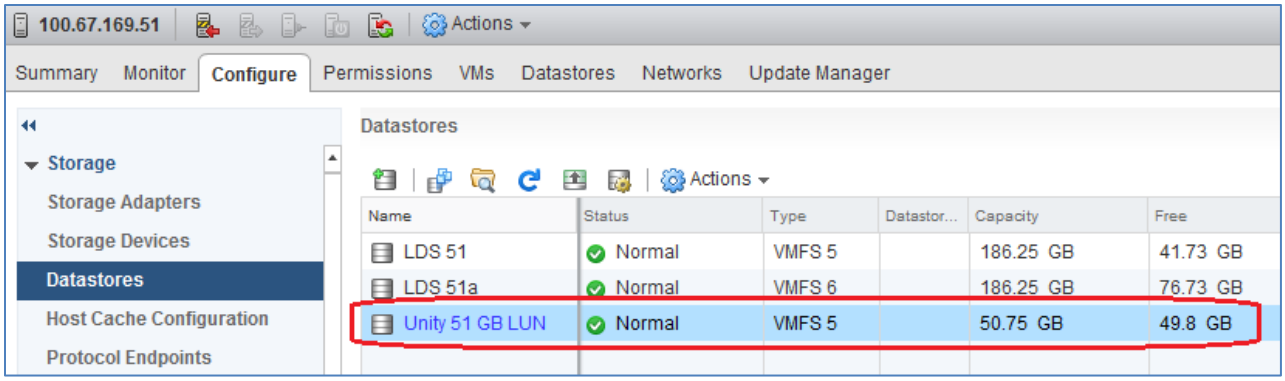

Figure 29 Datastore configured

The datastore is also accessible by going to **Home > Storage.** It is listed under the **Datacenter** object in the **Navigator** pane.

9.3 Use the datastore to create a virtual disk

Note: Virtual machine/guest operating system deployment steps are not included in this document. For instructions, see the VMware vSphere 6.5 Documentation. Guest operating systems can be any supported by ESXi 6.5.

In this example, the ESXi host with the datastore configured in the previous section contains a VM named **VM1** that is running a Windows Server guest OS.

To create a virtual disk on VM1 from the datastore:

1. Go to **Home > Hosts and Clusters**.
2. In the **Navigator** pane, right click on **VM1** and select **Edit Settings**.
3. Next to **New Device**, select **New Hard Disk** and click **Add**.
4. Click the ▶ icon next to **New Hard Disk** to view the configuration options.
5. Next to **Location,** select **Browse.** Select the previously configured datastore, e.g. **Unity 51 GB LUN**, and click **OK**. The screen will look similar to Figure 30.

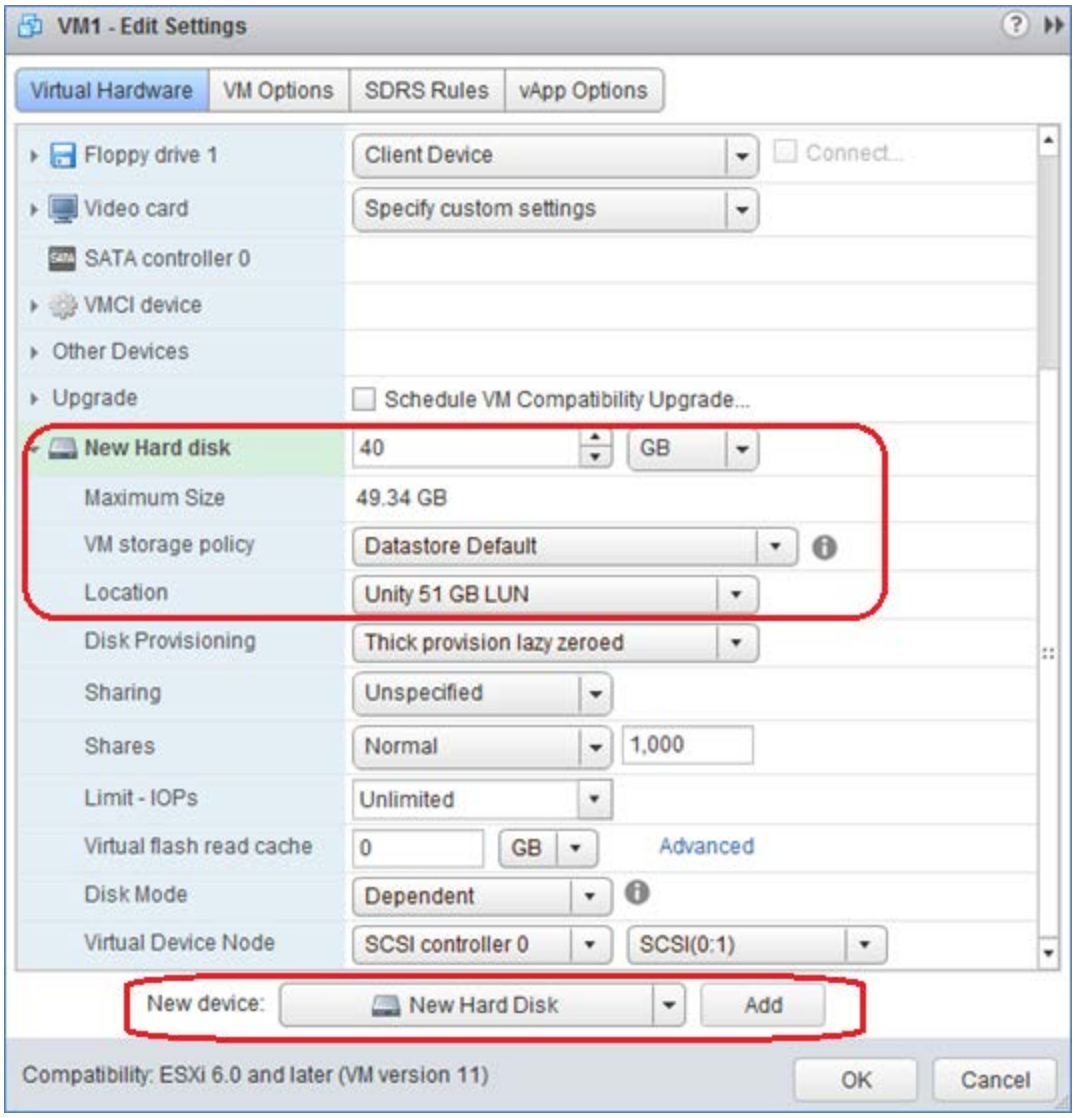

Figure 30 New hard disk configuration options

6. Next to **New Hard disk**, set the size in GB less than or equal to the **Maximum size** shown on the line below. The size is set to **40 GB** in this example.
7. Click **OK** to close the **Edit Settings** window and create the virtual disk.

9.4 Configure the virtual disk in the VM

Note: The following example is applicable for VMs running Windows Server 2008, 2012, or 2016. See your operating system documentation to configure virtual disks on other supported guest operating systems.

1. Power on the VM and log in to the Windows Server guest OS.
2. In Windows, go to **Server Manager > Tools > Computer Management > Storage > Disk Management.**
3. Right click on **Disk Management** and select **Rescan Disks**.
4. The new hard disk appears in the list (**Disk 1** in Figure 31).

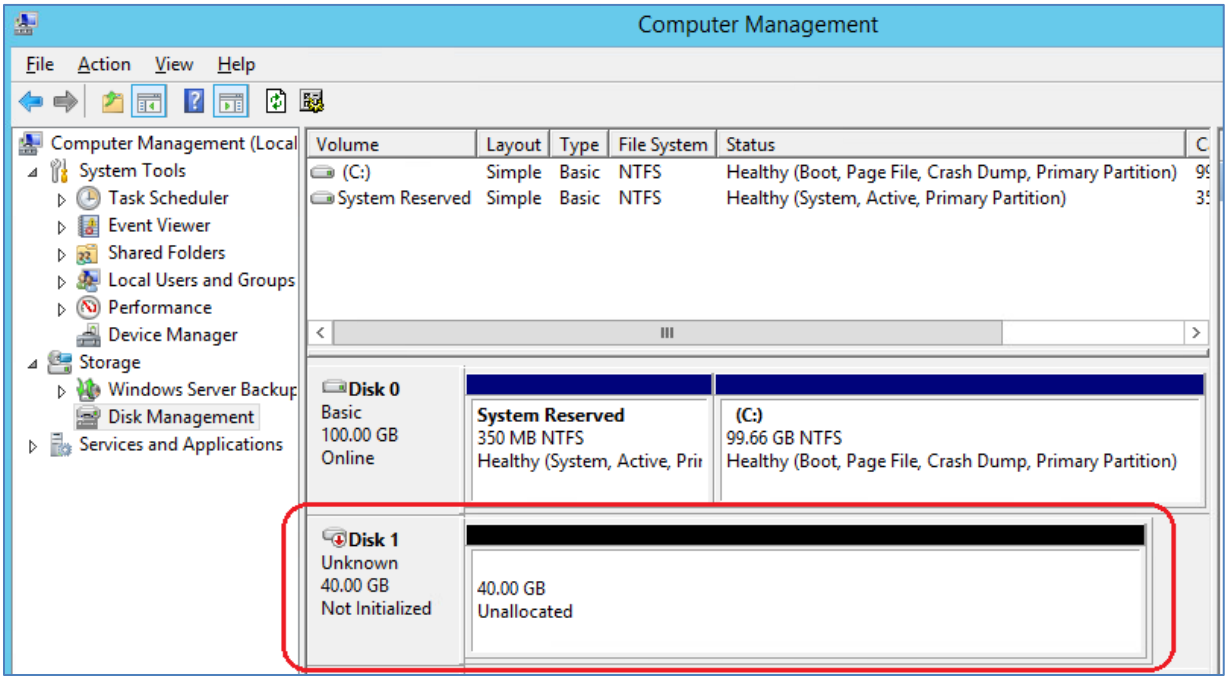

Figure 31 Windows Server Disk Management utility

5. Right click in the box containing the **Disk #**, e.g. **Disk 1**, and use the menu options to bring the disk online and initialize it. Format as needed.

10 Configure ESXi hosts for LAN traffic

In this section, ESXi hosts are configured for TCP/IP access to the production network. (Refer to the LAN topology shown in Figure 11 in section 3.2).

10.1 vSphere distributed switches

A vSphere Distributed Switch (also referred to as a VDS or a distributed switch) is a virtual switch that provides network connectivity to hosts and virtual machines. Unlike vSphere standard switches, distributed switches act as a single switch across multiple hosts in a cluster.

Distributed switches are configured in the web client and the configuration is populated across all hosts associated with the switch. They are used for connectivity to the production network in this guide.

Distributed Switches contain two different port groups:

- **Uplink port group** – an uplink port group maps physical NICs on the hosts (vmnics) to uplinks on the VDS. Uplink port groups act as trunks and carry all VLANs by default.
- **Distributed port group** - Distributed port groups define how connections are made through the VDS to the network. In this guide, one distributed port group is created for the production network.

Note: For consistent network configuration, you can connect the same vmnic on every host to the same uplink port on the distributed switch. For example, if you are adding two hosts, connect vmnic1 on each host to Uplink1 on the distributed switch.

10.2 Create a VDS

In this section a VDS named **FC630-VDS** is created for host access to the production network.

1. In the vSphere Web Client, go to **Home > Networking.**
2. Right click on **Datacenter**. Select **Distributed switch > New Distributed Switch**.
3. Provide a name for the VDS, **FC630-VDS**. Click **Next**.
4. On the **Select version** page, select **Distributed switch: 6.5.0 > Next**.
5. On the **Edit settings** page:
   a. Since a dual-port adapter (QLogic 57810) is used, set the **Number of uplinks** to **2**.
   b. Leave **Network I/O Control** set to **Enabled**.
   c. **Uncheck** the **Create a default port group** box.
6. Click **Next** followed by **Finish.**

The VDS is created with the uplinks port group shown beneath it as shown in Figure 32.

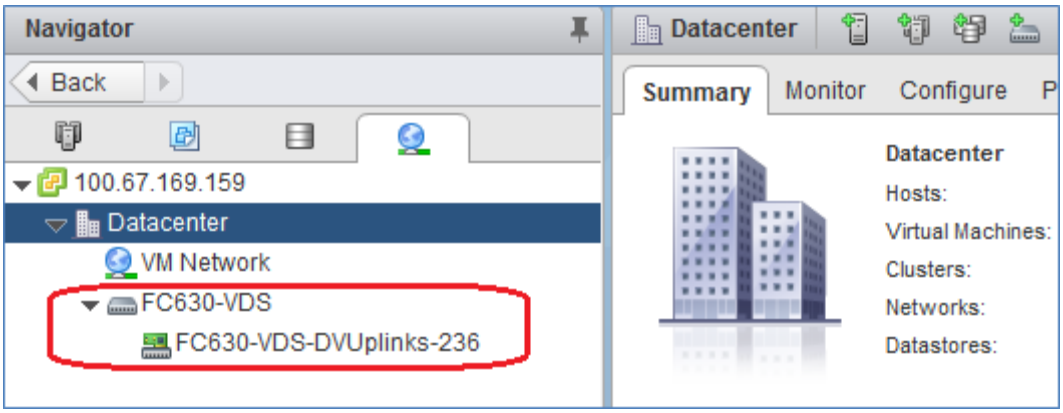

Figure 32 FC630-VDS created with Uplinks port group

10.3 Add a distributed port group

In this section, a distributed port group is created on FC630-VDS.

To create the port group:

1. On the web client **Home** screen, select **Networking**.
2. Right click on **FC630-VDS**. Select **Distributed Port Group > New Distributed Port Group.**
3. On the **Select name and location** page, provide a name for the distributed port group, for example, **Production**. Click **Next**.
4. On the **Configure settings** page, leave the values at their defaults and click **Next > Finish.**

When complete, the **Navigator** pane appears similar to Figure 33.

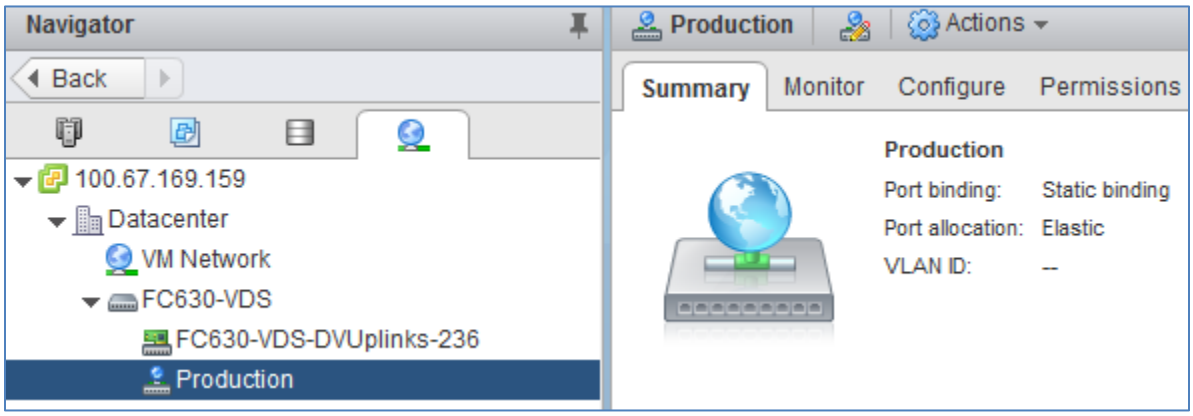

Figure 33 Production distributed port group created

10.4 Associate hosts and assign uplinks

In this section, ESXi hosts and their vmnics are associated with the VDS.

Note: Before starting this section, be sure you know the vmnic-to-physical adapter mapping for each host. This can be determined by going to **Home > Hosts and Clusters** and selecting the host in the **Navigator** pane. In the center pane go to **Configure > Networking > Physical adapters**. Vmnics and their MAC addresses are shown. Running the command `show lldp neighbors` on the FN2210S switches will display the MAC addresses of the connected QLogic adapters.
In this example, the two vmnics used on each host are numbered vmnic0 and vmnic1. Your vmnic numbering may vary.

To add hosts to FC630-VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on **FC630-VDS** and select **Add and Manage Hosts**.
3. In the **Add and Manage Hosts** dialog box:
   a. On the **Select task** page, make sure **Add hosts** is selected. Click **Next**.
   b. On the **Select hosts** page, Click the + icon. Select the check box next to each host in the FX2-FC630-FC cluster. Click **OK > Next**.
   c. On the **Select network adapters tasks** page, be sure the **Manage physical adapters** box is checked. Be sure all other boxes are unchecked. Click **Next**.
   d. On the **Manage physical network adapters** page, each host is listed with its vmnics beneath it.
      i. Select the first vmnic (vmnic0 in this example) on the first host and click the Assign uplink icon.
      ii. Select **Uplink 1 > OK**.
      iii. Select the second vmnic (vmnic1 in this example) on the first host and click the Assign uplink icon.
      iv. Select **Uplink 2 > OK**.
   e. Repeat steps i – iv for the remaining hosts. Click **Next** when done.
   f. On the **Analyze impact** page, **Overall impact status** should indicate No impact.
   g. Click **Next > Finish**.

When complete, the **Configure > Settings > Topology** page for FC630-VDS will look similar to Figure 34.

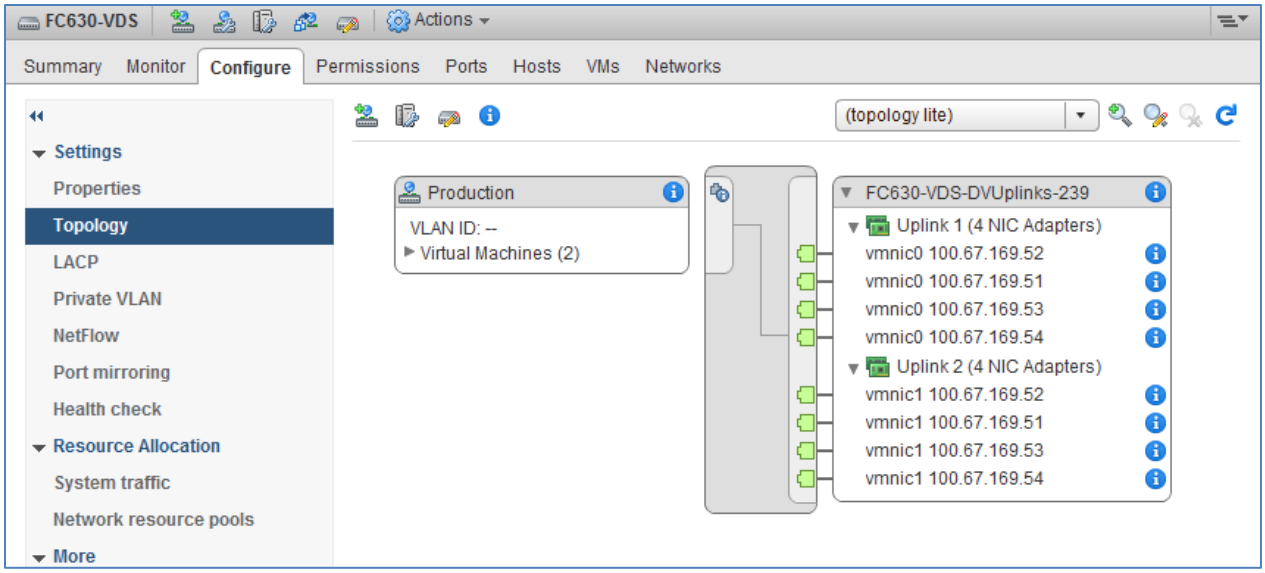

Figure 34 Uplinks configured on FC630-VDS

10.5 Add a virtual network adapter to VMs

Note: Virtual machine/guest operating system deployment steps are not included in this document. For instructions, see the VMware vSphere 6.5 Documentation. Guest operating systems can be any supported by ESXi 6.5. VMs should be deployed before proceeding with this section.

In this section, virtual network adapters (vNICs) are added to VMs for LAN traffic using the Production port group on the VDS.

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. Under the **FX2-FC630-FC** cluster, right click on a VM and click **Edit Settings**.
3. Next to **New Device**, select **Network.** Click **Add.**
4. Click the ▶ icon next to **New Network** to view configuration options.
   a. Next to **New Network**, expand the drop-down menu and select **Show more networks** to open the **Select Network** page.

Note: Be sure to click **Show more networks** to get to the **Select Network** page. If you simply select the VDS that is shown in the drop-down menu, e.g. FC630-VDS, a Port ID error will be displayed.

   b. On the **Select Network** page, select the **Production** port group on **FC630-VDS** and click **OK.**
   c. Make sure the **Connect at Power On** box is checked next to **Status.**
   d. Next to **Adapter Type**, select **VMXNET3** (since 10GbE adapters are used in the hosts).

The **Edit Settings** box now appears as shown in Figure 35. The key network settings are circled.

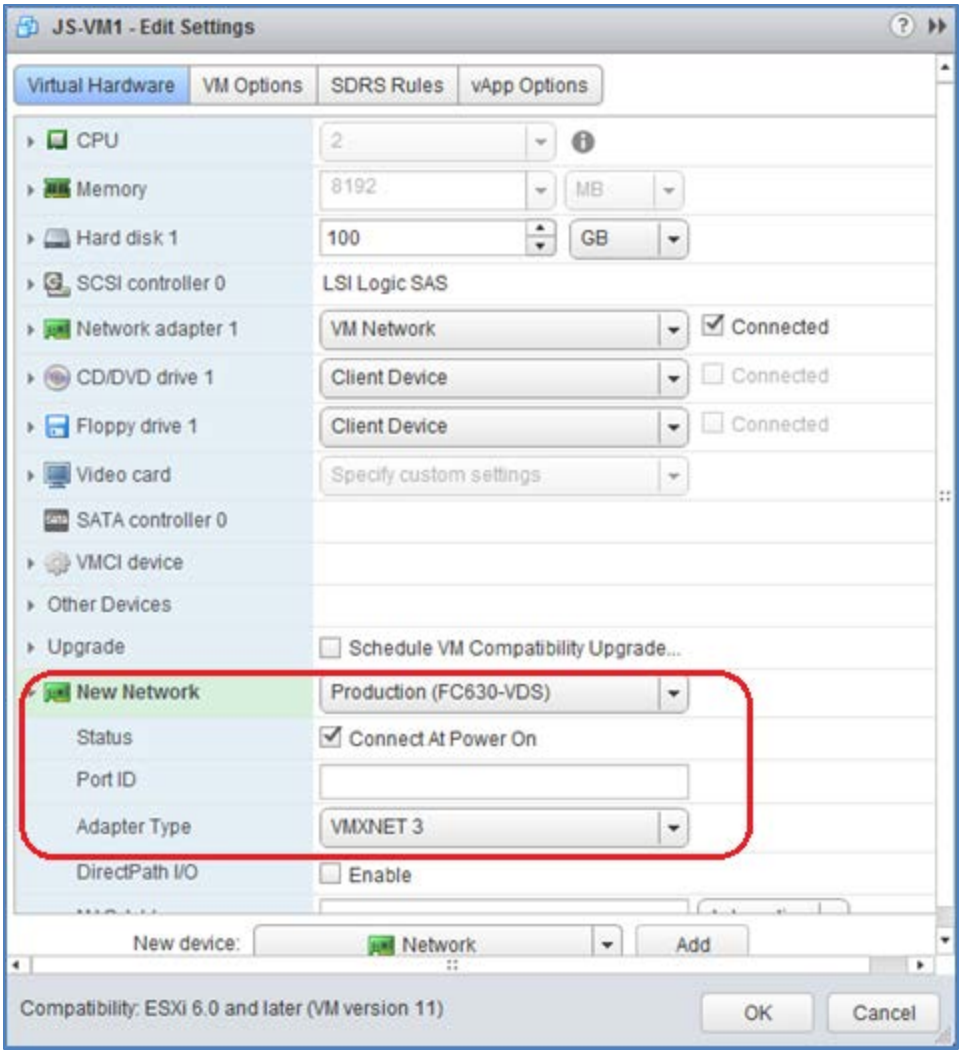

Figure 35 Virtual network adapter configured

5. Click **OK** to add the virtual network adapter.

Repeat steps 2-5 for remaining VMs that will access the production network.

10.6 Verify connectivity to the production network

Log in to a guest OS by right clicking on the VM and selecting **Open Console**. Use the procedure dictated by the guest OS to configure an IP address on the newly added vNIC.

Test connectivity by pinging other VMs in the cluster as well as systems available over the leaf-spine network.

Note: Guest operating system firewalls may need to be temporarily disabled or modified to allow responses to ICMP ping requests.

A FN IOM internal port mapping

A.1 Quarter-width servers with dual-port CNAs

The PowerEdge FC430 is an example of a quarter-width server. For quarter-width servers with dual-port CNAs, each CNA port maps to a single port on each of the two FN IOMs. The server slots in the top row are designated 1a through 1d, and 3a through 3d on the bottom row. See Figure 36 and Table 4.

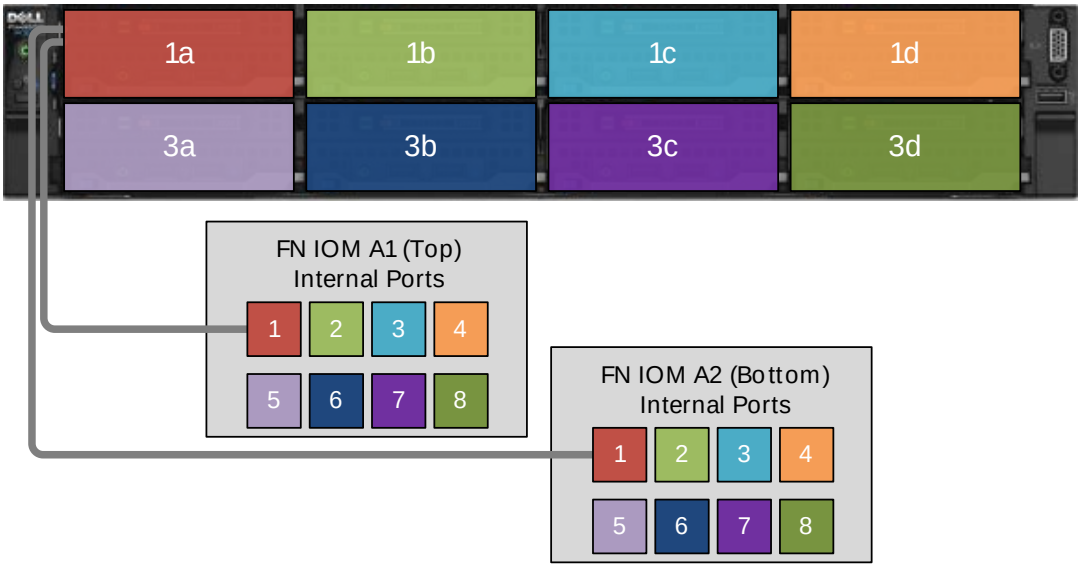

Figure 36 Quarter-width servers with dual-port CNAs

Table 4 Quarter-width servers with dual-port CNAs

Server Slot	FN IOM A1 (Top) Internal Port Numbers	FN IOM A2 (Bottom) Internal Port Numbers
1a	1	1
1b	2	2
1c	3	3
1d	4	4
3a	5	5
3b	6	6
3c	7	7
3d	8	8

Note: Quad-port CNAs are not available for quarter-width servers.

A.2 Half-width servers with dual-port CNAs

The PowerEdge FC630 is an example of a half-width server. For half-width servers with dual-port CNAs installed, the CNA ports map to a single port on each of the two FN IOMs. Figure 37 and Table 5 present the port mapping for half-width servers with dual-port CNAs.

Note: Ports 2, 4, 6 and 8 are not used when using half-width blades with dual-port adapters.

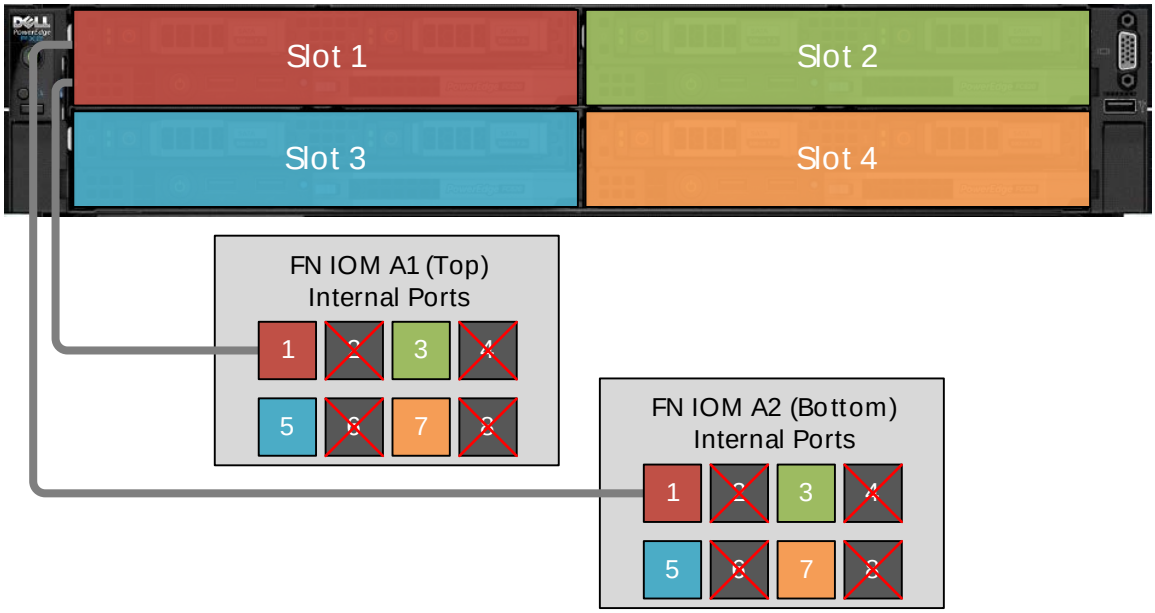

Figure 37 Half-width servers with dual-port CNAs

Table 5 Half-width servers with dual-port CNAs

Server Slot	FN IOM A1 (Top) Internal Port Numbers	FN IOM A2 (Bottom) Internal Port Numbers
1	1	1
2	3	3
3	5	5
4	7	7

A.3 Half-width servers with quad-port CNAs

For half-width servers with quad-port CNAs installed, the CNA ports map to two ports on each FN IOM. Figure 38 and Table 6 present the port mapping for half-width servers with quad-port CNAs.

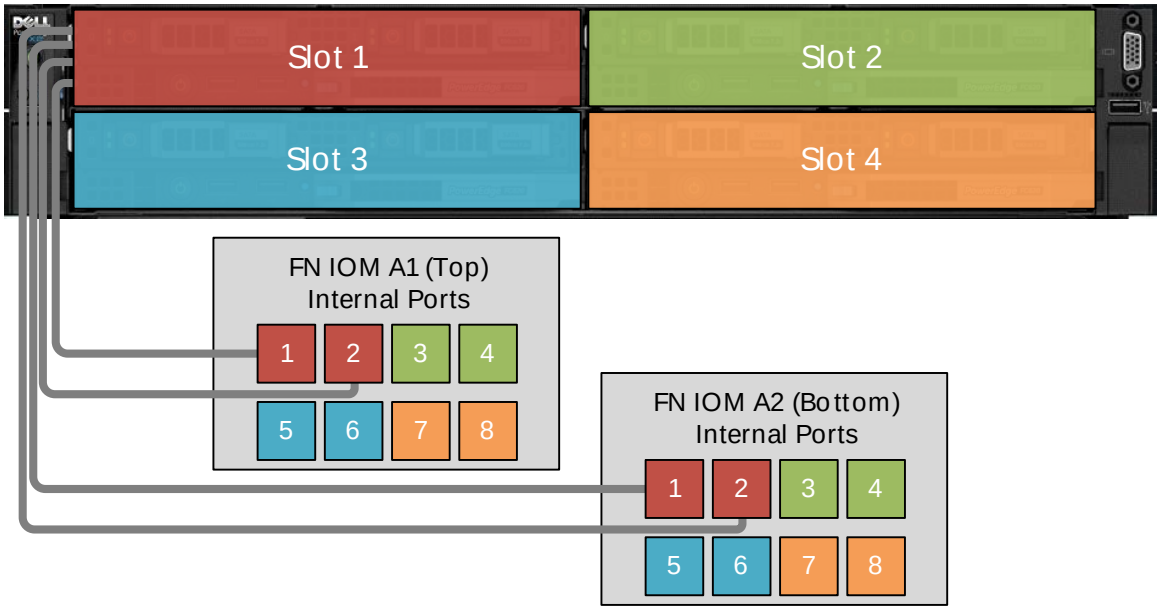

Figure 38 Half-width servers with quad-port CNAs

Table 6 Half-width servers with quad-port CNAs

Server Slot	FN IOM A1 (Top) Internal Port Numbers	FN IOM A2 (Bottom) Internal Port Numbers
1	1,2	1,2
2	3,4	3,4
3	5,6	5,6
4	7,8	7,8

A.4 Full-width servers with dual-port CNAs

The PowerEdge FC830 is an example of a full-width server. For full-width servers with dual-port CNAs installed, the CNA ports map to two ports on each FN IOM. Figure 39 and Table 7 present the port mapping for full-width servers with dual-port CNAs.

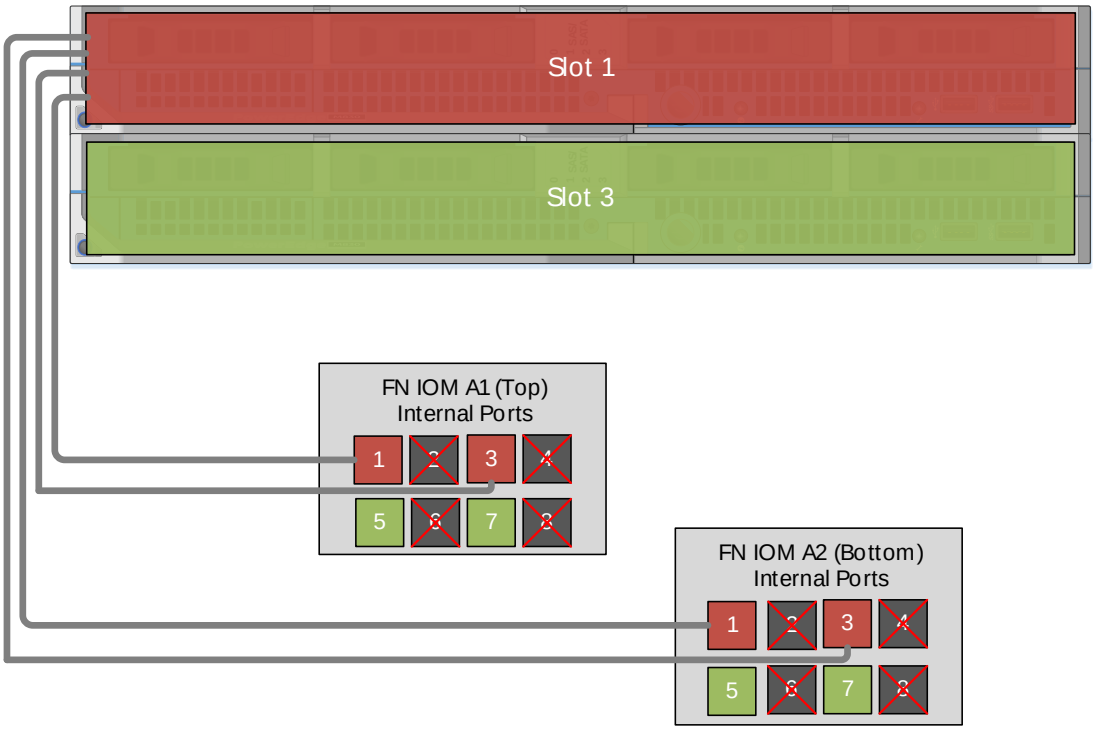

Figure 39 Full-width servers with dual-port CNAs

Table 7 Full-width servers with dual-port CNAs

Server Slot	FN IOM A1 (Top) Internal Port Numbers	FN IOM A2 (Bottom) Internal Port Numbers
1	1 , 3	1 , 3
3	5 , 7	5 , 7

A.5 Full-width servers with quad-port CNAs

For full-width servers with quad-port CNAs installed, the CNA ports map to four ports on each FN IOM. Figure 40 and Table 8 present the port mapping for full-width servers with quad-port CNAs.

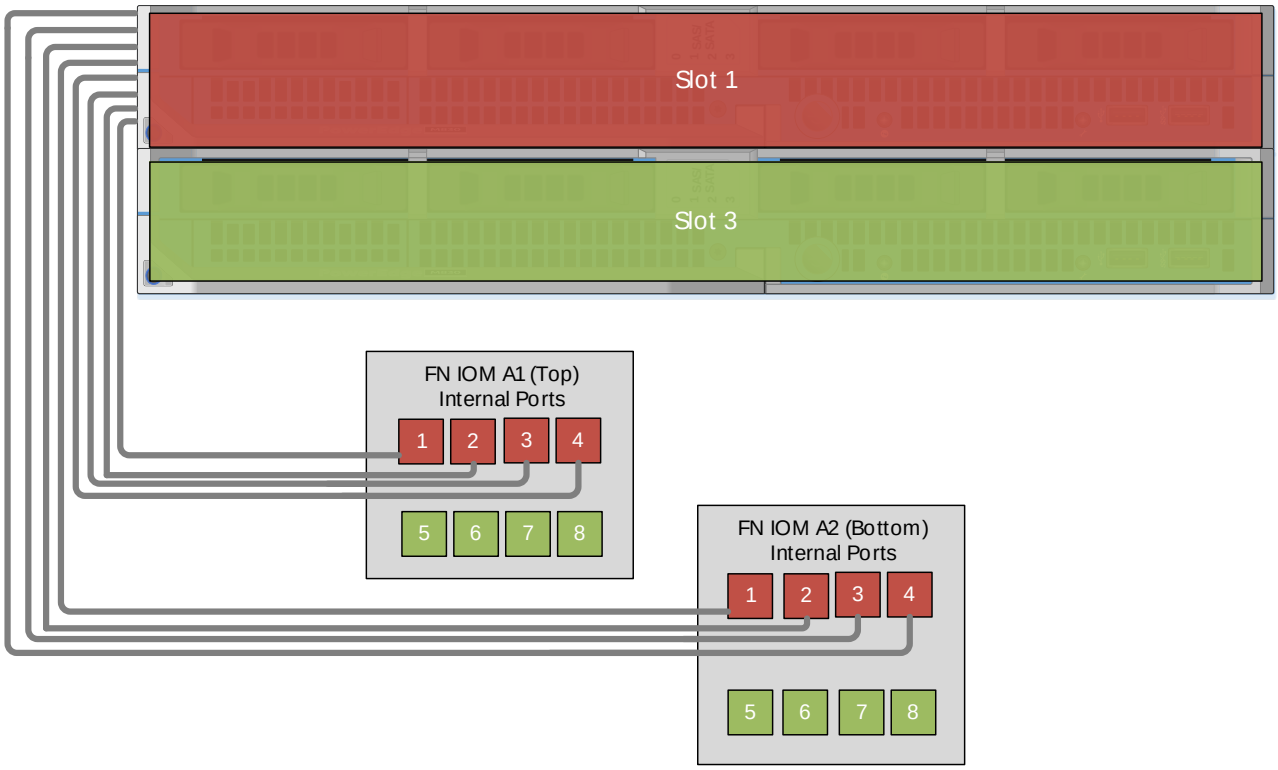

Figure 40 Full-width servers with quad-port CNAs

Table 8 Full-width servers with quad-port CNAs

Server Slot	FN IOM A1 (Top) Internal Port Numbers	FN IOM A2 (Bottom) Internal Port Numbers
1	1 , 2 , 3 , 4	1 , 2 , 3 , 4
3	5 , 6 , 7 , 8	5 , 6 , 7 , 8

B FN IOM operational modes

The FN IOM supports five operational modes: Standalone (SMUX), Virtual Link Trunking (VLT), Programmable MUX (PMUX), Stacking and Full-switch.

To enable a new operational mode, issue the command **`stack-unit 0 iom-mode Mode`** in configuration mode and reload the switch.

Table 9 FN IOM modes and descriptions

IOM mode	Description
Standalone mode (SMUX)	This is the factory default mode for the FN IOM. Standalone mode is a fully-automated, low-touch mode that allows VLAN memberships to be defined on the server-facing ports while all upstream ports are configured in port channel 128. Administrators cannot modify this setting.
VLT mode	This low-touch mode automates all configurations except VLAN membership. Port 9 is dedicated to the VLT interconnect in this mode.
Programmable MUX mode (PMUX)	PMUX mode allows the administrator to create multiple link aggregation groups (LAGs), configure VLANs on uplinks, and configure data center bridging (DCB) parameters on the server-facing ports.
Stacking mode (FN410S and FN410T only)	Stacking mode combines multiple switches to make a single logical switch, which is managed by a designated master unit in the stack.
Full-switch mode	Full-switch mode makes all switch features available. This mode requires the most configuration but allows for the most flexibility and customization.

C Dell EMC validated hardware and components

The following tables include the hardware and components used to configure and validate the example configurations in this guide.

C.1 Switches

Table 10 Switches and firmware versions

Qty	Item	Firmware Version
2	Dell EMC Networking S4048-ON Leaf switches	DNOS 9.11.2.6
1	Dell EMC Networking S3048-ON Management switch	DNOS 9.11.2.6
2	6510 FC switches	FOS 8.1.0a

C.2 PowerEdge FX2s chassis and components

This guide uses one FX2s chassis with four FC630 servers and two FN2210S IOMs.

Table 11 FX2s components and firmware versions

Qty per chassis	Item	Firmware Version
1	FX2s Chassis Management Controller	1.41.200
4	FC630 server blades. Each blade contains: 2 - Intel Xeon E5-2690 v3 2.6GHz CPU, 12 cores 8 - 16 GB DIMMS (128 GB total) 4 – 200 GB SATA SSD 2 - 16 GB Internal SD Cards 1 - PERC H730P Integrated RAID Controller 1 - QLogic 57810 10GbE DP CNA bNDC FC630 BIOS FC630 iDRAC with Lifecycle Controller	 - - - - 25.5.0.0018 08.07.26 2.3.5 2.41.40.40
2	FN2210S IOMs	DNOS 9.11.2.6
4	Intel I350-T Base-T 1GbE DP LP PCIe adapters	17.5.10

C.3 FC Storage Arrays

Storage arrays validated with FN2210S IOMs in NPG mode include the following:

- Dell EMC Unity 500F All-flash storage
- Dell EMC PowerVault MD3820F

This guide uses the Unity 500F storage platform as an example FC storage array. It delivers all-flash storage with up to 8.0 PB raw capacity. It has concurrent support for NAS, iSCSI, and FC protocols. The Disk Processing Enclosure (DPE) has a 2-RU form factor, has 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.

Each of the SPs in the Disk Processing Enclosure (DPE) has 2 on-board 12 Gbps SAS ports for connecting to Disk Access Enclosures (DAEs). Additionally, a 4-port 12 Gbps SAS I/O Module can be provisioned in order to provide additional back-end buses.

Figure 41 Dell EMC Unity 500F

D Validated software and required licenses

D.1 VMware software

The following table lists the VMware software components used to validate the examples in this guide.

Table 12 Software versions

Item	Version
VMware ESXi	6.5 U1 - Dell EMC customized image version A00 build 5969303
VMware vCenter Server Appliance	6.5 U1 – build 5973321 (Includes vSphere Web Client)

D.2 VMware licenses

The vCenter Server is licensed by instance. Other VMware product licenses are allocated based on the number of CPU sockets in the participating hosts.

Required licenses for the topology built in this guide are as follows:

- VMware vSphere 6 Enterprise Plus – 8 CPU sockets
- vCenter 6 Server Standard – 1 instance

E Technical support and resources

Dell TechCenter is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware and services.

E.1 Dell EMC product manuals and technical guides

Manuals and documentation for Dell EMC Networking S3048-ON

Manuals and documentation for Dell EMC Networking S4048-ON

Manuals and documentation for Dell EMC Networking S6010-ON

Manuals and documentation for Dell EMC FN IOM

Manuals and documentation for Dell EMC PowerEdge FX2/FX2s

DELL EMC Unity Technical Documentation

Dell EMC TechCenter Networking Guides

Dell EMC PowerEdge FN I/O Module (FN IOM) – Easy Deployment Guide

Dell EMC Leaf-Spine Deployment and Best Practices Guide

E.2 VMware product manuals and technical guides

VMware vSphere Documentation

vSphere Installation and Setup – This document includes ESXi 6.5 and vCenter Server 6.5.

VMware Compatibility Guide

VMware vSphere 6.5 Documentation - Configuring FCoE

E.3 6510 Fabric OS guides

The following guides are available at http://my.Brocade.com for Brocade 6510 or https://support.emc.com for Connectrix DS-6510B switches:

Brocade Fabric OS Administration Guide, 8.1.0

Brocade Fabric OS Command Reference, 8.1.0

Brocade Fabric OS Web Tools Administration Guide 8.1.0

F Support and Feedback

Contacting Technical Support

Support Contact Information	Web: http://support.dell.com/
	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.