

FX2 Storage Networking with iSCSI

Dell EMC Networking Solutions Engineering
May 2017

Revisions

Date	Revision	Description	Authors
May 2017	1.0	Initial Release	Shree Rathinasamy, Jordan Wilson

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND. Copyright © 2017 Dell Inc. or its subsidiaries. All rights reserved. Dell and the Dell EMC logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of contents

Revisions.....	2
1 Introduction.....	6
2 Hardware overview.....	8
2.1 Dell EMC PowerEdge FX2s enclosure and supported modules.....	8
2.1.1 PowerEdge FC630 server.....	9
2.1.2 PowerEdge FN410S I/O Module.....	9
2.2 Dell EMC Networking S4048-ON.....	10
2.3 Dell EMC Networking S6010-ON.....	10
2.4 Dell EMC Networking S3048-ON.....	10
3 iSCSI and example environment overview.....	11
3.1 Storage topology.....	11
3.2 Ethernet topology.....	13
3.3 LAN and Storage topology.....	13
3.3.1 Traditional topology example.....	14
3.3.2 Converged topology example.....	16
3.4 Management network.....	16
4 Server and VMware environment preparation.....	18
4.1 Servers.....	18
4.1.1 Confirm network adapters are at factory default settings.....	18
4.1.2 Enable iSCSI offloading on QLogic 57810 adapters.....	18
4.2 VMware ESXi 6.5.....	19
4.3 VMware vCenter 6.5.....	19
4.4 vCenter Host and Cluster settings.....	19
4.5 VLAN/SAN.....	19
5 Traditional deployment configurations for physical switches.....	21
5.1 Factory default settings.....	21
5.2 Dell EMC FN410S IOM Configuration.....	22
5.2.1 FN410S switch Full-switch mode configuration.....	22
5.3 S4048-ON iSCSI SAN switch configuration.....	25
5.4 Leaf switch configuration.....	27
5.4.1 Leaf S4048-ON switch configuration.....	27
5.5 S3048-ON management switch configuration.....	30

- 6 Physical switch configuration for converged deployments.....31
 - 6.1 FN410S configuration31
 - 6.1.1 FN410S in Standalone mode31
 - 6.1.2 FN410S in Full-switch mode.....32
 - 6.2 S4048-ON Leaf switch configuration33
 - 6.3 S6010-ON Spine switch configuration.....36
 - 6.4 iSCSI SAN S4048-ON switch configuration38
- 7 Virtual Network configuration41
 - 7.1 Information on vSphere Standard Switch.....41
 - 7.2 Information on vSphere Distributed Switch42
 - 7.3 Create an iSCSI VDS for the FX2-FC630-iSCSI cluster42
 - 7.4 Set the iSCSI VDS MTU to 9000 and enable LLDP.....43
 - 7.5 Add distributed port groups44
 - 7.6 Configure teaming and failover.....47
 - 7.7 Associate hosts and assign uplinks.....49
 - 7.8 Add VMkernel adapters for iSCSI.....49
 - 7.9 Increase the MTU to 9000 on iSCSI VMkernel adapters52
 - 7.10 Bind iSCSI adapters with VMkernel ports52
 - 7.11 Configure dynamic discovery53
- 8 Configure iSCSI Storage55
 - 8.1 Create a Storage Pool56
 - 8.2 Add vCenter to Unisphere57
 - 8.3 Create a LUN.....58
 - 8.4 Configure iSCSI Interfaces59
 - 8.5 Rescan storage on hosts.....60
 - 8.6 Create a datastore61
 - 8.7 CHAP Authentication.....63
- A FN IOM Internal Port Mapping64
 - A.1 Quarter-Width Servers - Dual Port CNAs.....64
 - A.2 Half-Width Servers - Dual Port CNAs.....65
 - A.3 Half-Width Servers - Quad Port CNAs65
 - A.4 Full-Width Servers - Dual Port CNAs66
 - A.5 Full-Width Servers - Quad Port CNAs.....67

- B FN Series Operational Modes69
- C Dell EMC validated hardware and components70
 - C.1 Switches70
 - C.2 PowerEdge FX2s chassis and components70
 - C.3 Validated iSCSI Storage arrays71
- D Dell EMC validated software and required licenses72
 - D.1 Software72
 - D.2 Licenses72
- E Technical support and resources73
 - E.1 Dell EMC product manuals and technical guides73
 - E.2 VMware product manuals and technical guides73
- F Support and Feedback74

1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company in the data center - not only for 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 also to 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 address evolving workloads precisely. 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, dedicated hosting organizations, and enterprises addressing growth with easily scalable platforms. 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: The Dell EMC Networking FN-series of switches runs internally to the FX2 chassis. An integrated networking solution, the FN-series provides Ethernet, as well as LAN/SAN convergence with iSCSI and FCoE support. The S-series switches are multi-layer Ethernet switches that scale beyond multiple chassis. The switches run on Dell EMC's own operating system or the operating system from one of Dell EMC's 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 iSCSI 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. These instructions for deploying iSCSI using Dell EMC hardware and software target a network administrator or engineer. The instructions assume that the administrator or engineer has traditional networking and VMware ESXi experience.

1.1 Typographical conventions

This document uses the following typographical conventions:

Monospace text	Command Line Interface (CLI) examples
Bold monospace text	Commands entered at the CLI prompt
<i>Italic monospace text</i>	Variables in CLI examples

2 Hardware overview

This section briefly describes the primary hardware used to validate the deployment of iSCSI. [Dell EMC validated hardware and components](#) section provides a complete listing of hardware validated for this guide.

2.1 Dell EMC PowerEdge FX2s enclosure and supported modules

The Dell EMC FX Architecture is a great way to optimize workloads, maximize efficiency, and simplify complexity in today's data center.

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 storage sleds. The FX2s enclosure used for the example 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 one Chassis Management Controller (CMC), two FN IO Modules (FN IOMs), eight Peripheral Component Interconnect Express (PCIe) expansion slots, and redundant power supplies as shown in Figure 2:



Figure 2 Dell EMC PowerEdge FX2s (back)

Note: The Dell EMC PowerEdge FX enclosure is currently available in two models: FX2 and FX2s. The FX2 enclosure is similar to the FX2s but does not support storage sleds and PCIe slots.

2.1.1 PowerEdge FC630 server

The PowerEdge FC630 server is a half-width, 2-socket server. Four FC630 servers (see Figure 3) in the FX2s enclosure form the compute cluster for the example deployment in this guide.



Figure 3 PowerEdge FC630

2.1.2 PowerEdge FN410S I/O Module

PowerEdge FN IOMs are network switches housed in the back of the FX2s enclosure. Dell EMC offers three FN IOM options. Each one provides 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 Fibre Channel (FC) traffic. It uses 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 provide eight 10GbE internal, server-facing ports and four external ports. The FX2s enclosure for the example deployment in this guide uses two PowerEdge FN410S IOMs.



FN410S

4-port SFP+ I/O Module

Provides four SFP+ 10GbE ports. Supports optical and Direct Attach Copper (DAC) cable media.



FN410T

4-port 10GBASE-T I/O Module

Provides four 10GBASE T ports. Supports cost-effective, copper media up to 100 meters.



FN2210S

4-port Combo FC/Ethernet I/O Module

Provides four ports. Up to two ports can be configured for 2, 4, or 8 Gbit/s FC. The remaining ports support SFP+ 20GbE to provide Ethernet connectivity. Supports optical and DAC cable media.

Figure 4 PowerEdge FN410S

2.2 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. This deployment uses four S4048-ON switches.



Figure 5 Dell EMC Networking S4048-ON

2.3 Dell EMC Networking S6010-ON

The S6010-ON is a 1-RU, layer 2/3 switch with thirty-two 40GbE QSFP+ ports. The example deployment in this guide uses one S6010-ON as a spine/core switch.

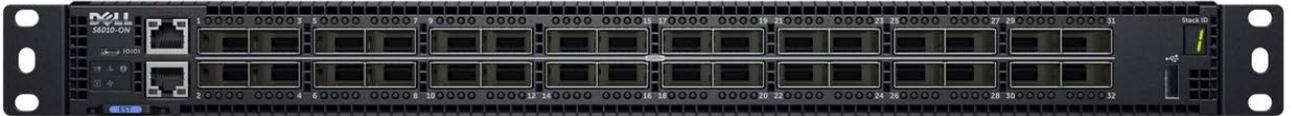


Figure 6 Dell EMC Networking S6010-ON

2.4 Dell EMC Networking S3048-ON

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

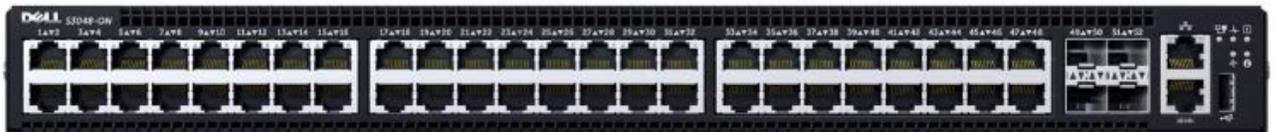


Figure 7 Dell EMC Networking S3048-ON

3 iSCSI and example environment overview

In a traditional data center environment, we see a defined storage network as well as a defined LAN. Choosing iSCSI as the storage protocol allows for a more cost-effective solution for the implementation of a storage network. The use of iSCSI also gives the ability to access a SAN without dedicated networking equipment. In the example in Figure 8, we use the FN410S to access both networks simultaneously. Each FC630 server includes a dual-port converged network adapter (CNA).

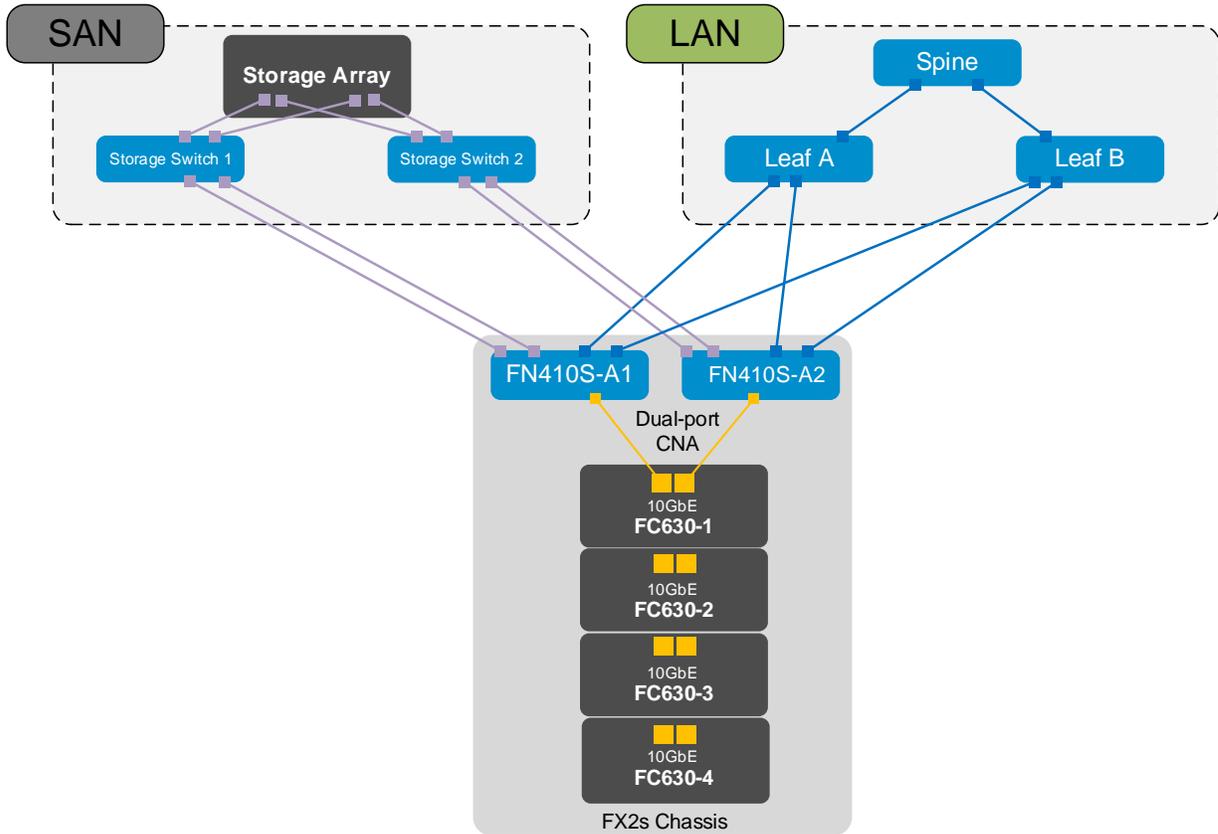


Figure 8 LAN and SAN using iSCSI

3.1 Storage topology

The storage solution is simplistic in design. As Figure 9 shows, the FC630 server uses iSCSI to transmit storage frames to the FN410S. To validate the solution, S4048-ON switches were used as the upstream switches and a storage storage array was used as the storage target.

This deployment uses four FC630 servers running VMware ESXi 6.5 as compute nodes. All the servers use QLogic BCM57810 dual-Port CNA.

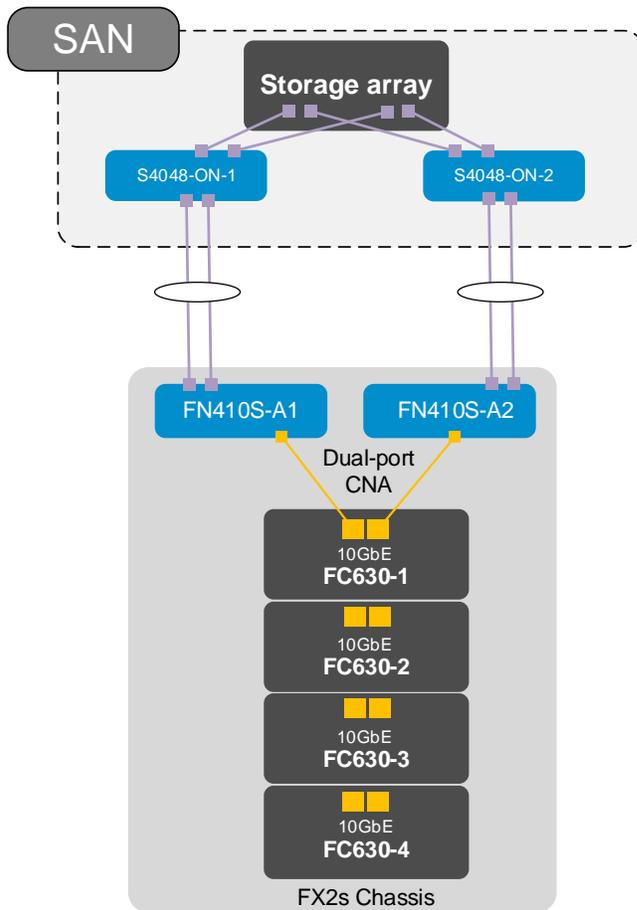


Figure 9 iSCSI storage topology showing only one FC630 connection

3.2 Ethernet topology

Figure 10 illustrates the topology used to connect an FC630 server to an Ethernet LAN using FN410S. Leaf switch configuration in the figure creates a Virtual Link Trunking (VLT) pair. This allows the port channel from each FN410S to span the two leaf switches. To validate the solution, S6010-ON switch was used as a Spine switch and S4048-ON switches were used as leaf switches. Figure 10 shows the Ethernet topology.

Note: The use of a leaf-spine network in the datacenter as shown in Figure 10 is considered a best practice. See [Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage](#) for a detailed description and configuration instructions for a leaf-spine network.

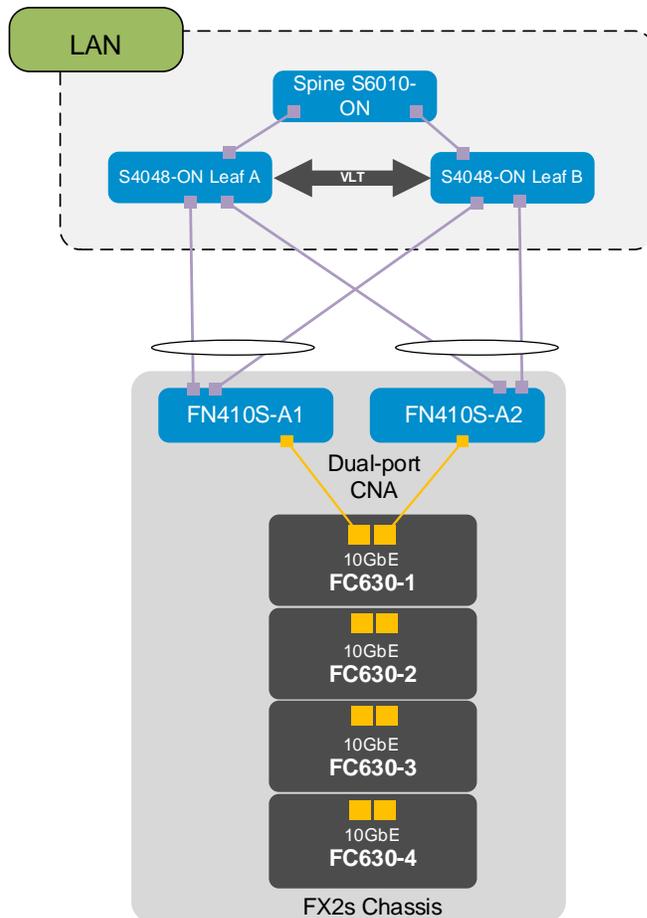


Figure 10 Ethernet topology

3.3 LAN and Storage topology

This deployment guide illustrates two iSCSI implementations. One example uses the traditional, non-converged topology and the other uses a converged environment.

3.3.1 Traditional topology example

Figure 11, shows a non-converged, traditional topology with segregation of the LAN and SAN.

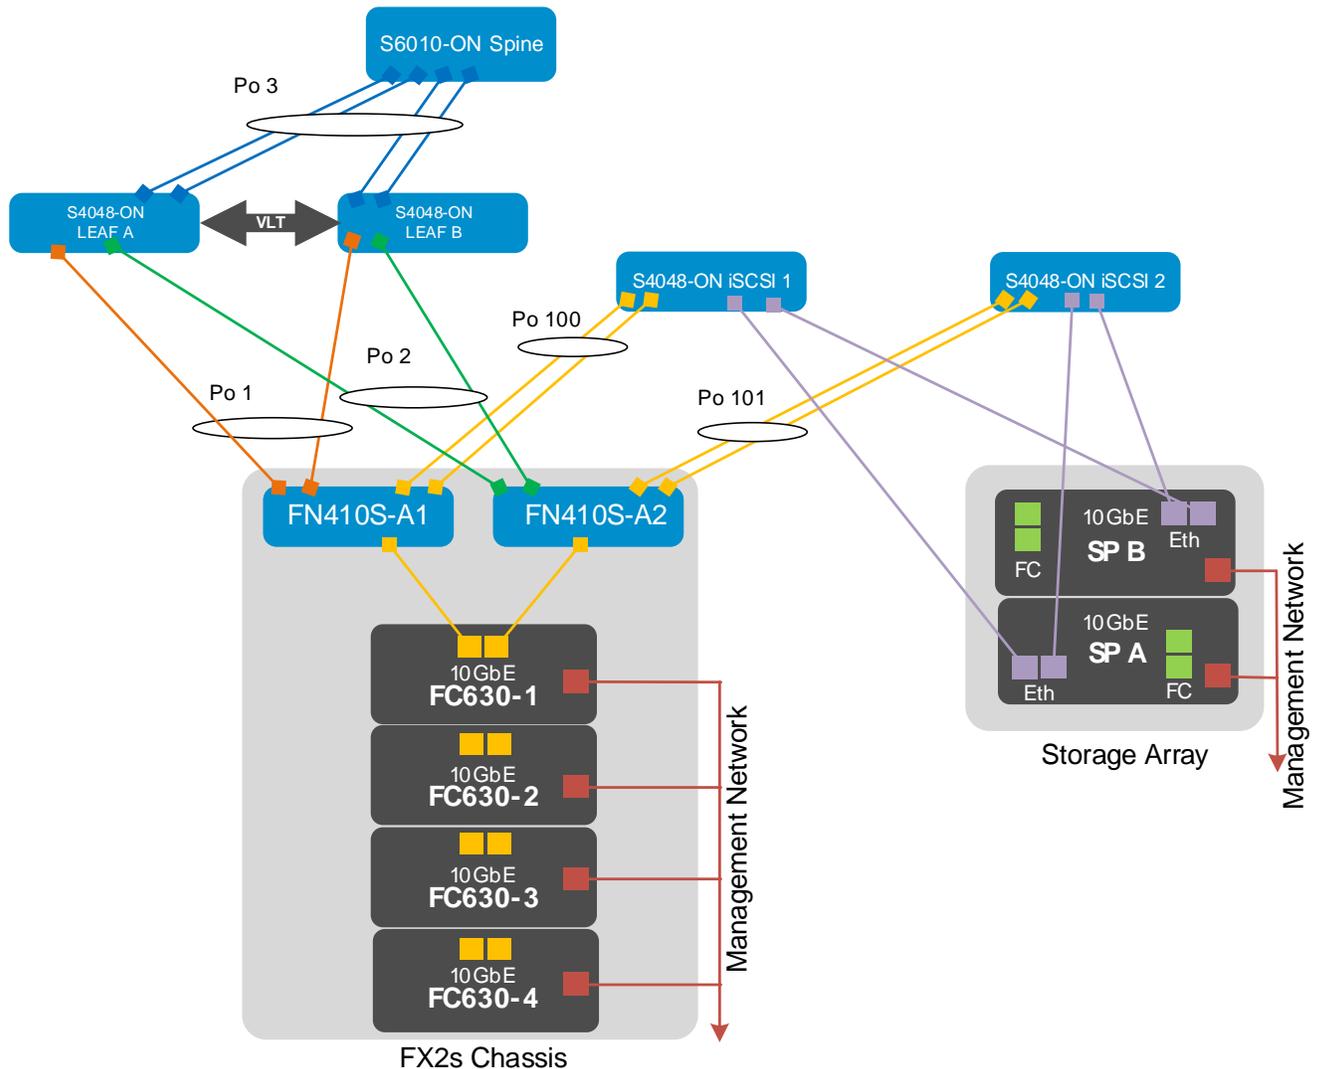


Figure 11 Topology for traditional, non-converged iSCSI example

In this example, port channels connect each FN410S to S4048-ON switches that behave as leaf switches. The leaf switches use VLT for increased bandwidth and resilience. Port channels from each leaf switch connect to an S6010-ON switch acting as a Spine switch. Each of these port channel consists of two members with 40GbE capacity each. The comprehensive configuration and validation of the spine switch, including routing, is beyond the scope of this document. This document captures Layer 2 configuration of the Spine switch for the converged topology example.

Note: This topology represents the port channel in Full-switch mode. In Standalone mode, the port channels for external ports are 128 by default.

Dual port CNAs on the FC630 servers are internally mapped to FN410S IOM in slots A1 and A2. The port channel from ports Te 0/9 and Te 0/10 connects to S4048-ON storage switches. Ports Te 1/1 and Te 1/2 from each S4048-ON connect to storage processors A and B on the storage array. Table 1 below lists the mapping of the port from S4048-ON-1 and S4048-ON-2:

Table 1 iSCSI VLANs and networks

Port number	S4048-ON - 1	S4048-ON - 2
Port 1	To SP A Ethernet Port 0	To SP A Ethernet Port 1
Port 2	To SP B Ethernet Port 0	To SP B Ethernet Port 1
Port 9	To IOM A1 Port 9	To IOM A2 Port 9
Port 10	To IOM A1 Port 10	To IOM A2 Port 10

Since there are two VLANs used for iSCSI, please see [VLAN/SAN Assignment Planning](#) for further information on VLAN planning. Zero-touch Standalone mode automatically applies the same iSCSI VLAN ID. This happens by default and requires no configuration. The process works identically for the dual-port and quad-port configurations.

3.3.2 Converged topology example

The second example for deploying iSCSI involves using a converged topology as shown in Figure 12. This topology removes the segregation of the LAN and SAN, allowing for a much more manageable and scalable solution.

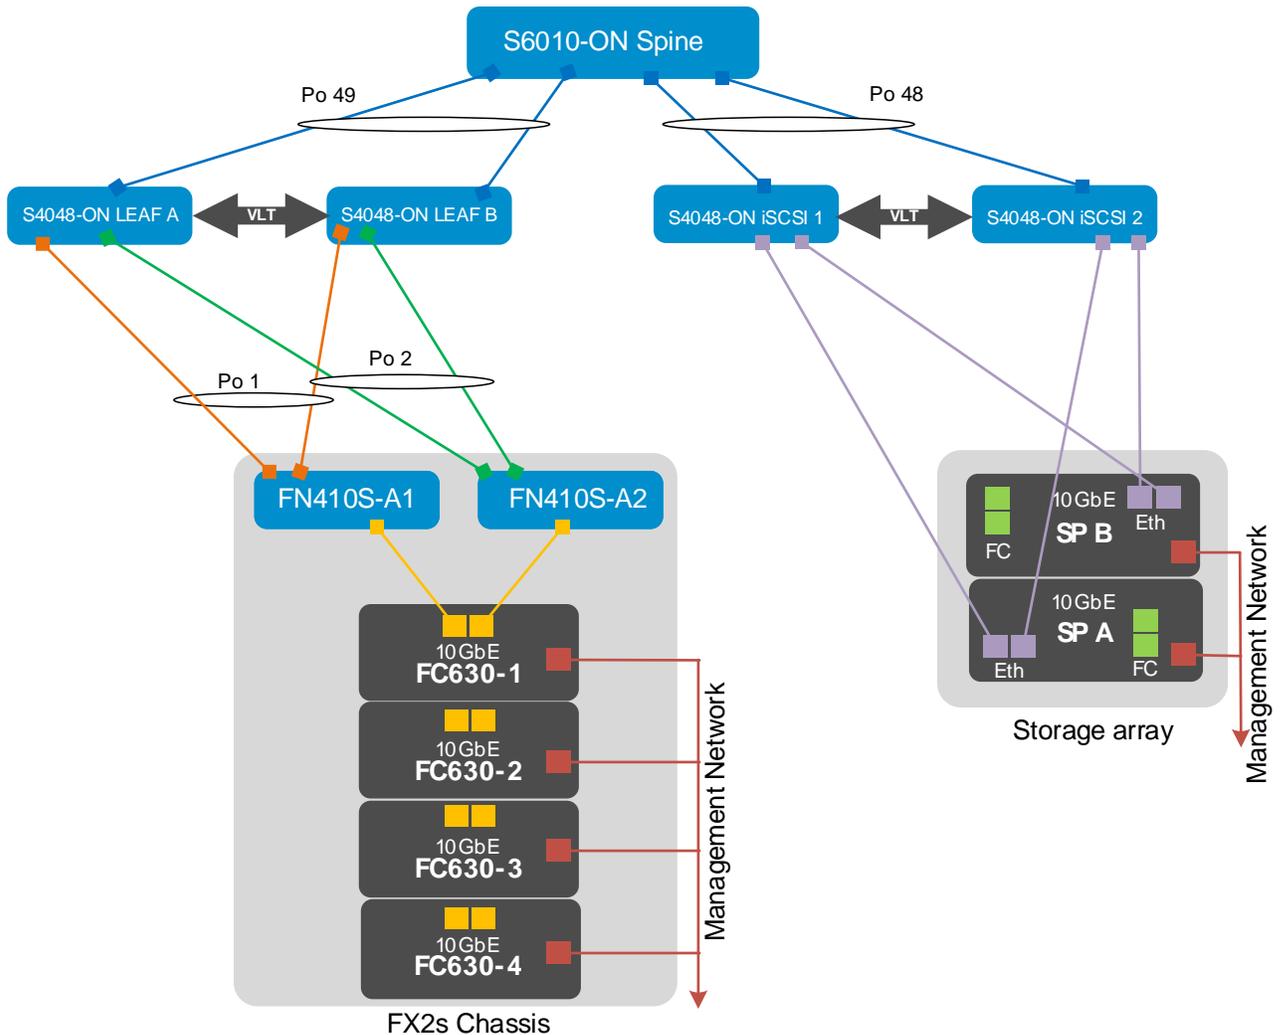


Figure 12 Topology for converged iSCSI example

3.4 Management network

This guide uses a single management traffic network that is isolated from the LAN and SAN. It is consistent in both traditional and converged examples. An S3048-ON switch installed in each rack provides connectivity to the management network. The following components connect to the S3048-ON management network:

- Four S4048-ON switches
- One S6010-ON switch
- One FX2s CMC
- Four FC630 servers connected using PCIe slots 2, 4, 6, and 8
- Two Storage Processors (A and B)

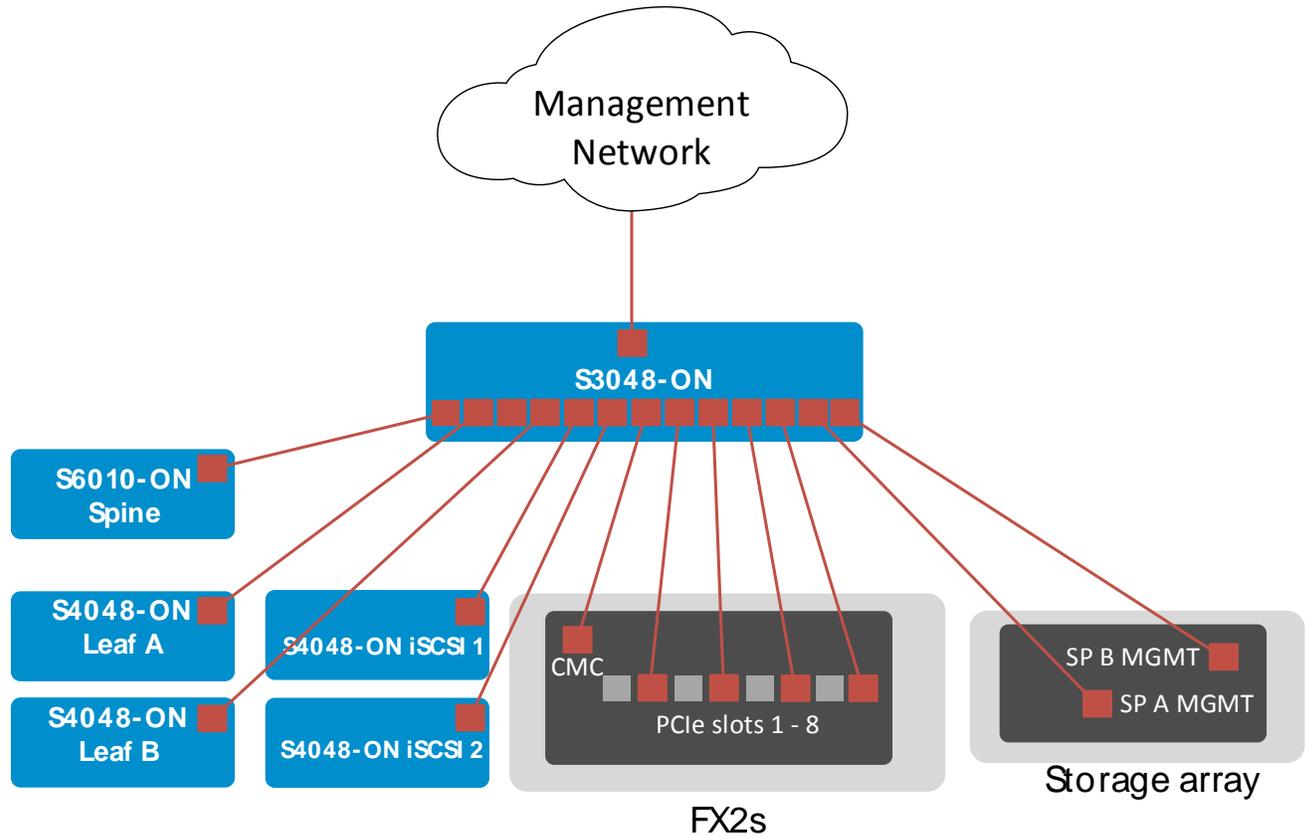


Figure 13 Management network

4 Server and VMware environment preparation

4.1 Servers

This section covers basic PowerEdge server preparation. Installation of guest operating systems (Microsoft Windows Server, Red Hat Linux, and so on) is outside the scope of this document.

Note: Exact iDRAC console 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 virtual console.

4.1.1 Confirm network adapters are at factory default settings

Note: These steps are only necessary if installed network adapters have been modified from their factory default settings. For more information on internal port mapping, see [Appendix A](#)

1. Connect to the 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 NIC in the list.
6. From the **Main Configuration Page**, click **Default** followed by **Yes** to load the default settings. Click **OK**.
7. To save changes to the settings, click **Finish** then **Yes**. Click **OK**.
8. Repeat for each NIC and port listed on the **Device Settings** page.

4.1.2 Enable iSCSI offloading on QLogic 57810 adapters

For each FC630 in the cluster:

1. From the System Setup Main Menu, select **Device Settings**.
On the **Device Settings** page, click the first port of the QLogic 57810 CNA to be used for iSCSI connectivity. This opens the **Main Configuration Page** for the port.
Select **Device Level Configuration** and change the **Virtualization Mode** to **NPar**. Click **Back**.
Select **NIC Partitioning Configuration**.
Select **Partition 1 Configuration**. Set **NIC Mode** and **iSCSI Offload Mode** to **Enabled**. Leave **FCoE Mode** set to **Disabled**. Click **Back**.
Select **Partition 2 Configuration** and set all three modes (NIC, iSCSI, FCoE) to **Disabled**. Click **Back**. Repeat on partitions 3 and 4 to disable modes.
After configuring all partitions, click **Back > Finish**. Answer **Yes** when prompted to save changes.
Click **OK** to return to the **Device Settings** page.
Click the second port of the QLogic 57810 CNA to be used for iSCSI connectivity and repeat step 5 above to configure partitioning.
Click **Finish > Finish** and answer the confirmation prompts as needed to save all changes and reboot the system.

4.2 VMware ESXi 6.5

Install VMware ESXi 6.5 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.

4.3 VMware vCenter 6.5

This guide does not cover the installation and initial configuration of VMware's vCenter 6.5. For information on the installation and configuration of vCenter 6.5 see the [VMware vSphere 6.5 Documentation Center](#).

4.4 vCenter Host and Cluster settings

The vSphere Web Client is a service running on vCenter Server. In the vSphere Web Client, a datacenter object named Datacenter is created for this deployment and the ESXi hosts are added to it. Create a cluster named FX2-FC630-ISCSI and add it to the datacenter object. Add the four FC630 hosts to the FX2-FC630-ISCSI cluster. When complete, the vSphere Web Client Navigator pane appears as shown in Figure 14.

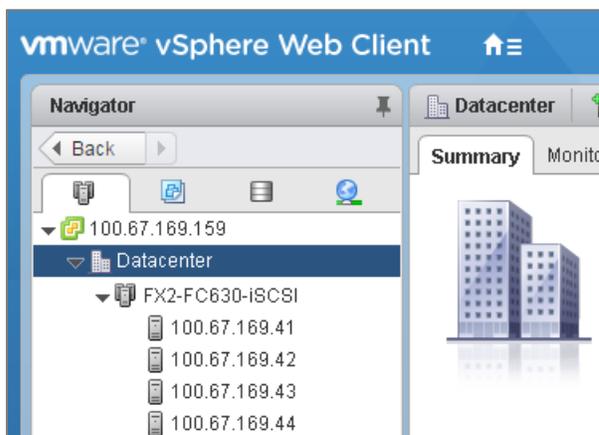


Figure 14 Datacenter and Cluster

4.5 VLAN/SAN

To ensure consistency across all switches, it is a best practice to plan all the required VLANs for the network ahead of time. This deployment uses two VLANs, one for each iSCSI SAN. Table 2 lists these VLANs and the purpose of each. For additional information on VLANs and their use inside of a VMware-enabled environment, see [VMware vSphere 6.5 Documentation Center](#). This deployment uses the following addressing scheme for the iSCSI networks to make the deployment faster and more efficient. It also helps with troubleshooting later.

Table 2 iSCSI VLANs and associated networks in CIDR notation

VLAN ID	Network	Used For
100	192.168.100.0/24	iSCSI-1
101	192.168.101.0/24	iSCSI-2

Table 3 Host iSCSI IP addresses

Host	iSCSI-1	iSCSI-2
FC630-1	192.168.100.41	192.168.101.41
FC630-2	192.168.100.42	192.168.101.42
FC630-3	192.168.100.43	192.168.101.43
FC630-4	192.168.100.44	192.168.101.44

Note: FN410S IOM configuration needs only the iSCSI VLANs in Full Switch mode. FN410S in standalone mode automatically selects the iSCSI VLAN, eliminating the need for configuration. The storage network S4048-ON requires VLAN configuration regardless of the IOM mode used.

5 Traditional deployment configurations for physical switches

This section contains example switch configuration for non-converged, traditional deployment with explanations for one switch in each major role on the production network. See Figure 11. This section details the following switches:

- FN410S-A1
- S4048-ON: iSCSI 1
- S4048-ON: Leaf A
- S3048-ON: Management

The remaining switches use configurations very similar to the configuration detailed in this section.

Notes: The attachments contain individual switch configurations for the switches described by the examples in this guide.

5.1 Factory default settings

The configuration commands in the sections below assume switches are at their factory default settings. As per your need, use the following procedure to reset all switches in this guide to factory defaults:

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

These commands restore factory settings and reload the switch. After reload, enter **A** at the [A/C/L/S] prompt as shown below to exit Bare Metal Provisioning (BMP) 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: BMP mode does not appear for IOMs after restoring factory settings and reloading the switch.

5.2 Dell EMC FN410S IOM Configuration

This section covers FN410S in Full-switch mode and provides the configuration used in this example. This modes allow for more granular control over the FN410S and allow the IOM to behave more like a traditional switch.

Note: The topology this example uses is feasible for Full-switch mode. Hence, Standalone mode is not used for the example in non-converged topology.

5.2.1 FN410S switch Full-switch mode configuration

The default mode is Standalone mode. In order to change FN410S switches from Standalone mode to Full-switch mode, use the following command. Note that a reload is required to ensure that the change takes effect.

```
Dell(conf)#stack-unit 0 iom-mode full-switch
% You are about to configure the Full Switch Mode.
Please reload to effect the changes

Dell(conf)#exit
Dell#Feb 18 04:16:29: %STKUNIT0-M:CP %SYS-5-CONFIG_I: Configured from
console
reload
System configuration has been modified. Save? [yes/no]: y
!
Feb 18 04:16:33: %STKUNIT0-M:CP %FILEMGR-5-FILESAVED: Copied running-config
to startup-config in flash by default

Proceed with reload [confirm yes/no]: y
```

Initial configuration involves setting the hostname and enabling Link Layer Discovery Protocol (LLDP), which is useful for troubleshooting. Finally, configure the management interface and default gateway.

```
enable
configure

hostname FN410S-A1

protocol lldp
  advertise management-tlv management-address system-name
  no advertise dcbx-tlv ets-reco

interface ManagementEthernet 0/0
ip address 100.67.169.145/24
no shutdown

management route 0.0.0.0/0 100.67.169.254
```

This deployment example uses a non-DCB environment. Hence, DCB must be manually disabled using the command:

```
no dcb enable
```

Configure a password to enable serial console. Change the default SSH/telnet password to root. Reset the password on each of the first two lines as desired. Disable telnet.

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**

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

Continue configuring the FN410S by changing the internal interfaces to portmode hybrid and switchport.

Notes:1. MTU - Dell EMC recommends setting the MTU to 9216 for best performance on switches used in iSCSI SANs.
2. Port channel numbering - LACP port channel numbers can be any number from 1-128.

```
interface range TenGigabitEthernet 0/1-8
no ip address
mtu 9216
portmode hybrid
switchport
no shutdown
```

The example adds Te0/9 and Te 0/10 to Link Aggregation Control Protocol (LACP)-enabled port channel 100. Similarly, this example uses LACP to add Te 0/9 and Te 0/10 to port channel 101 in FN410S slot A2.

```
interface TenGigabitEthernet 0/9
description Connection to S4048-ON iSCSI
no ip address
mtu 9216
port-channel-protocol LACP
port-channel 100 mode active
no shutdown
```

```
interface TenGigabitEthernet 0/10
description Connection to S4048-ON iSCSI
no ip address
mtu 9216
port-channel-protocol LACP
port-channel 100 mode active
no shutdown
```

```

interface TenGigabitEthernet 0/11
  description Connection to S4048-ON LEAF
  no ip address
  port-channel-protocol LACP
    port-channel 1 mode active
  no shutdown

interface TenGigabitEthernet 0/12
  description Connection to S4048-ON LEAF
  no ip address
  port-channel-protocol LACP
    port-channel 1 mode active
  no shutdown

interface Port-channel 1
  description Port Channel for LAN traffic
  no ip address
  portmode hybrid
  switchport
  no shutdown

interface Port-channel 100
  description Port Channel for iSCSI 1
  no ip address
  mtu 9216
  portmode hybrid
  switchport
  no shutdown

```

The example uses two VLAN interfaces, VLAN 100 and VLAN 101, for iSCSI. Since FN410S is in Full-switch mode, the example requires tagging ports Te 1/1-8 to VLAN 100 in FN410S slot A1 and VLAN 101 in FN410S in slot A2. Similarly, the example tags Port Channel 100 to VLAN 100 in the FN410S in slot A1 and Port Channel 101 to VLAN 101 in FN410S slot A2. Although the example configures all internal interfaces for switchport, portmode hybrid and jumbo frames, only internal interfaces that are mapped to the dual port adapters require switchport, portmode hybrid and jumbo frames capability. See [FN IOM Internal Port Mapping](#) for internal port mappings.

```

interface Vlan 100
  no ip address
  mtu 9216
  tagged TenGigabitEthernet 0/1-8
  tagged Port-channel 100
  no shutdown

```

Save the configuration.

```
end
write
```

5.3 S4048-ON iSCSI SAN switch configuration

The deployment example in this guide uses two iSCSI SAN switches. The following configuration details are specific to switch S4048-ON iSCSI-1. Switch S4048-ON iSCSI-2's configuration is similar. The switches start at their factory default settings. See [Factory default settings](#). Initial configuration involves setting the hostname and enabling LLDP, which is useful for troubleshooting. Configure the management interface, default gateway and serial-console enable password.

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_passwordenable

```
enable
configure

hostname iSCSI-1

protocol lldp
  advertise management-tlv management-address system-description system-name
  advertise interface-port-desc

interface ManagementEthernet 1/1
ip address 100.67.169.35/24
no shutdown

management route 0.0.0.0/0 100.67.169.254

enable sha256-password <enable_password>

no ip telnet server enable
```

Use the next set of commands to configure the two downstream interfaces (connected to the FN410S). Issue the interfaces portmode hybrid and switchport commands and set the MTU to 9216 for performance. Add these two interfaces to port channel 100.

```
interface TenGigabitEthernet 1/9
  description FN410S A1
  no ip address
  mtu 9216
  port-channel-protocol LACP
  port-channel 100 mode active
```

```
no shutdown

interface TenGigabitEthernet 1/10
description FN410S A1
no ip address
mtu 9216
port-channel-protocol LACP
port-channel 100 mode active
no shutdown

interface Port-channel 100
no ip address
mtu 9216
portmode hybrid
switchport
no shutdown
```

Configure the two upstream interfaces connected to Storage array. Place the interfaces in Layer 2 mode with the switchport command and set the MTU to 9216 for performance.

```
interface TenGigabitEthernet 1/1
description SPA EP0
no ip address
mtu 9216
switchport
no shutdown

interface TenGigabitEthernet 1/2
description SPB EP0
no ip address
mtu 9216
switchport
no shutdown
```

Create VLAN interface 100 and tag all interfaces used in the SAN to VLAN 100. Attached configuration for iSCSI switch 2 non-converged shows use of VLAN 101 in this example.

```
interface Vlan 100
no ip address
mtu 9216
tagged TenGigabitEthernet 1/1-1/2
tagged Port-channel 100
no shutdown
```

Save the configuration.

```
end
write
```

5.4 Leaf switch configuration

The leaf switch configuration in this deployment guide outlines the basic configuration, such as port channels and VLT (VLT is specific to leaf switches). Spine switch configuration is beyond the scope of this document. For comprehensive configuration details, including routing for S4048-ON leaf switches and Z9100-ON spine switches, see [Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage](#).

5.4.1 Leaf S4048-ON switch configuration

The following configuration is for the Leaf A switch. See the attachment for configuration of the Leaf B switch. Initial configuration involves setting the hostname, configuring the management interface, default gateway and configuring the serial console enable password. If needed, restore the switch to factory defaults. See [Factory default settings](#).

```
enable
configure

hostname S4048-LF-A-U31

protocol lldp
  advertise management-tlv management-address system-description system-name
  advertise interface-port-desc

interface ManagementEthernet 1/1
  ip address 100.67.169.31/24
  no shut

management route 0.0.0.0/0 100.67.169.254

enable sha256-password <enable_password>

no ip telnet server enable
```

The two leaf switches are in VLT, using ports 1/53 and 1/54 for the VLT port channel.

Note: To configure VLT, the VLT domain must have the back-up destination as the management IP address of the secondary switch in the pair. In addition, the unit-id must be 0 and the secondary switch's unit-id must be 1.

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

vlt domain 127
  peer-link port-channel 127
  back-up destination 100.67.169.30
```

```

unit-id 0

interface fortyGigE 1/53
description VLTi
no ip address
no shutdown

interface fortyGigE 1/54
description VLTi
no ip address
no shutdown

```

Add the interfaces as members of respective port channels using LACP.

```

interface TenGigabitEthernet 1/11
description To FN410S A1
no ip address
port-channel-protocol LACP
port-channel 1 mode active
no shutdown

interface TenGigabitEthernet 1/12
description To FN410S A2
no ip address
port-channel-protocol LACP
port-channel 2 mode active
no shutdown

interface fortyGigE 1/49
description To S6010-ON Spine
no ip address
port-channel-protocol LACP
port-channel 3 mode active
no shutdown

interface fortyGigE 1/50
description To S6010-ON Spine
no ip address
port-channel-protocol LACP
port-channel 3 mode active
no shutdown

```

Port channels 1 and 2 connect to the IOMs. Port channel 3 connects to the S6010-ON spine switch.

```

interface Port-channel 1
description Port Channel to FN IOM A1
no ip address
portmode hybrid

```

```

switchport
vlt-peer-lag port-channel 1
no shutdown

interface Port-channel 2
description Port Channel to FN IOM A2
no ip address
portmode hybrid
switchport
vlt-peer-lag port-channel 2
no shutdown

interface Port-channel 3
description Connection to S6010 Spine
no ip address
portmode hybrid
switchport
vlt-peer-lag port-channel 3
no shutdown

```

Save the configuration.

```

end
write

```

To verify the VLT details, issue the following command. Ensure that the ICL Link Status, Heartbeat Status, and VLT Peer Status are Up.

```

S4048-LF-A-U31#show vlt brief
VLT Domain Brief
-----
Domain ID:                127
Role:                     Secondary
Role Priority:            32768
ICL Link Status:         Up
HeartBeat Status:        Up
VLT Peer Status:         Up
Local Unit Id:           0
Version:                  6(7)
Local System MAC address: 14:18:77:e0:69:31
Remote System MAC address: 14:18:77:7c:c4:e8
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

```

Verify port channels constructed using LACP with the following command.

```
S4048-LF-A-U31#show lacp summary
```

```
Port-channel    Member Ports
1                Te 1/11
2                Te 1/12
3                Fo 1/49,Fo 1/50
```

```
S4048-LF-A-U31#show lacp 1
```

```
Port-channel 1 admin up, oper up, mode lacp
LACP Fast Switch-Over Disabled
Actor   System ID:  Priority 32768, Address 1418.777c.c4e8
Partner System ID:  Priority 32768, Address f48e.383d.4e86
Actor Admin Key 1, Oper Key 1, Partner Oper Key 1, VLT Peer Oper Key 1
LACP LAG 1 is an aggregatable link
LACP LAG 1 is a VLT LAG
```

A - Active LACP, B - Passive LACP, C - Short Timeout, D - Long Timeout
E - Aggregatable Link, F - Individual Link, G - IN_SYNC, H - OUT_OF_SYNC
I - Collection enabled, J - Collection disabled, K - Distribution enabled
L - Distribution disabled, M - Partner Defaulted, N - Partner Non-defaulted,
O - Receiver is in expired state, P - Receiver is not in expired state

```
Port Te 1/11 is enabled, LACP is enabled and mode is lacp
```

```
Port State: Bundle
```

```
Actor   Admin: State ACEHJLMP Key 1 Priority 32768
        Oper: State ACEGIKNP Key 1 Priority 32768
Partner Admin: State BDFHJLMP Key 0 Priority 0
        Oper: State ACEGIKNP Key 1 Priority 32768
```

5.5 S3048-ON management switch configuration

For the S3048-ON management switches, configure all ports used in Layer 2 mode and in the default VLAN. They require no additional configuration and remain the same in both the traditional and converged topologies.

6 Physical switch configuration for converged deployments

6.1 FN410S configuration

The converged deployment example in this guide shows two methods. The first method uses the FN410S in Standalone mode, which requires minimal configuration on the IOM. The second method uses the FN410S in Full-switch mode, which does require configuration on the IOM. See [converged deployment topology diagram](#).

6.1.1 FN410S in Standalone mode

Even though the FN410S switches are fully operational in their factory default state, the following administrative commands may be run as needed. The configuration of FN410S in this example starts from their factory defaults state. See [Factory default settings](#).

Note: Configure the FNIOM (only for Standalone mode) via the Dell Blade I/O Manager. Access the Dell Blade I/O Manager through the FX2s CMC. See the attachment, [Dell Blade IO Manager V1.0.pdf](#), for more information.

```
enable
configure

hostname FN410S-A1

interface ManagementEthernet 0/0
 ip address 100.67.169.145/24
 no shutdown

management route 0.0.0.0/0 100.67.169.254
```

This deployment example uses a non-DCB environment. Hence, DCB must be manually disabled using the command:

```
no dcb enable
```

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

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

Note: SSH and Telnet are both enabled by default. It is a best practice to use SSH and disable Telnet for security.

6.1.2 FN410S in Full-switch mode

The following configuration is for the FN410S in slot A1. For configuration of the FN410S in slot A2, see the attachment “FN410S A2 FS mode Converged”. To change the mode of FN410S to Full-switch mode, see [FN410S switch Full-switch mode configuration](#). The configuration of FN410S in this example starts from their factory defaults state. See [Factory default settings](#). Initial configuration involves setting the hostname and configuring both the management interface and default gateway.

```
enable
configure

hostname FN410S-A1

protocol lldp
advertise management-tlv management-address system-name
no advertise dcbx-tlv ets-reco

interface ManagementEthernet 0/0
ip address 100.67.169.145/24
no shutdown

management route 0.0.0.0/0 100.67.169.254
```

This deployment example uses a non-DCB environment. Hence, DCB must be manually disabled using the command:

```
no dcb enable
```

Configure a serial console enable password and change the default SSH/telnet password for root. Replace password on each of the first two lines with the desired password. Disable telnet.

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**.

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

Configure internal and external ports along with the port channels and VLAN.

Note: Although all internal interfaces include configuration for switchport, portmode hybrid, and jumbo frames, only internal interfaces that are mapped to the dual-port adapters require switchport, portmode hybrid, and jumbo frames capability. See [FN IOM Internal Port Mapping](#) for internal port mappings.

```
interface Range TenGigabitEthernet 0/1-8
no ip address
```

```

mtu 9216
portmode hybrid
switchport
no shutdown

interface TenGigabitEthernet 0/11
no ip address
mtu 9216
port-channel-protocol LACP
port-channel 1 mode active
no shutdown

interface TenGigabitEthernet 0/12
no ip address
mtu 9216
port-channel-protocol LACP
port-channel 1 mode active
no shutdown

interface Port-channel 1
no ip address
mtu 9216
portmode hybrid
switchport
no shutdown

interface Vlan 100
no ip address
mtu 9216
tagged TenGigabitEthernet 0/1-8
tagged Port-channel 1
no shutdown

```

Save the configuration.

```

end
write

```

6.2 S4048-ON Leaf switch configuration

The following configuration is for the S4048-ON Leaf A switch. For configuration of the S4048-ON Leaf B switch, see the attachment. Initial configuration involves setting the hostname, configuring the management interface and default gateway. If needed, restore the switch to factory defaults. Refer [Factory default settings](#)

```

enable
configure

```

```

hostname Leaf-A

protocol lldp
  advertise management-tlv management-address system-description system-name
  advertise interface-port-desc

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

management route 0.0.0.0/0 100.67.169.254

enable sha256-password <enable_password>

no ip telnet server enable

```

The two leaf switches are in VLT, using ports 1/53 and 1/54 for the VLT interconnect (VLTi).

Note: To configure VLT, the VLT domain needs to have the management IP address of the secondary switch in the pair configured as the back-up destination. In addition, assign unit-id 0 to the primary switch and the unit-id 1 to the secondary switch.

```

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

interface fortyGigE 1/53
  description VLTi
  no ip address
  no shutdown

interface fortyGigE 1/54
  description VLTi
  no ip address
  no shutdown

interface Port-channel 127
  description VLTi
  no ip address
  channel-member fortyGigE 1/53,1/54
  no shutdown

```

Add the interfaces as members of the respective port channels using LACP.

```

interface TenGigabitEthernet 1/11
  description U25 FN410S A1 port 11

```

```

no ip address
mtu 9216
port-channel-protocol LACP
  port-channel 1 mode active
no shutdown

interface TenGigabitEthernet 1/12
description U25 FN410S A2 port 11
no ip address
mtu 9216
port-channel-protocol LACP
  port-channel 2 mode active
no shutdown

interface fortyGigE 1/49
description Spine Connection
no ip address
mtu 9216
port-channel-protocol LACP
  port-channel 49 mode active
no shutdown

```

Connect port channels 1 and 2 to the IOMs. Connect port channel 49 to the spine switch, S6010-ON.

```

interface Port-channel 1
description Port Channel to FN IOM A1
no ip address
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 1
no shutdown

interface Port-channel 2
description Port Channel to FN IOM A2
no ip address
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 2
no shutdown

interface Port-channel 49
description Spine Connection
no ip address
mtu 9216

```

```
portmode hybrid
switchport
vlt-peer-lag port-channel 49
no shutdown
```

Add the port channels to the respective VLANs for iSCSI.

```
interface Vlan 100
description iSCSI-1
no ip address
mtu 9216
tagged Port-channel 1-2,49
no shutdown

interface Vlan 101
description iSCSI-2
no ip address
mtu 9216
tagged Port-channel 1-2,49
no shutdown
```

Save the configuration.

```
end
write
```

6.3 S6010-ON Spine switch configuration

This section illustrates the configuration used for the Spine switch in this example. The configuration that this section describes covers the Layer 2 configurations. For Layer 3 routing configurations, refer [Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage](#)

```
enable
configure

hostname S6010-SPINE

protocol lldp
advertise management-tlv management-address system-description system-name
advertise interface-port-desc

interface ManagementEthernet 1/1
ip address 100.67.169.33/24
no shutdown

management route 0.0.0.0/0 100.67.169.254

enable sha256-password <enable_password>
```

```

no ip telnet server enable

interface fortyGigE 1/25
  description iSCSI 1 Connection
  no ip address
  mtu 9216
  port-channel-protocol LACP
  port-channel 48 mode active
  no shutdown

interface fortyGigE 1/26
  description iSCSI 2 Connection
  no ip address
  mtu 9216
  port-channel-protocol LACP
  port-channel 48 mode active
  no shutdown

interface fortyGigE 1/29
  description Leaf-A Connection
  no ip address
  mtu 9216
  port-channel-protocol LACP
  port-channel 49 mode active
  no shutdown

interface fortyGigE 1/30
  description Leaf-B Connection
  no ip address
  mtu 9216
  port-channel-protocol LACP
  port-channel 49 mode active
  no shutdown

```

Port channel 49 and port channel 48 connect to leaf switches and iSCSI storage switches.

```

interface Port-channel 48
  description iSCSI Connection
  no ip address
  mtu 9216
  portmode hybrid
  switchport
  no shutdown

interface Port-channel 49
  description Leaf Connection

```

```
no ip address
mtu 9216
portmode hybrid
switchport
no shutdown
```

Add the port channels to the respective VLANS for iSCSI.

```
interface Vlan 100
description iSCSI 1
no ip address
mtu 9216
tagged Port-channel 48-49
no shutdown
```

```
interface Vlan 101
description iSCSI 2
no ip address
mtu 9216
tagged Port-channel 48-49
no shutdown
```

Save the configuration.

```
end
write
```

6.4 iSCSI SAN S4048-ON switch configuration

There are two iSCSI SAN switches used in the deployment example in this guide. The following configuration details are specific to switch S4048-ON iSCSI-1. Refer attachment for configuration of Switch S4048-ON iSCSI-2. The switches start at their factory default settings. Refer to [Factory default settings](#)

```
enable
configure

hostname ISCSI-1

protocol lldp
advertise management-tlv management-address system-description system-name
advertise interface-port-desc

interface ManagementEthernet 1/1
ip address 100.67.169.35/24
no shutdown

management route 0.0.0.0/0 100.67.169.254
```

```
enable sha256-password <enable_password>
```

```
no ip telnet server enable
```

The two iSCSI switches are in VLT and ports 1/53 and 1/54 are utilized for the VLT port channel.

```
vlt domain 127
peer-link port-channel 127
back-up destination 100.67.169.34
unit-id 0
```

```
interface fortyGigE 1/53
description VLTi
no ip address
no shutdown
```

```
interface fortyGigE 1/54
description VLTi
no ip address
no shutdown
```

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

Interfaces Te 1/1 and 1/2 are connected to the Ethernet ports on the Storage Processors on The Storage array. Port Channel 48 runs upto the Spine switch.

```
interface TenGigabitEthernet 1/1
description SPA EP0
no ip address
mtu 9216
switchport
no shutdown
```

```
interface TenGigabitEthernet 1/2
description SPB EP0
no ip address
mtu 9216
switchport
no shutdown
```

```
interface fortyGigE 1/49
description Spine Connection
no ip address
mtu 9216
```

```
port-channel-protocol LACP
  port-channel 48 mode active
no shutdown

interface Port-channel 48
description Spine Connection
no ip address
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 48
no shutdown

interface Vlan 100
no ip address
mtu 9216
tagged TenGigabitEthernet 1/1-1/2
tagged Port-channel 48
no shutdown

interface Vlan 101
no ip address
mtu 9216
tagged TenGigabitEthernet 1/1-1/2
tagged Port-channel 48
no shutdown
```

Save the configuration.

```
end
write
```

7 Virtual Network configuration

This deployment guide uses FC630 servers running ESXi 6.5 and strives to highlight the implementation of iSCSI using vSphere web client. While iSCSI can be implemented using Windows server, this deployment guide focuses on the use case where VMware ESXi 6.5 is running on servers.

7.1 Information on vSphere Standard Switch

A vSphere standard switch (also referred to as a VSS or a standard switch) is a virtual switch that handles network traffic at the host level in a vSphere deployment. Standard switches provide network connectivity to hosts and virtual machines.

A standard switch named vSwitch0 is automatically created on each ESXi host during installation to provide connectivity to the management network.

Standard switches may be viewed, and optionally configured, as follows:

1. Go to the web client **Home** page, select **Hosts and Clusters**, and select a host in the **Navigator** pane.
2. In the center pane, select **Configure > Networking > Virtual switches**.
3. Standard switch **vSwitch0** appears in the list. Click on it to view details as shown in Figure 15.

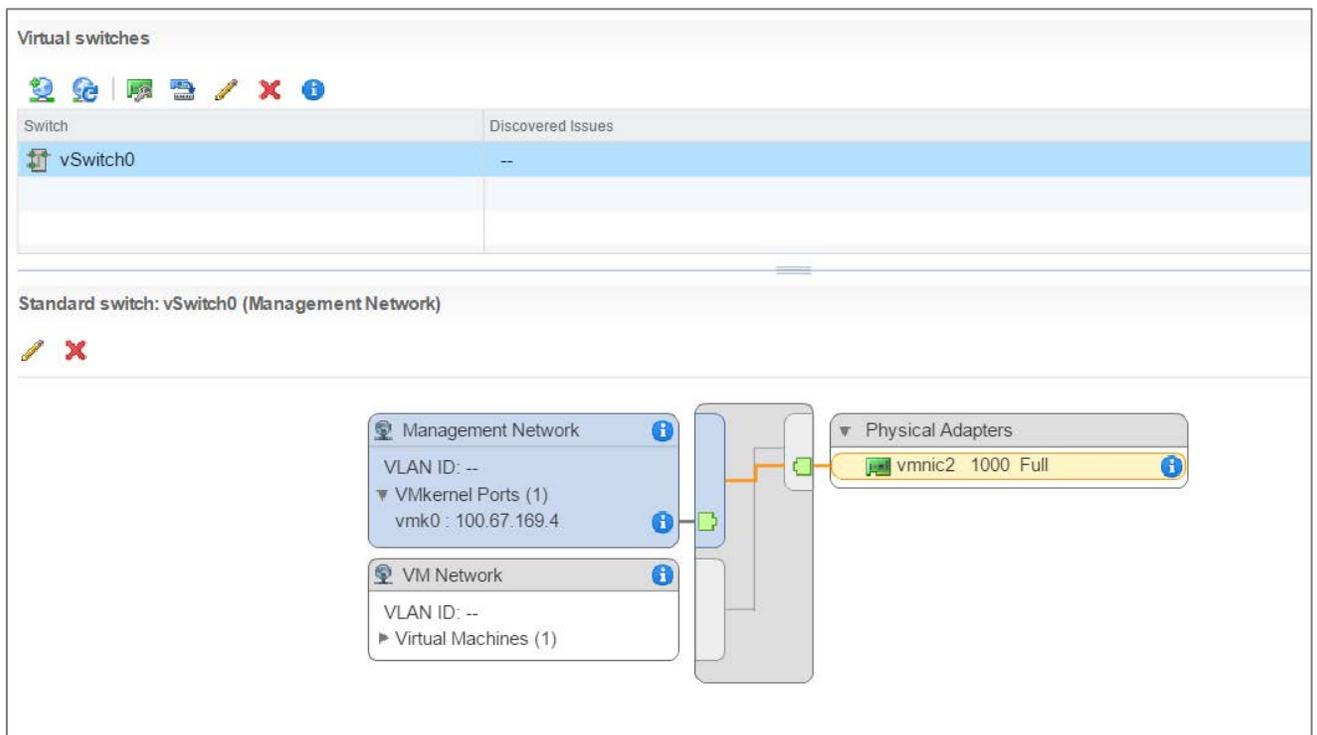


Figure 15 Example of vSphere standard switch

Note: For this guide, only the default configuration is required on the standard switches. Standard switches are only used in this deployment for connectivity to the management network. Distributed switch covered in the next section, are used for Ethernet and storage traffic.

7.2 Information on vSphere Distributed Switch

A vSphere distributed switch (also referred to as a VDS or a distributed switch) is a virtual switch that handles network traffic at the vCenter level in a vSphere deployment. Distributed switches provide network connectivity to hosts and virtual machines. Distributed switches must be manually created at the vCenter and associated with each ESXi host.

7.3 Create an iSCSI VDS for the FX2-FC630-iSCSI cluster

In the vSphere web client, 4 hosts that were added to a cluster named FX2-FC630-iSCSI. To create a VDS for iSCSI traffic for FX2-FC630-iSCSI cluster:

1. On the web client **Home** screen, select **Networking**.
Right click on Datacenter. Select **Distributed switch > New Distributed Switch**.
Provide a name for the VDS, e.g. **VDS-iSCSI**. Click **Next**.
On the Select version page, select **Distributed switch: 6.5.0** and click **Next**.
On the **Edit settings** page:
 - Change the **Number of uplinks** to **2**.
 - Leave **Network I/O Control** set to **Enabled**.
 - Uncheck the **Create a default port group** box.Click **Next > Finish**.

The VDS is created with the uplink port group shown beneath it. When complete, the **Navigator** pane should look similar Figure 16.

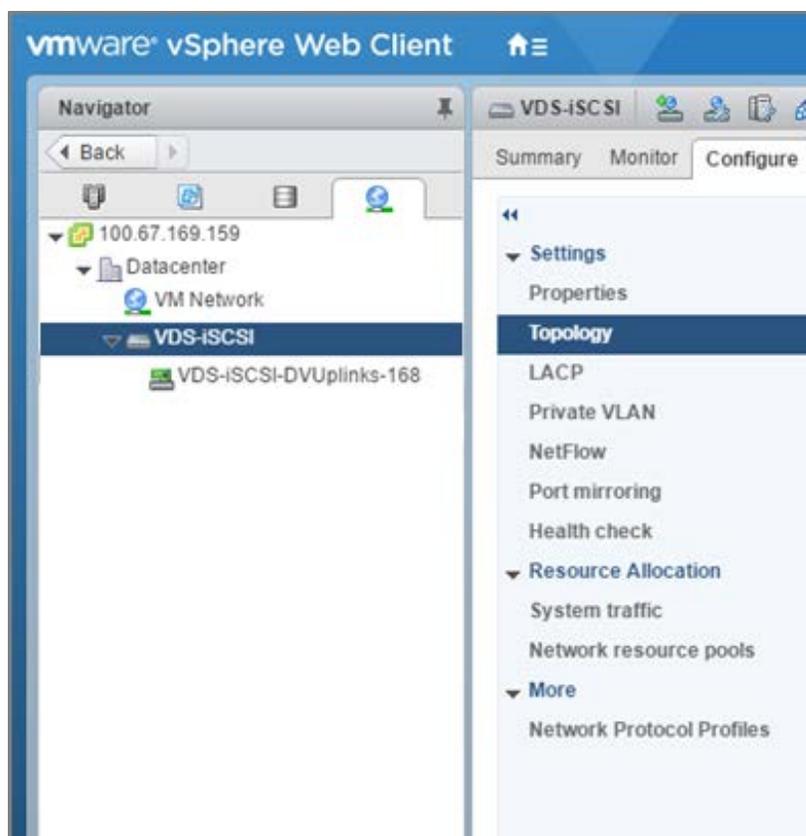


Figure 16 VDS-iSCSI created

Note: The VDS that were already existing in the environment are displayed along with the newly created VDS-iSCSI. In this deployment guide, only VDS-iSCSI will be used.

7.4 Set the iSCSI VDS MTU to 9000 and enable LLDP

Dell EMC recommends increasing the Maximum Transmission Unit (MTU) of devices handling storage traffic to 9000 bytes for best performance. LLDP may be configured at the same time.

To configure the VDS-iSCSI:

1. Go to **Home > Networking**.
2. Right click on, VDS-iSCSI and select **Settings > Edit Settings**.
3. In the left pane of the **Edit Settings** page, click **Advanced**.
4. Set **MTU (Bytes)** to 9000.
5. Under Discovery protocol, set Type to **Link Layer Discovery Protocol** and Operation to **Both**.

The window should appear similar to Figure 17.

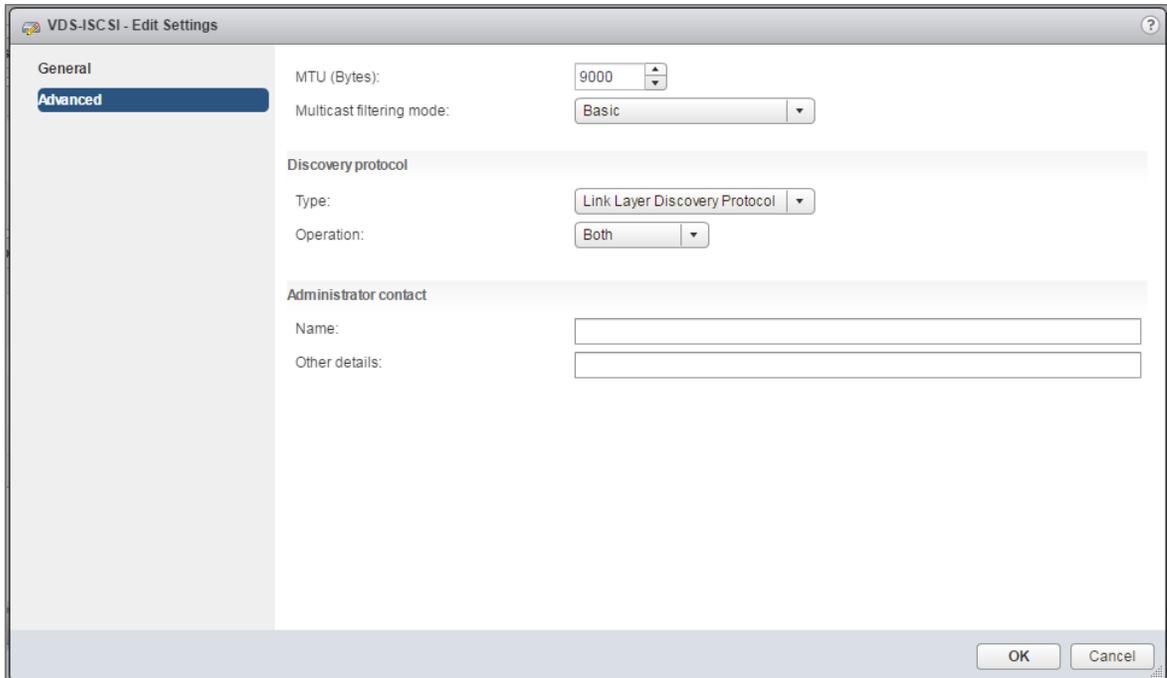


Figure 17 VDS-ISCSI: Edit Settings window

6. Click **OK** to apply the settings.

7.5 Add distributed port groups

In this section, two distributed port groups for iSCSI traffic are added to VDS-ISCSI created in the previous section. In addition, a distributed port group DPG-Production is added for the non-iSCSI traffic.

To create the distributed port groups on the **VDS-ISCSI**:

1. On the web client **Home** screen, select **Networking**.
Right click on VDS-ISCSI. Select **Distributed Port Group > New Distributed Port Group**.
On the **Select name and location** section, provide a name for the first distributed port group, e.g. **DPG-
ISCSI-1**. Click **Next**.
On the **Configure settings** page, next to **VLAN type**, select **VLAN**. Set the **VLAN ID** to **100**. Leave other values at their defaults as shown in Figure 18.

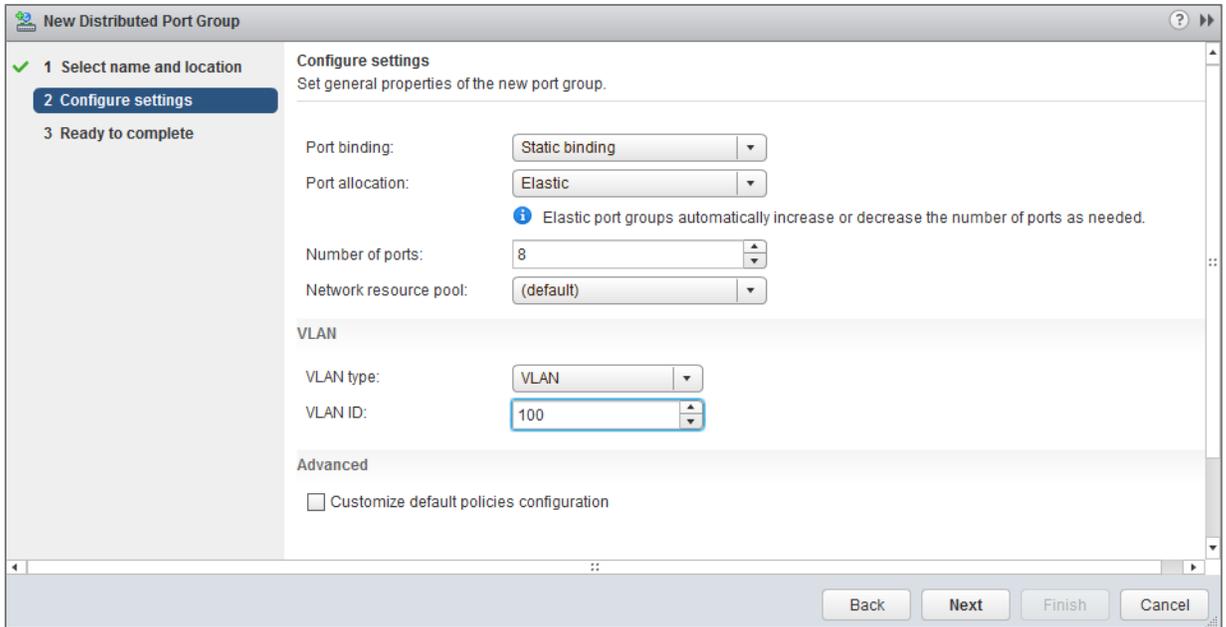


Figure 18 DPG-ISCSI-1 Distributed Port Group settings page

Click **Next > Finish**.

Repeat steps 2-5 above for the second port group. Provide a unique name, e.g. **DPG-ISCSI-2** and set its VLAN ID to **101**. When complete, the Navigator pane will appear similar to Figure 19.

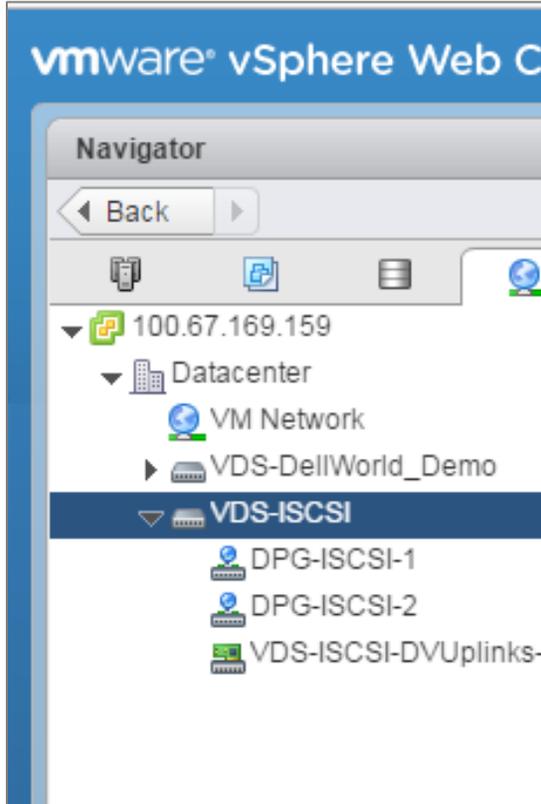


Figure 19 iSCSI distributed port groups created

Repeat the steps above to create a port group named DPG-Production. In step 4, use the default VLAN settings (no VLAN ID specified). The DPG-production port group is used for non-iSCSI traffic between VMs. When complete, the Navigator pane will appear similar to Figure 20.

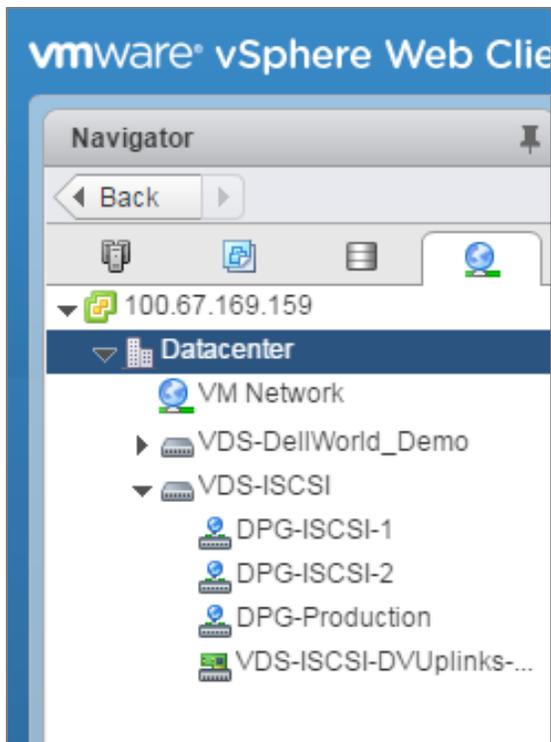


Figure 20 Distributed Port Groups created in total

7.6 Configure teaming and failover

1. On the web client **Home** screen, select **Networking**.
 Right click on VDS-ISC SI. Select **Distributed Port Group > Manage Distributed Port Groups**.
 Select only the **Teaming and failover** checkbox. Click **Next**.
 Click **Select distributed port groups**. Check the box next to DPG-ISC SI-1. Click **OK > Next**.
 On the **Teaming and failover** page, click **Uplink 1** and move it up to the **Active uplinks** section by clicking the up arrow. Move **Uplink 2** down to the **Unused uplinks** section. Leave other settings at their defaults. The **Teaming and failover** page should look similar to Figure 21 when complete.

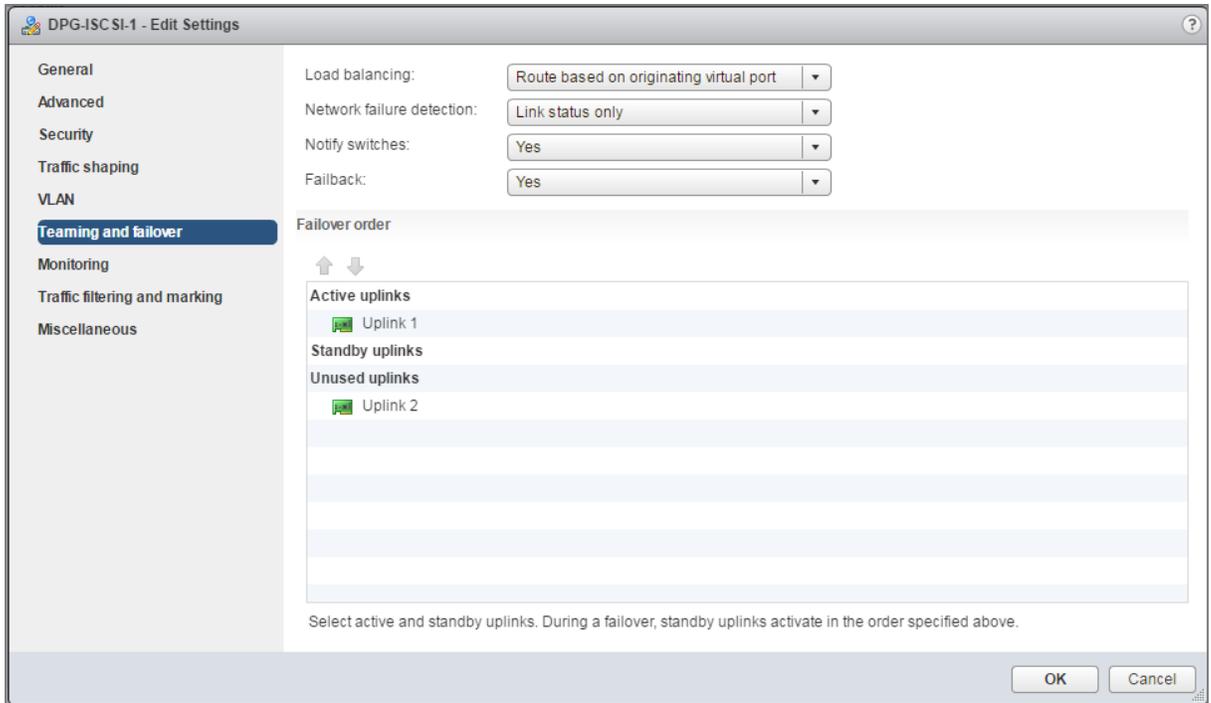


Figure 21 Teaming and failover settings for the DPG-ISCSI-1 port group

Click **Next > Finish** to apply the settings.

Repeat steps 2-6 above for **DPG-ISCSI-2** port group and **DPG-Production** port group. For **DPG-ISCSI-2** port group, make sure that **Uplink 2** is moved to **Active** and **Uplink 1** is moved to **Unused**. For **DPG-Production**, in this example, **Uplink 2** is moved to **Standby** and **Uplink 1** is moved to **Active**.

7.7 Associate hosts and assign uplinks

Hosts and their vmnics must be associated with VDS-ISCSI.

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 select **Configure > Networking > Physical adapters**. Adapter MAC addresses can be determined by connecting to the iDRAC. In this example, vmnics used are numbered vmnic0 and vmnic1. vmnic numbering will vary depending on adapters installed in the host.

To add hosts to the VDS-ISCSI:

1. On the web client **Home** screen, select **Networking**.
Right click VDS-ISCSI and select **Add and Manage Hosts**.
2. In the Add and Manage Hosts dialog box:
On the **Select task** page, make sure **Add hosts** is selected. Click **Next**.
On the **Select hosts** page, Click the **+ New hosts** icon. Select the check box next to each host in the cluster FX2-FC630-ISCSI. Click **OK > Next**.
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**.
On the **Manage physical network adapters** page, each host is listed with its vmnics beneath it.
Select the first vmnic (vmnic0 in this example) on the first host and click  **Assign uplink**.
Select **Uplink 1 > OK**.
Select the second vmnic (vmnic1 in this example) on the first host and click  **Assign uplink**.
Select **Uplink 2 > OK**.
Repeat steps i – iv for the remaining hosts. Click **Next** when done.
On the Analyze impact page, Overall impact status should indicate  **No impact**.
Click **Next > Finish**.

7.8 Add VMkernel adapters for iSCSI

In this section, two iSCSI VMkernel adapters (also referred to as VMkernel ports) are added to each ESXi host to allow for multipath iSCSI traffic.

IP addresses can be statically assigned to VMkernel adapters upon creation, or DHCP may be used. Static IP addresses are used in this guide.

This deployment uses the following addressing scheme for the iSCSI networks:

Table 4 iSCSI VLANs and networks

VLAN ID	Network	Used For
100	192.168.100.0/24	iSCSI-1
101	192.168.101.0/24	iSCSI-2

Table 5 Host iSCSI IP Addresses

Host	iSCSI-1	iSCSI-2
FC630-1	192.168.100.41	192.168.101.41
FC630-2	192.168.100.42	192.168.101.42
FC630-3	192.168.100.43	192.168.101.43
FC630-4	192.168.100.44	192.168.101.44

To add a VMkernel adapter to each host connected to the VDS-iSCSI:

- On the web client **Home** screen, select **Networking**.
 Right click on **VDS-iSCSI**, and select **Add and Manage Hosts**.
 In the Add and Manage Hosts dialog box:
 On the **Select task** page, make sure **Manage host networking** is selected. Click **Next**.
 On the **Select hosts** page, click **+ Attached hosts**. Select all hosts. Click **OK > Next**.
 On the **Select network adapter tasks** page, make sure the **Manage VMkernel adapters** box is checked and all other boxes are unchecked. Click **Next**.
 The **Manage VMkernel network adapters** page opens.
 To add the first iSCSI adapter, select the first host and click **+ New Adapter**.
 On the **Select target device** page, click the radio button next to **Select an existing network** and click **Browse**.
 Select **DPG-iSCSI-1**. Click **OK > Next**.
 On the **Port properties** page, leave **IPv4** selected and make sure no boxes are checked next to **Enable services**. Click **Next**.
 On the **IPv4 settings page**, if DHCP is not used, select **Use static IPv4 settings**. Set the IP address, for example **192.168.100.41**, and subnet mask, **255.255.255.0**, for the host on the first iSCSI network. Click **Next > Finish**.
 Repeat steps i-v to add a second iSCSI VMkernel adapter on network **DPG-iSCSI-2** with an appropriate IP address, for example **192.168.101.41**. Repeat for the remaining hosts resulting in two VMkernel adapters per host, with one on each network. Click **Next**.
 On the **Analyze impact** page, **Overall impact status** should indicate **✔ No impact**.
 Click **Next > Finish**.

When complete, the **VMkernel adapters** page for each ESXi host in the vSphere data center should look similar to Figure 22. This page is visible by going to **Hosts and Clusters**, selecting a host in the **Navigator** pane, then selecting **Configure > Networking > VMkernel adapters** in the center pane.

Make sure the adapters are configured correctly on each host.

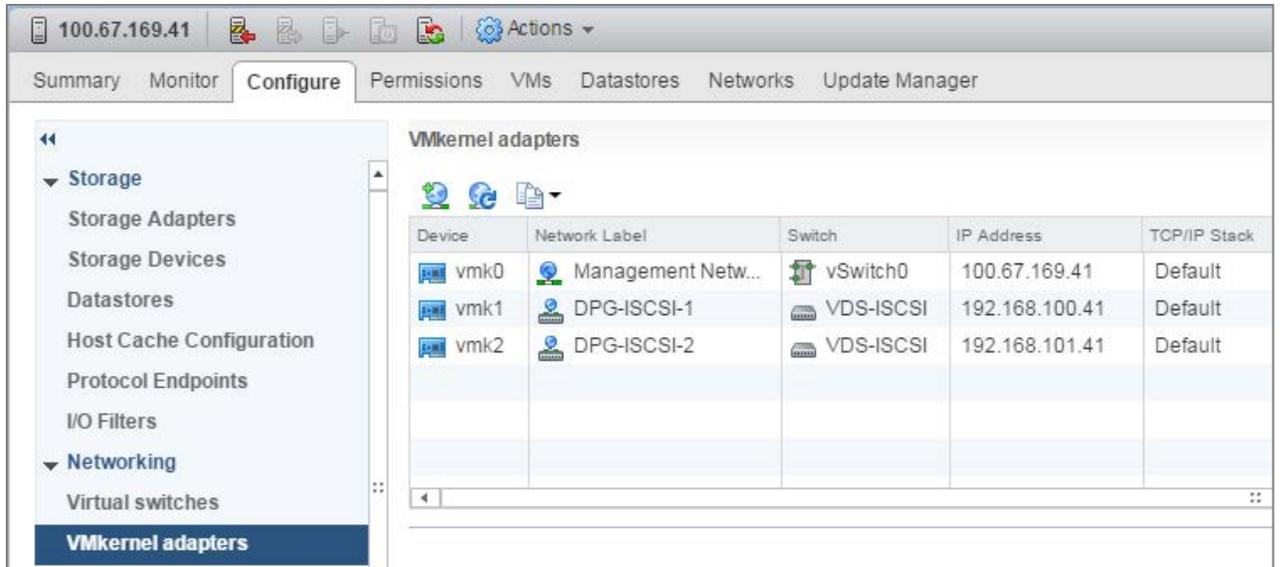


Figure 22 Host VMkernel adapters page with iSCSI networking configured

When complete, the **Configure > Settings > Topology** page for **VDS-ISCSI** should look similar to Figure 23.

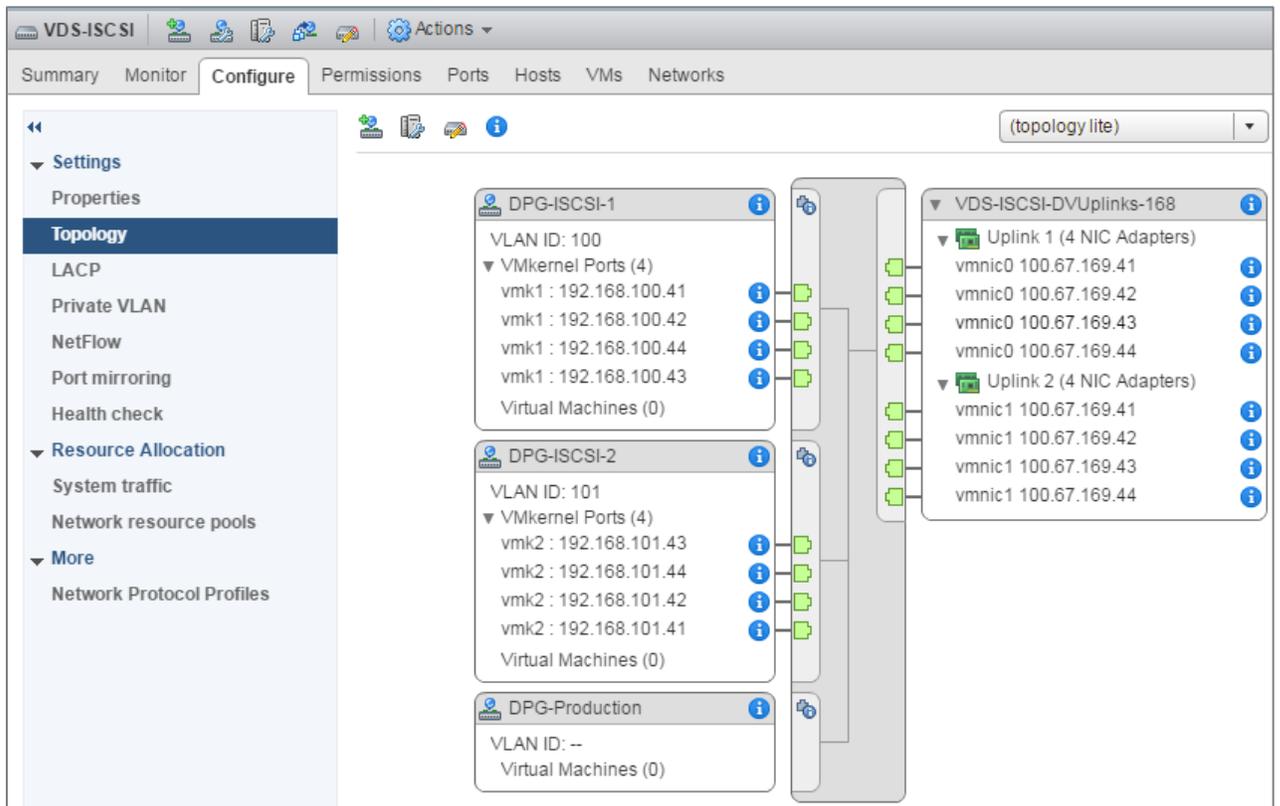


Figure 23 VDS-ISCSI vmnic configuration

7.9 Increase the MTU to 9000 on iSCSI VMkernel adapters

Dell EMC recommends increasing the Maximum Transmission Unit (MTU) of devices handling storage traffic to 9000 bytes for best performance.

To set the MTU to 9000 bytes on iSCSI VMkernel adapters:

1. Go to **Home > Hosts and Clusters** and select the first host in the compute cluster.
2. In the center pane, select **Configure > Networking > VMkernel adapters**.
3. Select the **vmk1** VMkernel adapter. Click  **Edit settings**.
4. Select **NIC settings** and change the MTU to **9000**.
5. Click **OK**.
6. Repeat for the **vmk2** VMkernel adapter.

Repeat steps 1-6 on the remaining hosts in the compute cluster.

7.10 Bind iSCSI adapters with VMkernel ports

1. Go to **Home > Hosts and Clusters**.

In the **Navigator** pane, select a host.

In the center pane, select **Configure > Storage > Storage adapters**.

Select a connected physical adapter listed under **QLogic 57810 10 Gigabit Ethernet Adapter**, e.g. **vmhba33**.

Under **Adapter Details**, click the Network Port Binding tab.

Click the  icon to open the **Bind vmhbaxx with VMkernel Adapter** window.

Select the **DPG-ISCSI-1(VDS-ISCSI) port group**.

Click **OK**

Click the  icon to rescan the host's storage adapter.

Repeat steps 4 through 7 for the second vmhba (e.g. vmhba34) on the host. In step 6.a., connect it to the **DPG-ISCSI-2** port group.

Repeat the steps above for the remaining hosts in the FX2-FC630-ISCSI cluster.

When complete, each host's **Storage Adapters** page should look similar to Figure 24.

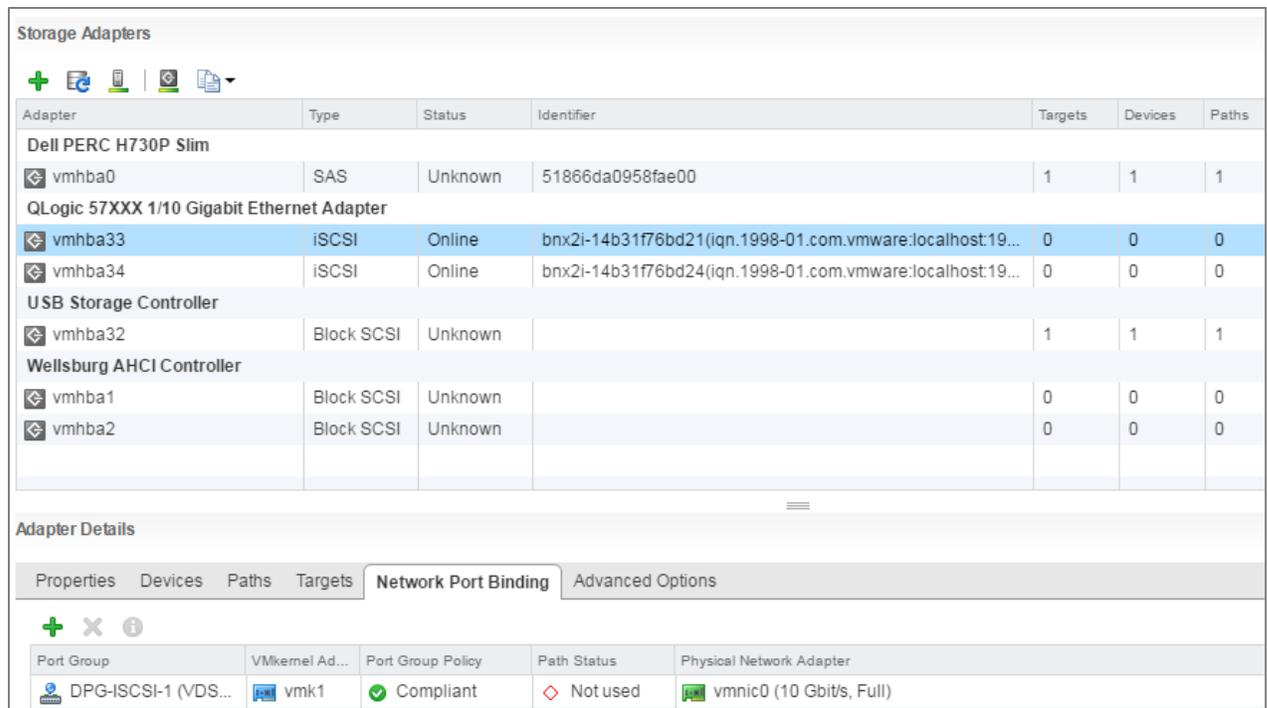


Figure 24 Storage adapters details page

7.11 Configure dynamic discovery

With dynamic discovery, each time the initiator contacts the storage system, it sends a send targets request to the system. The Storage array responds by supplying a list of available targets to the initiator.

To configure dynamic discovery for iSCSI:

1. Go to **Home > Hosts and Clusters**.

In the **Navigator** pane, select the first host in the **FX2-FC630-ISC**SI cluster.

In the center pane, select **Configure > Storage > Storage adapters**.

Select host's first connected iSCSI adapter listed under **QLogic 57810 10 Gigabit Ethernet Adapter**, vmhba33 for example.

Under **Adapter Details**, click the **Targets** tab

Select Dynamic Discovery and click **Add** to open the Add Send Target Server window.

Next to **iSCSI Server**, enter the IP address for the iSCSI target created in Unisphere. In this example,

192.168.100.1 or 192.168.100.3 is provided. Refer [Configure iSCSI Interfaces](#)

Leave the other settings at their default values and click **OK**.

Click the  icon to rescan the host's storage adapter.

Repeat steps 4-9 for the host's second connected iSCSI adapter. In this example, the IP address configured for the iSCSI target, **192.168.101.2 or 192.168.101.4** is used.

Repeat the above for the remaining hosts in the compute cluster.

Note: The storage LUN presented by the Storage array will be available, once the configuration of Storage device is completed. See the following section on configuring iSCSI on the Storage device.

8 Configure iSCSI Storage

This section covers the iSCSI configuration of the storage device using the Unisphere web interface. Dell EMC Unisphere presents a new approach to unified storage management through a simple, flexible, and integrated user experience. Information is consolidated and visible through a single lens and managing storage is simplified by providing an intuitive, context-based approach.

Note: The below configuration is used as an example and is specific to the storage device used in this example. See your storage configuration guide to configure your storage array for iSCSI.

Figure 25 references the SAN connectivity from the FN410S to the Storage Processors A and B, using two S4048-ON switches. Table 6 provides connection details from the S4048-ON switches to the Storage Processors A and B.

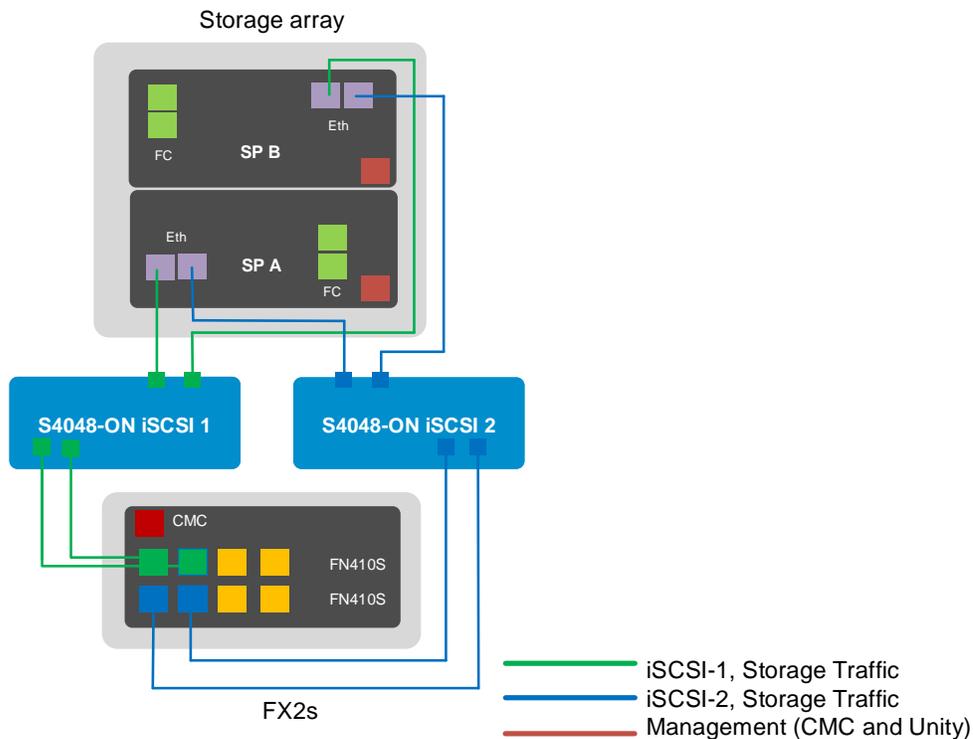


Figure 25 iSCSI SAN

Table 6 Switch to storage connections

Switch	Switch port #	VLAN	Storage Processor (SP)	port #	Controller port IP address
S4048-ON iSCSI-1	Te 1/1	100	SP A	Ethernet Port 0	192.168.100.1
S4048-ON iSCSI-1	Te 1/2	100	SP B	Ethernet Port 0	192.168.100.3
S4048-ON iSCSI-2	Te 1/1	101	SP A	Ethernet Port 1	192.168.101.2
S4048-ON iSCSI-2	Te 1/2	101	SP B	Ethernet Port 1	192.168.101.4

8.1 Create a Storage Pool

1. In the left pane, under **STORAGE**, choose **Pools**.

Click on the **+** icon to create a new Pool. In the dialog box that opens up, provide a **Name** for the pool.

Provide a **Description** if needed. Click **Next**.

Select the Storage Tiers for the Pool. Choose from **Extreme Performance Tier**, **Performance Tier** and **Capacity Tier** as per your requirement. Once you choose the Tier(s), if needed, change the RAID Configuration for chosen tiers. Click **Next**.

In the **Disks** section, select the **Amount of Storage**. The total number of disks and the total capacity will be displayed next to **Totals** label. Click **Next**.

Leave the **Capability Profile Name** section as it is and click **Next**.

Review your selections in the **Summary** section and click Finish. The **Results** section displays the **Overall status** of Storage Pool being created, in the form of a percentage number. Once the **Overall status** shows 100% Completed, click **Close**.

The newly created Storage Pool is now visible under the **Pools** section, as shown in Figure 26.

Note: If you close the **Results** section before the **Overall status** shows 100% completed, it continues to run in the background.

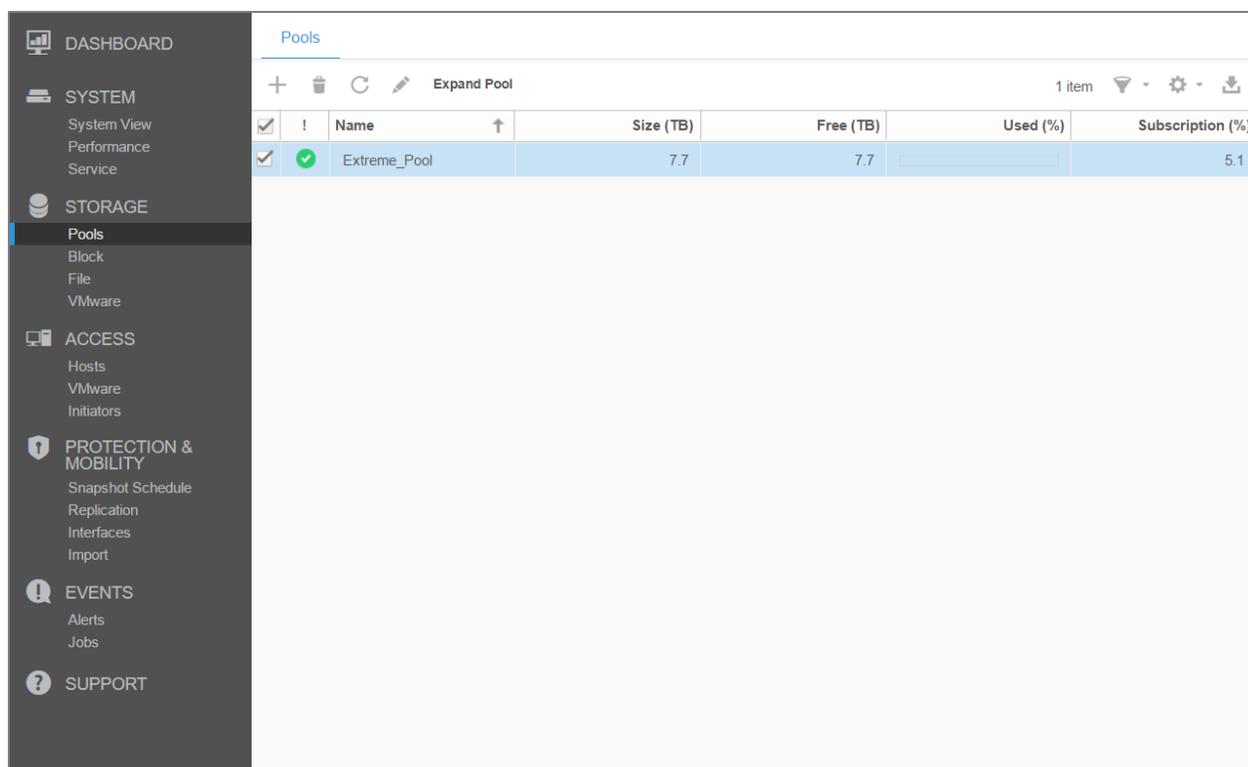


Figure 26 Storage pool created

8.2 Add vCenter to Unisphere

1. Launch the Unisphere GUI in a web browser.
2. In the left pane, under **ACCESS**, choose **VMware > vCenters**.
3. Click on the **+** icon to open the **Add vCenter** dialog box.
4. Enter the **Network Name or Address**, **User name** and **Password** for the vCenter Server and click **Find**.
5. The list of **ESXi hosts** that can be imported from the vCenter is displayed. Choose the needed ESXi hosts and click **Next**.
6. In the **Summary** section, review the ESXi Hosts managed by the vCenter Server that will be added to VMware Hosts. Click **Finish**.
7. The **Results** section displays the **Overall status** of discovering the vCenter, in the form of a percentage number. Once the **Overall status** shows 100% Completed, click **Close**.
8. The vCenter server should be displayed as shown in Figure 27.

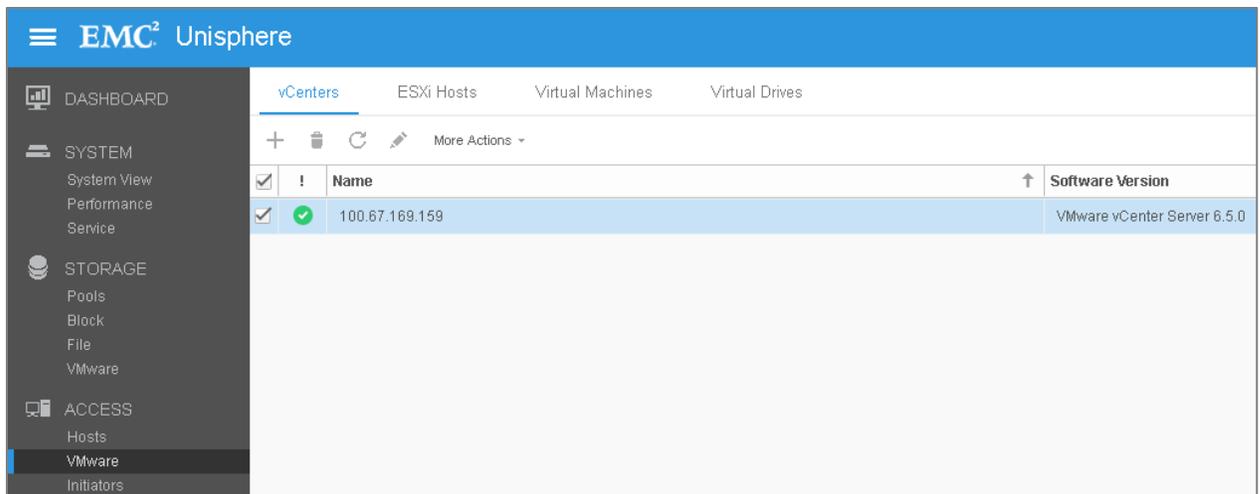


Figure 27 vCenter server added to Unisphere

9. The list of added of ESXi hosts should be displayed under the ESXi Hosts tab, as shown in Figure 28.

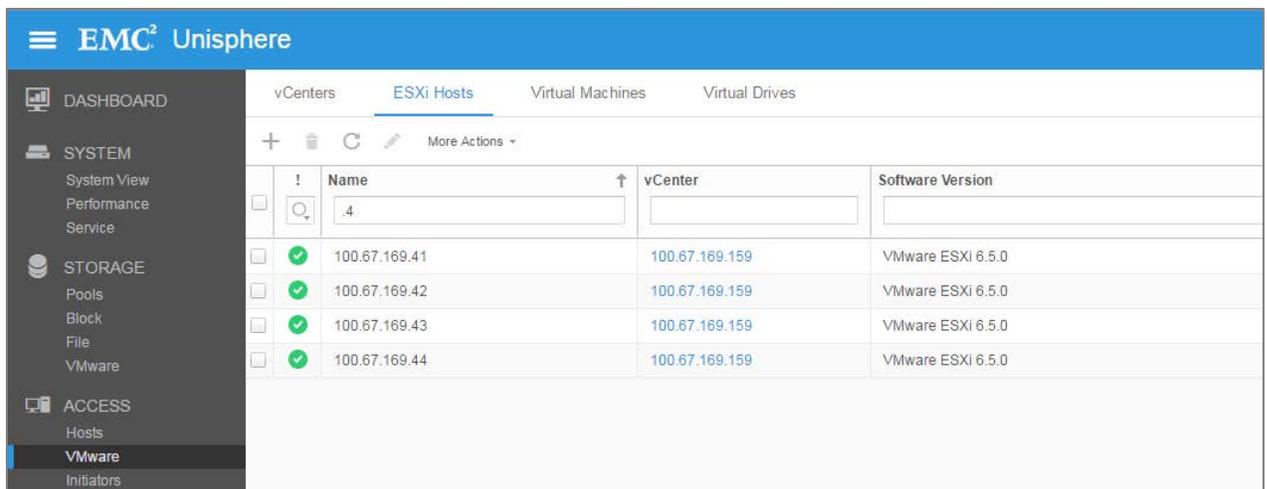


Figure 28 ESXi Hosts

8.3 Create a LUN

1. In the left pane of Unisphere web client, under **STORAGE**, choose **Block > LUNs**.
2. To create a new LUN for iSCSI, click on the **+** icon. A **Create LUN** dialog box opens up.
3. In the **Configure** section, select the **Number of LUNs**. Provide a **Name** and choose the **Storage Pool** from which the LUN will be created. Modify the size of the **LUN** as required. Click **Next**.
4. In the **Access** section, select the hosts that can access the storage resource. Click on the **+** icon and choose the hosts. If required, you can also add hosts at a later point in time. Click **Next**.

5. **Snapshot and Replication** section are not configured in this deployment guide. Click **Next** to move through these sections.
6. In the **Summary** section, review the details and click **Finish**.
7. The **Results** section displays the **Overall status** of LUNs being created, in the form of a percentage number. Once the **Overall status** shows 100% Completed, click **Close**.

The newly created LUN, named iSCSI LUN, is now visible under the **LUNs** section, as shown in Figure 29.

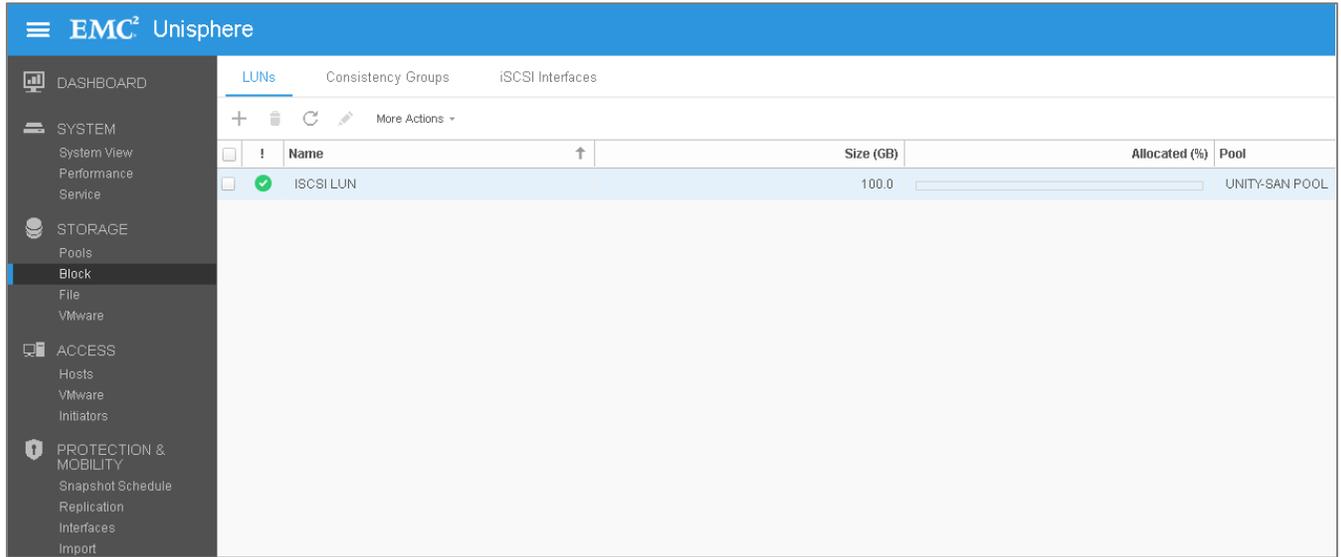


Figure 29 LUN created

8.4 Configure iSCSI Interfaces

1. Launch the Unisphere GUI in a web browser.
In the left pane, under **STORAGE**, choose **Block > iSCSI Interfaces**.

Click on the **+** icon to add an iSCSI interface. Provide the **IP Address**, **Subnet Mask / Prefix Length** and **VLAN ID**. Click **OK**.

Add a total of 4 iSCSI Interfaces. When completed, it should look similar to Figure 30.

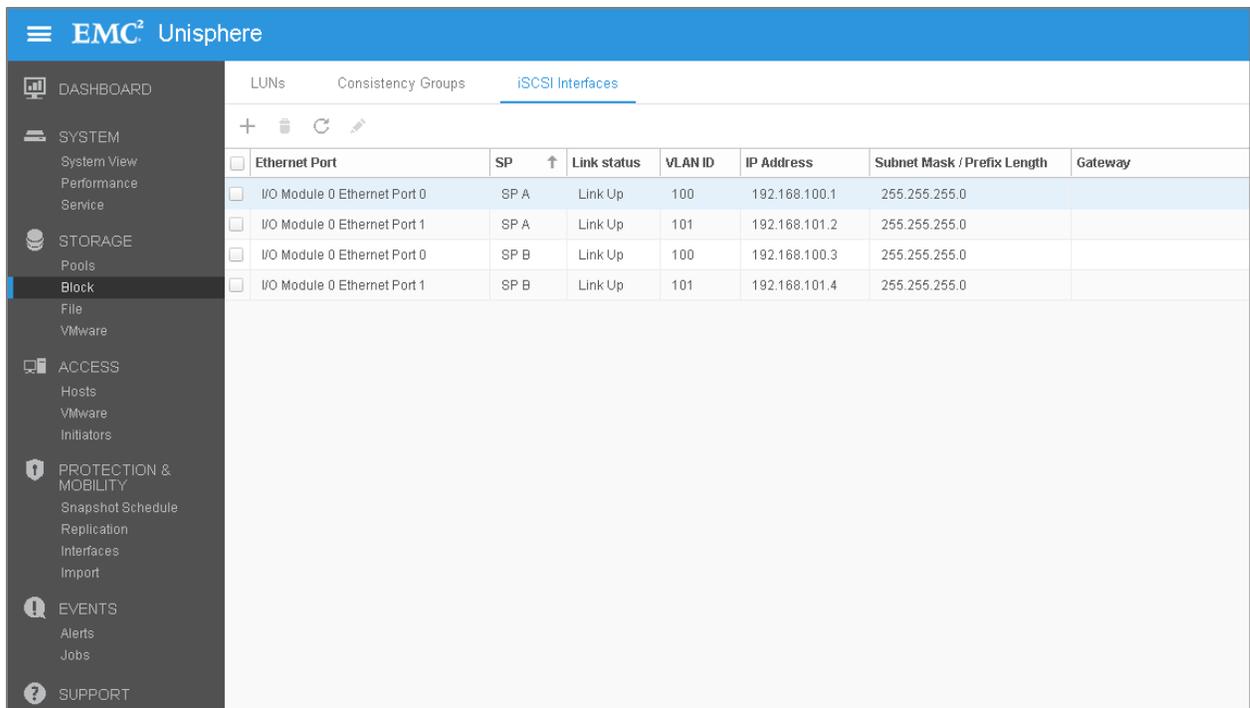


Figure 30 iSCSI interfaces configured

8.5 Rescan storage on hosts

Note: Return to the vSphere Web Client starting with this section.

1. On the vSphere Web Client **Home** screen, select **Hosts and Clusters**. In the **Navigator** pane, select the first host in the **FX2-FC630-ISCSI** cluster. In the center pane, select **Configure > Storage > Storage adapters** and select the host's first storage adapter (e.g. vmhba33).

Click the  icon to rescan for newly added storage devices.

Under **Adapter Details**, select the **Devices** tab. The volume appears as shown in Figure 31.

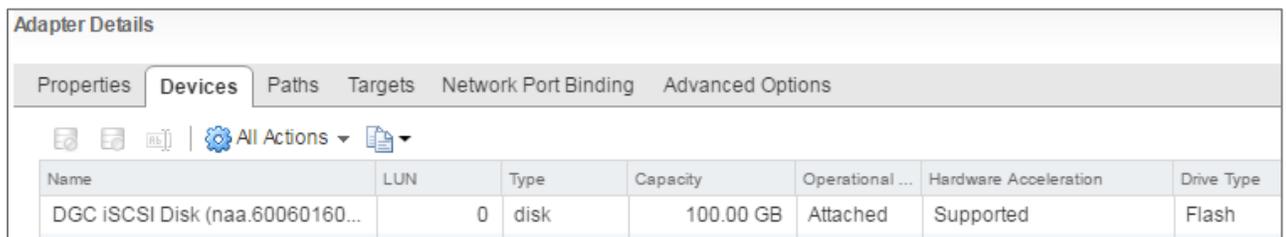


Figure 31 Devices tab

Select the **Paths** tab. The target name, LUN number and status are shown. The status field is marked either **Active** or **Active (I/O)** as shown in Figure 32.

Adapter Details			
Properties Devices Paths Targets Network Port Binding Advanced Options			
Enable		Disable	
Runtime Name	Target	LUN	Status
vmhba34:C0:T0:L0		0	◆ Active
vmhba34:C0:T3:L0		0	◆ Active

Figure 32 Paths tab

Repeat steps 3-6 above for the host's second storage adapter (e.g. vmhba34).

Repeat the above for the remaining hosts in the cluster. All hosts should have two active connections to the shared storage volume.

8.6 Create a datastore

Create a datastore that uses the shared storage volume.

To create the datastore:

1. Go to **Home > Storage**.

In the **Navigator** pane, right click on **Datacenter** and select **Storage > New Datastore**.

In the **New Datastore** window, for **Location**, **Datacenter** is selected. Click **Next**.

Leave the **Type** set to **VMFS** and click **Next**.

On the Name and device selection page:

Provide a **Datastore name**, e.g. **Datastore iSCSI**.

From the dropdown menu, select *any* host in the cluster. When created, this datastore will be accessible to all configured hosts in the cluster (refer to the note on the screen next to the  icon).

Click on the LUN and click **Next**.

In the VMFS version section, choose **VMFS 6** or **VMFS 5**. In this example, VMFS 6 is selected as show in Figure 33. Click **Next**.

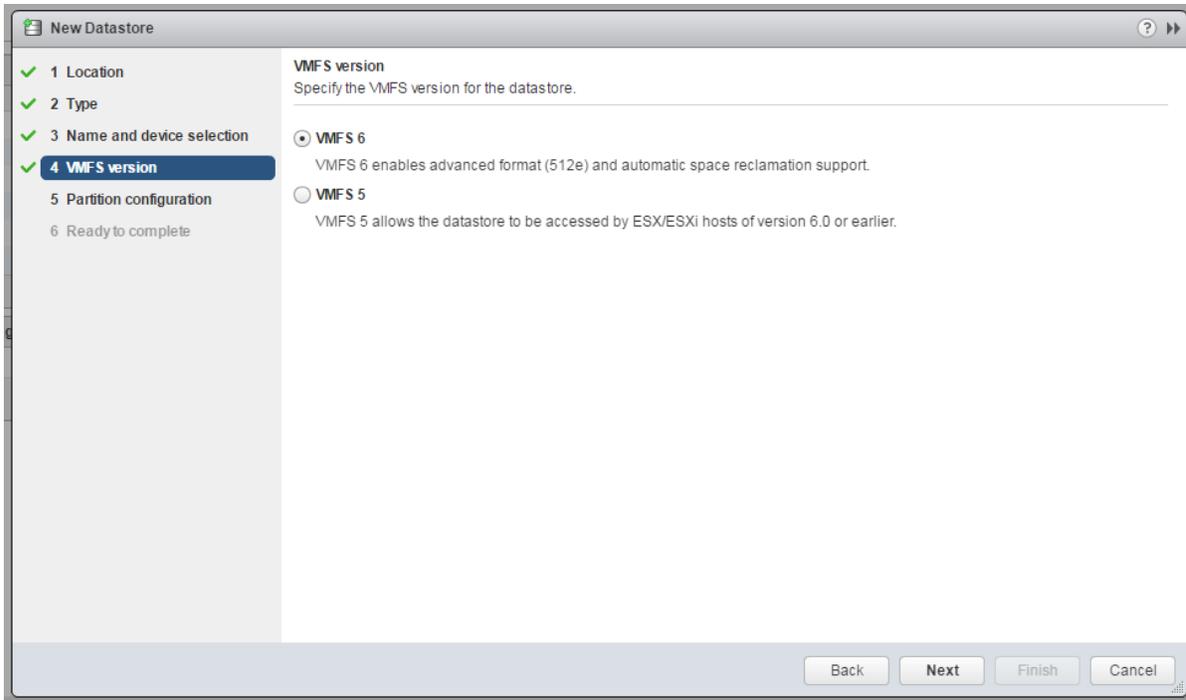


Figure 33 VMFS version

Leave the **Partition configuration** at its default settings and click **Next > Finish** to create the datastore.

To verify the datastore has been mounted by all hosts in the compute cluster:

1. Go to **Home > Storage**.

In the **Navigator** pane, select the newly created datastore, **Datastore iSCSI**.

In the center pane, select **Configure > Connectivity and Multipathing**.

All hosts in the compute cluster are listed with the datastore status shown as **Mounted** and **Connected** as per Figure 34.

Connectivity and Multipathing			
Select a host to view the multipathing details for its devices.			
Mount		Unmount	
Host	Datastore Mounted	Datastore Connectivity	Mount Point
100.67.169.44	Mounted	Connected	/vmfs/volumes/58c184c6-f088cbc4-e3df-e0db55117da4
100.67.169.42	Mounted	Connected	/vmfs/volumes/58c184c6-f088cbc4-e3df-e0db55117da4
100.67.169.43	Mounted	Connected	/vmfs/volumes/58c184c6-f088cbc4-e3df-e0db55117da4
100.67.169.41	Mounted	Connected	/vmfs/volumes/58c184c6-f088cbc4-e3df-e0db55117da4

Figure 34 Datastore mounted and connected

The path status to the LUN is verified from each host by selecting a host (from the list in Figure 34) and expanding the **Paths** item near the bottom of the window. In this example, each host has four active paths to the LUN as shown at the bottom of Figure 35.

Multipathing Details

Paths

Owner Plugin: NMP

Refresh Enable Disable

Runtime Name	Status	Target	LUN
vmhba33:C0:T2:L0	Active		0
vmhba34:C0:T0:L0	Active		0
vmhba33:C0:T0:L0	Active		0
vmhba34:C0:T3:L0	Active		0

Figure 35 Path status to LUN

The multipathing policy may be changed (e.g. to Round Robin) using the **Edit Multipathing** button.

8.7 CHAP Authentication

Unisphere provides Challenge Handshake Authentication Protocol (CHAP) for security measures. CHAP is a security protocol that defines a method for Authentication between iSCSI initiators and targets. The target authenticates a password from the initiator when the initiator wants to create a connection. If a valid password is not provided by the initiator, iSCSI connection will not be established. Types of CHAP authentication include one-way or mutual CHAP. In the one-way CHAP, the target seeks to authenticate the initiator. In mutual CHAP, both the iSCSI targets and the initiators seek to authenticate one another. Configuring CHAP Authentication is not covered in this deployment guide as it is beyond the scope of this document.

A FN IOM Internal Port Mapping

A.1 Quarter-Width Servers - Dual Port CNAs

The 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. Figure 36 and Table 7 present the port mapping for quarter-width servers with dual port CNAs.

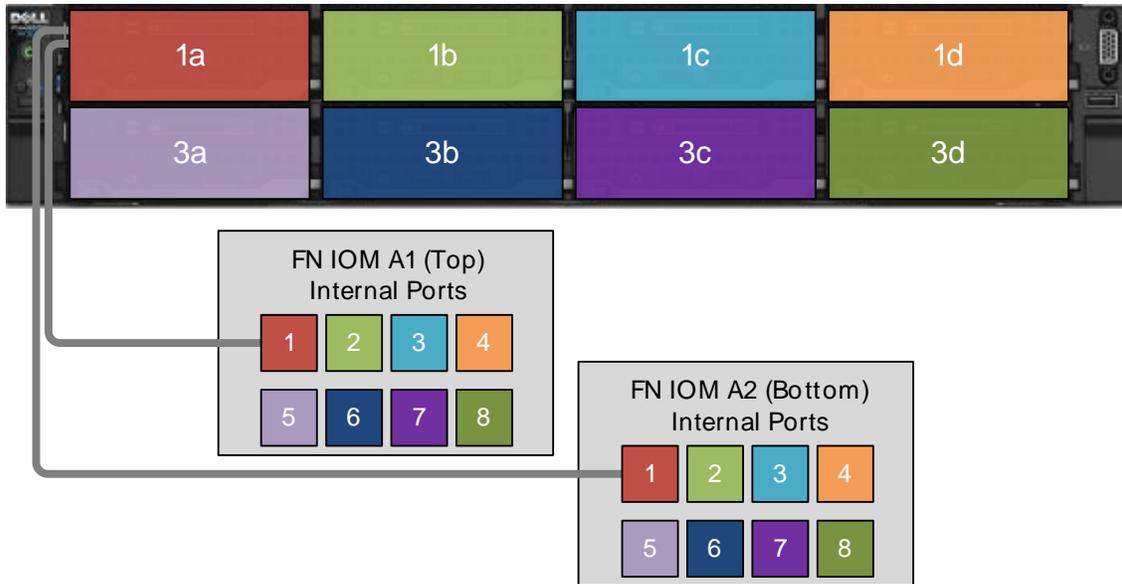


Figure 36 Quarter-width servers with dual port CNAs

Table 7 Quarter-width servers with dual port CNAs

Slot	FN IOM A1 (Top) Port Numbers	FN IOM A2 (Bottom) 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 - Dual Port CNAs

The 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 8 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.

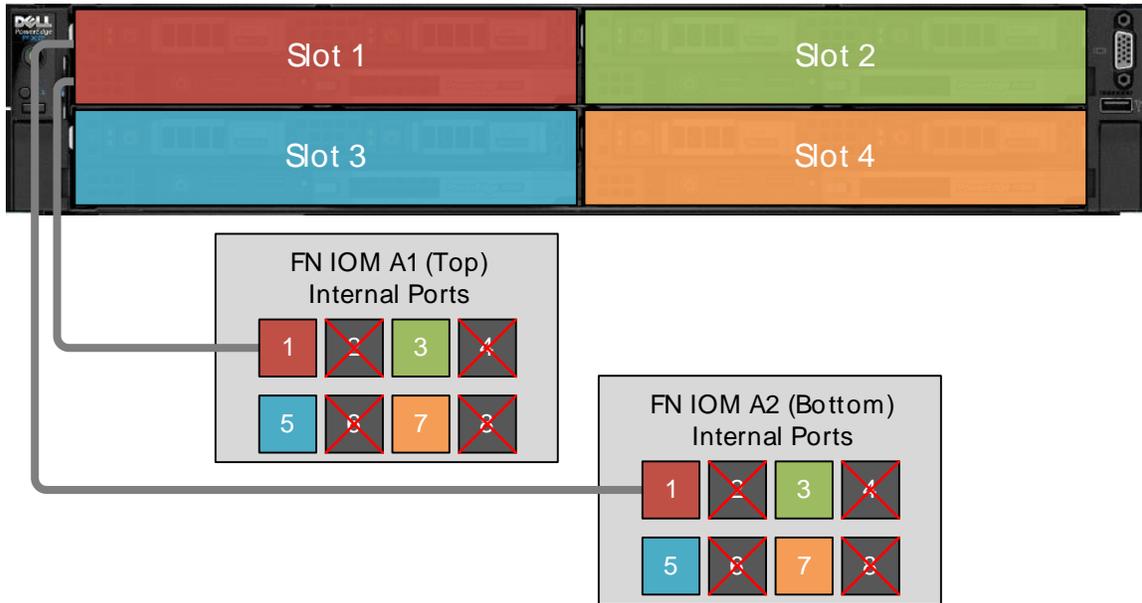


Figure 37 IOM Port Mapping half-width servers with dual port CNAs

Table 8 Half-width servers with dual port CNAs

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

A.3 Half-Width Servers - 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 9 present the port mapping for half-width servers with quad-port CNAs.

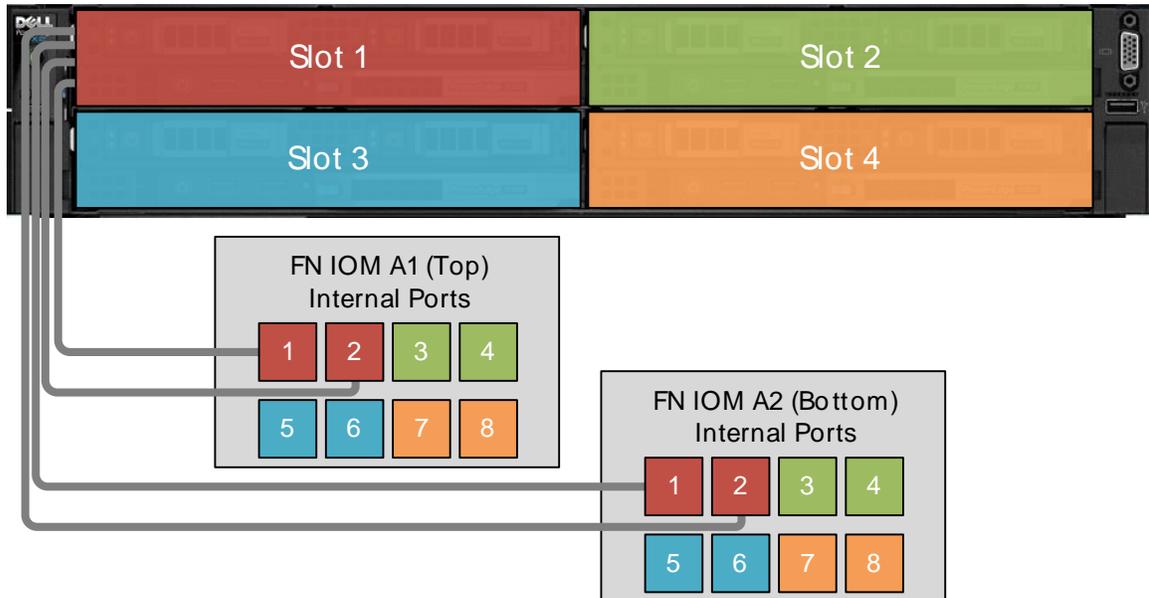


Figure 38 Half-width slots with quad port CNAs

Table 9 Half-width slots with quad port CNAs

Slot	FN IOM A1 (Top) Port Numbers	FN IOM A2 (Bottom) 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 - Dual Port CNAs

For full-width servers with dual-port CNAs installed, the CNA ports map to two ports on each FN IOM. Figure 14 and Table 4 present the port mapping for full-width servers with dual-port CNAs.

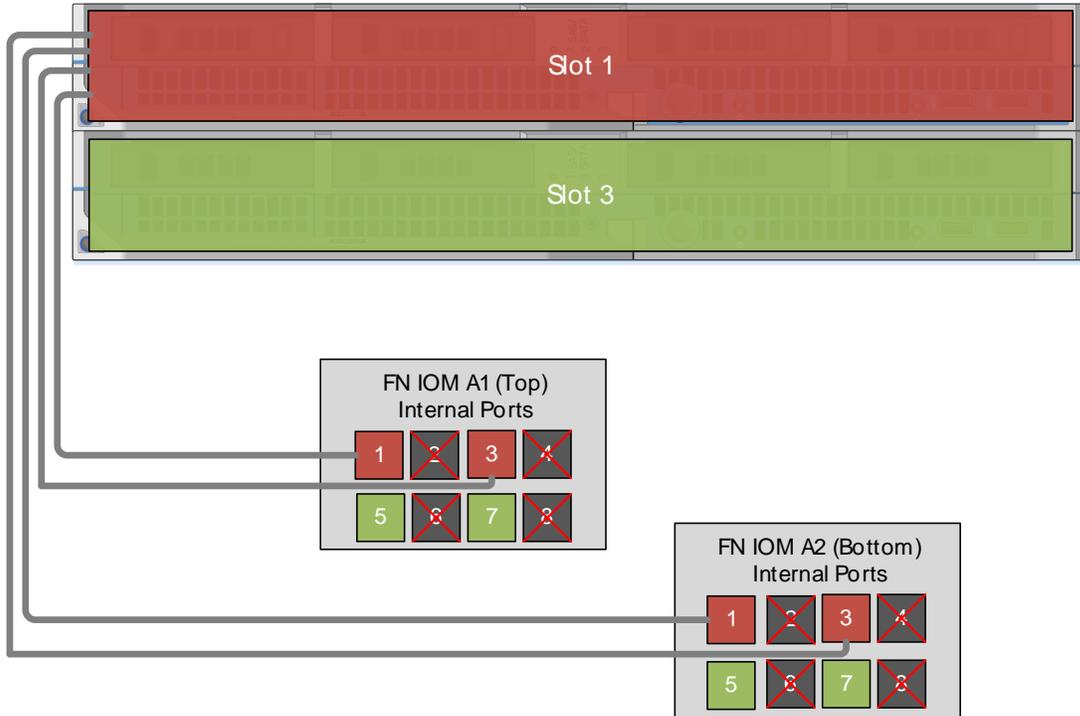


Figure 39 Full-width slots with dual port CNAs

Table 10 Full-width servers with dual port CNAs

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

A.5 Full-Width Servers - 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 11 present the port mapping for full-width servers with quad-port CNAs

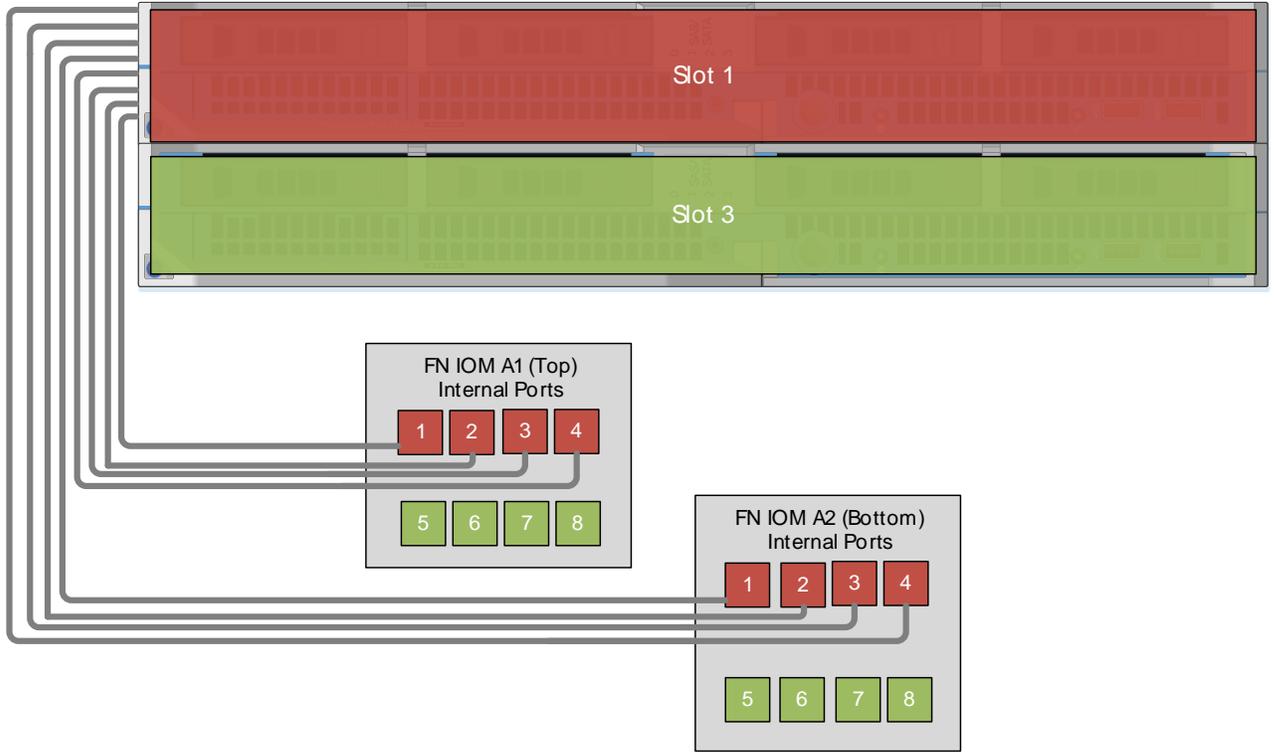


Figure 40 Full-width slots with quad port CNAs

Table 11 Full-width servers with quad port CNAs

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

B FN Series Operational Modes

The FN IOM supports five operational modes: Standalone (SMUX), Virtual Link Trunking (VLT), Stacking, Programmable MUX (PMUX) and Full-switch. Refer Table 12 for detailed descriptions of each mode. To enable a new operational mode issue the command `stack-unit 0 iom-mode IOM_Mode` in configuration mode. Administrators must reload the switch after enabling a new operational mode.

Table 12 FN series operational modes

IOM mode	Description
Standalone mode (SMUX)	This is the factory default mode for the IOM. SMUX is a fully automated, zero-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 provides operational flexibility by allowing 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. This mode simplifies management and redundant configurations.
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 present the hardware and components used to configure and validate the example configurations in this guide.

C.1 Switches

Qty	Item	Firmware Version
1	S6010-ON Spine switch	DNOS 9.11.0.0 P4
2	S4048-ON Leaf switch	DNOS 9.11.0.0 P4
2	S4048-ON iSCSI SAN switch	DNOS 9.11.0.0 P4
1	S3048-ON Management switch	DNOS 9.11.0.0 P4

C.2 PowerEdge FX2s chassis and components

This guide uses one FX2s chassis with four FC630 servers in the Compute cluster.

Qty per chassis	Item	Firmware Version
1	FX2s Chassis Management Controller	1.40.200.201608133205
2	FC630 servers. <ul style="list-style-type: none">• 2 - Intel Xeon E5-2683 v4 2.10GHz,16 Cores• 8 - 16GB DIMMS (128 GB total)• 2 - 200 GB SATA SSD• 2 - 16 GB Internal SD Cards• 1 - PERC H730P Slim Storage Controller• 2 - QLogic 577xx/578xx 10 Gb Ethernet BCM57810• FC630 BIOS• FC630 iDRAC with Lifecycle Controller	<ul style="list-style-type: none">• -• -• -• -• 25.5.0.0018• 08.07.26• 2.3.5• 2.41.40.40
1	FC630 server. <ul style="list-style-type: none">• 2 - Intel Xeon E5-2680 v3 2.50GHz,12 Cores• 8 - 16GB DIMMS (128 GB total)• 2 - 300 GB SAS HDD• 2 - 16 GB Internal SD Cards• 1 - PERC H330P Mini Storage Controller• 2 - QLogic 577xx/578xx 10 Gb Ethernet BCM57810• FC630 BIOS• FC630 iDRAC with Lifecycle Controller	<ul style="list-style-type: none">• -• -• -• -• 25.5.0.0019• 08.07.26• 2.3.5• 2.41.40.40
1	FC630 server.	

	<ul style="list-style-type: none"> • 2 - Intel Xeon E5-2683 v4 2.10GHz,16 Cores • 8 - 16GB DIMMS (128 GB total) • 3 - 200 GB SATA HDD • 2 - 16 GB Internal SD Cards • 1 - PERC H730P Slim Storage Controller • 2 - QLogic 577xx/578xx 10 Gb Ethernet BCM57810 • FC630 BIOS • FC630 iDRAC with Lifecycle Controller 	<ul style="list-style-type: none"> • - • - • - • - • 25.5.0.0018 • 08.07.26 • 2.3.5 • 2.41.40.40
2	FN410S IOM	DNOS 9.11.0.0P4
4	Intel I350-T Base-T 1GbE DP LP PCIe adapter (for Management)	17.5.10

C.3 Validated iSCSI Storage arrays

Dell EMC PowerEdge FX2s with FN410S modules were validated with the storage arrays shown below:

- Dell EMC Unity 500F
- Dell EMC Unity SC4020

Dell EMC Unity 500F delivers all-flash storage from 7.7 TB 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 two Storage Processors (SPs) and supports up to twenty-five 2.5-inch drives. Additional 2-RU Disk Array Enclosures (DAEs) may be added providing twenty-five additional drives each.

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



Figure 41 Dell EMC Unity Disk Processing Enclosure (DPE) (front)

Note: See [Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage](#) for more information on configuring SC4020.

D Dell EMC validated software and required licenses

The Software table presents the versions of the software components used to validate the example configurations in this guide. The Licenses section presents the licenses required for the example configurations in this this guide.

D.1 Software

Item	Version
VMware ESXi	6.5.0 Build 4564106 Dell Customized A00
VMware vCenter Server Appliance	VMware VCenter Server 6.5.0 Build 4602587
vSphere Web Client	6.5.0 Build 4602587 (included with VCSA above)

D.2 Licenses

The vCenter Server is licensed by instance. The remaining 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.com/support is focused on meeting customer needs with proven services and support.

[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 Networking S3048-ON](#)

[Manuals and documentation for Dell Networking S4048-ON](#)

[Manuals and documentation for Dell Networking S6010-ON](#)

[Manuals and documentation for the FN IO Module](#)

[Manuals and Documentation for PowerEdge FX2/FX2s](#)

[Dell TechCenter Networking Guides](#)

[Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage](#)

[FNIOM Easy deployment guide](#)

[DELL EMC Unity Technical Documentation](#)

[Dell EMC NSX Reference Architecture - FC630 Compute Nodes with iSCSI Storage.](#)

E.2 VMware product manuals and technical guides

[VMware vSphere 6.5 Documentation Center](#)

[VMware Compatibility Guide](#)

F Support and Feedback

Contacting Technical Support

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.