

Dell EMC Networking FCoE-to-Fibre Channel Deployment with S4148U-ON in F_port Mode

Connecting server CNAs to FC storage and a leaf-spine network using two S4148U-ON switches running OS10

Dell EMC Networking Infrastructure Solutions
June 2018

Revisions

Date	Rev.	Description	Authors
June 2018	1.0	Initial release	Jim Slaughter, Andrew Waranowski

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

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

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [6/5/2018] [Deployment Guide]

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

Table of contents

Revisions.....	2
1 Introduction.....	6
1.1 Typographical conventions.....	7
1.2 Attachments.....	7
2 Hardware overview.....	8
2.1 Dell EMC Networking S4148U-ON.....	8
2.2 Dell EMC Networking Z9100-ON.....	8
2.3 Dell EMC Networking S3048-ON	8
2.4 Dell EMC PowerEdge R640 server	9
2.5 Dell EMC Unity 500F storage array.....	9
3 Topology overview	10
3.1 FC SAN topology detail	11
3.2 LAN topology detail	12
3.3 OOB management network	13
4 Preparation	14
4.1 Reset server CNAs to factory defaults	14
4.2 Configure NIC partitioning on CNAs.....	14
4.3 Determine CNA FCoE port WWPNS	16
4.4 Determine Unity 500F storage array FC WWPNS.....	18
4.5 VMware preparation	19
4.5.1 VMware ESXi download and installation.....	19
4.5.2 Install and configure VMware vCenter Server 6.5 U1	19
4.5.3 Add ESXi hosts to vCenter Server	19
4.5.4 Add VMkernel adapters for FCoE.....	20
4.5.5 Increase MTU size for FCoE	21
4.5.6 Add software FCoE adapters to ESXi Hosts	22
5 S4148U-ON switch configuration	23
5.1 Prepare switches	23
5.1.1 Factory default configuration	23
5.1.2 Set switch port profile	23
5.2 Configure switches	24
5.2.1 Initial configuration settings	24

5.2.2	Fibre Channel configuration	24
5.2.3	QoS Configuration	27
5.2.4	LAN configuration for IP traffic.....	29
6	S4148U-ON validation.....	33
6.1	General commands	33
6.1.1	show interface status	33
6.1.2	show vlan.....	33
6.1.3	show vlt domain_id	34
6.1.4	show vlt domain_id backup link.....	34
6.2	FC commands	34
6.2.1	show fc switch.....	34
6.2.2	show fc ns switch.....	35
6.2.3	show fc zoneset active	36
6.2.4	show fcoe sessions	36
6.2.5	show vfabric.....	37
6.3	QoS commands	38
6.3.1	show qos maps.....	38
6.3.2	show policy-map	38
6.3.3	show class-map.....	39
6.3.4	show qos interface ethernet interface_number	39
6.4	DCBX commands	39
6.4.1	show lldp dcbx interface ethernet interface_number	39
6.4.2	show lldp dcbx interface ethernet interface_number pfc detail	41
6.4.3	show lldp dcbx interface ethernet interface_number ets detail	41
7	Configure Unity FC storage.....	43
7.1	Create a storage pool	43
7.2	Add ESXi hosts.....	44
7.3	Create LUNs and configure host access.....	44
8	Configure storage on ESXi hosts	46
8.1	Rescan storage.....	46
8.2	Create a datastore	47
8.3	Create a virtual disk.....	48
8.4	Configure the virtual disk	50

- 9 Configure ESXi hosts for LAN traffic51
 - 9.1 vSphere distributed switches51
 - 9.2 Create a VDS.....51
 - 9.3 Add a distributed port group52
 - 9.4 Configure load balancing53
 - 9.5 Add hosts to the VDS53
 - 9.6 Add a virtual network adapter to VMs55
 - 9.7 Verify connectivity to the production network57
- A Validated components58
 - A.1 Switches58
 - A.2 PowerEdge R640 servers.....58
 - A.3 Storage58
 - A.4 VMware software59
 - A.5 VMware licenses.....59
- B Technical support and resources60
 - B.1 Dell EMC product manuals and technical guides60
 - B.2 VMware product manuals and technical guides60
- C Support and feedback61

1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company in the data center for today's applications and for the cloud-native world we are entering. To attain that vision, the Dell EMC portfolio focuses on making every component of data center infrastructure (servers, storage, and networking) compelling by making the value of the integrated solution greater than the sum of the parts.

This guide provides a step-by-step deployment example of a Fibre Channel (FC) Storage Area Network (SAN) in a leaf-spine data center environment. Two Dell EMC Networking S4148U-ON switches running OS10 are used as a leaf pair with Ethernet ports connected to Z9100-ON spine switches for production TCP/IP traffic. The same leaf pair is also used in F_port mode for direct FC connections to a Dell EMC Unity 500F storage array and Fibre Channel over Ethernet (FCoE) connections to PowerEdge servers.



This guide includes configuration of the S4148U-ON leaf pair, VMware ESXi hosts, Converged Network Adapters (CNAs), and Unity FC storage. The goal of this guide is to enable a network administrator or deployment services engineer to deploy a pair of S4148U-ONs in a converged environment using the hardware and software described.

1.1 Typographical conventions

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

Monospace Text	CLI examples
<u>Underlined Monospace Text</u>	CLI examples that wrap the page. This text is entered as a single command.
<i>Italic Monospace Text</i>	Variables in CLI examples
Bold Monospace Text	Commands entered at the CLI prompt
Bold text	GUI fields and information entered in the GUI

1.2 Attachments

This .pdf includes switch configuration file attachments. To access attachments in Adobe Acrobat Reader, click the  icon in the left pane halfway down the page, then click the  icon.

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 A.

2.1 Dell EMC Networking S4148U-ON

The S4148U-ON enables converging LAN and SAN traffic in a single 1-RU, multilayer switch. It includes twenty-four 10GbE ports, two 40GbE ports, four 10/25/40/50/100GbE or FC8/16/32 ports, and twenty-four 10GbE or FC8/16 ports. Two S4148U-ONs are used as leaf switches in this guide for Ethernet connectivity to the leaf-spine network and Fibre Channel connectivity to the SAN.



Figure 1 Dell EMC Networking S4148U-ON

2.2 Dell EMC Networking Z9100-ON

The Z9100-ON is a 1-RU, multilayer switch with thirty-two ports supporting 10/25/40/50/100GbE plus two 10GbE SFP+ ports. Two Z9100-ON switches are used as spines in this guide.



Figure 2 Dell EMC Networking Z9100-ON

2.3 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 is used per rack for out-of-band (OOB) management traffic.



Figure 3 Dell EMC Networking S3048-ON

2.4 Dell EMC PowerEdge R640 server

The PowerEdge R640 is a 1-RU, two-socket server platform with support for up to 56 processor cores, 3TB memory, and up to twelve SAS/SATA HDD/SSD drives or eight NVMe drives. Two R640 servers are used in the deployment in this guide.



Figure 4 Dell EMC PowerEdge R640

2.5 Dell EMC Unity 500F storage array

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



Figure 5 Dell EMC Unity 500F front view



Figure 6 Dell EMC Unity 500F rear view

3 Topology overview

Note: In this deployment guide, "LAN" is used to broadly refer to the data center's leaf-spine production TCP/IP network, and "SAN" is used to refer to FCoE and FC storage networks.

A pair of S4148U-ONs is installed in a rack to forward converged LAN and SAN traffic for all devices in the rack. Each S4148U-ON provides universal ports that are configured as either FC or Ethernet. The mix of FC and Ethernet ports allows the switches to simultaneously connect directly to an FC SAN for storage traffic and act as leaves in a leaf-spine network for LAN traffic.

Note: FC SAN traffic does not traverse the leaf-spine network and is contained within the rack.

The combined LAN/SAN topology is shown in the diagram below.

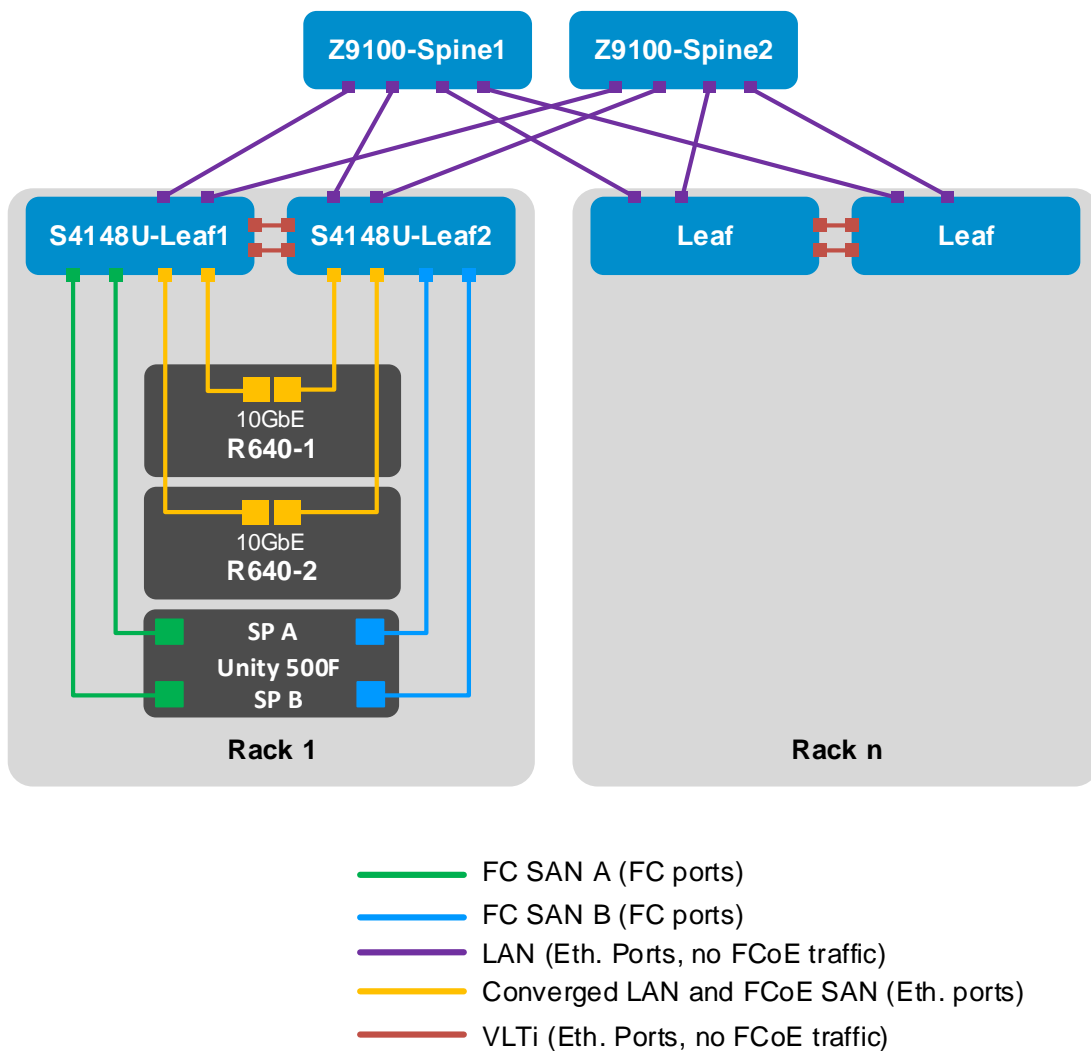


Figure 7 Combined LAN/SAN topology

Note: Using a leaf-spine network in the data center is considered a best practice. For detailed leaf-spine network configuration instructions, including spine switch configuration, refer to [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).

3.1 FC SAN topology detail

For the SAN portion of the topology, each S4148U-ON switch is placed in F_port mode. This enables direct connections to N_port devices such as FC storage without the need for an additional dedicated FC switch.

The FC SAN topology used in this guide is shown in Figure 8.

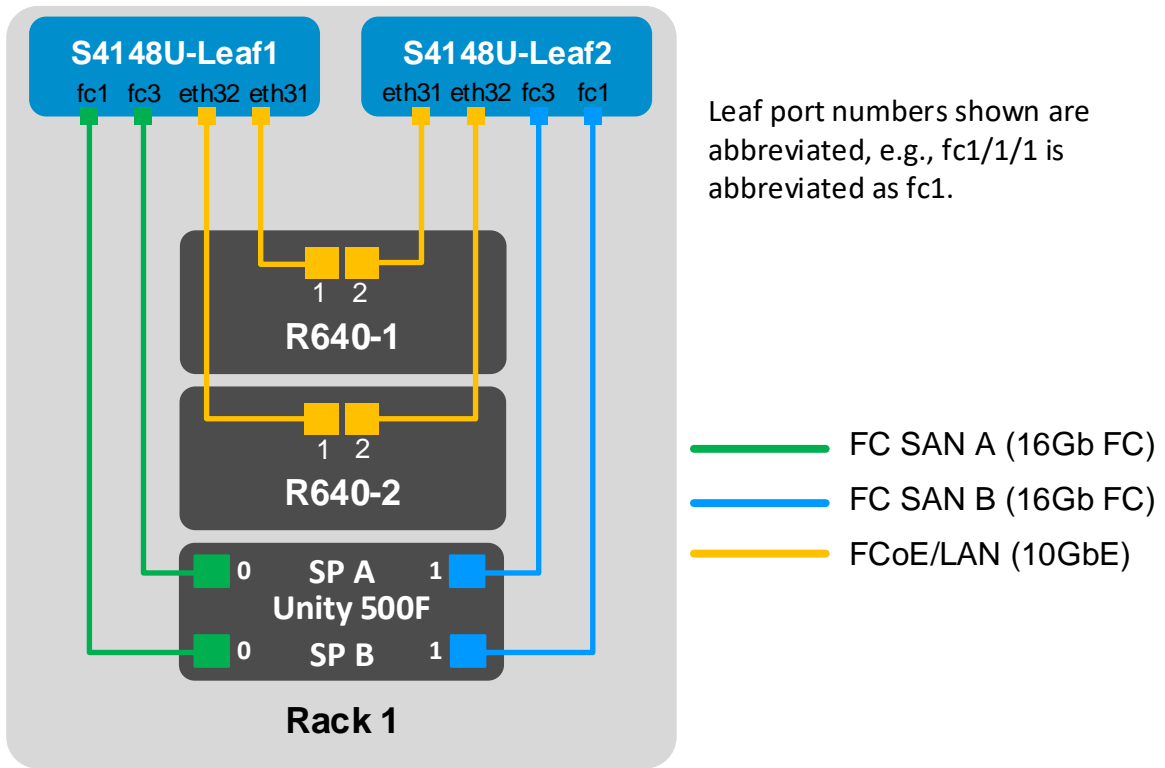


Figure 8 FC SAN topology

Each PowerEdge R640 server in this deployment contains one QLogic 57800 rack server network daughter card (rNDC) with two 10GbE CNA ports. These connections carry converged LAN and SAN (FCoE) traffic. Each CNA port is connected to an Ethernet port on an S4148U-ON as shown.

For SAN traffic, S4148U-ON switches serve as Fibre Channel Forwarders (FCFs). They decapsulate the FC payload from incoming FCoE frames and forward them to the storage array. Incoming FC frames from storage are encapsulated using FCoE and forwarded to the appropriate CNA.

The Unity 500F storage array contains redundant storage processors, SP A and SP B. FC ports from the Unity SPs are directly connected to FC ports on the S4148U-ON switches as shown.

3.2 LAN topology detail

For the LAN portion of the topology, the S4148U-ON switches are leaves in the data center's leaf-spine network. The topology includes the converged connections from the server CNAs and leaf-spine network connections as shown. The leaf pair forwards production TCP/IP traffic for all devices in the rack. Traffic destined for other racks is forwarded by the leaves to the spines.

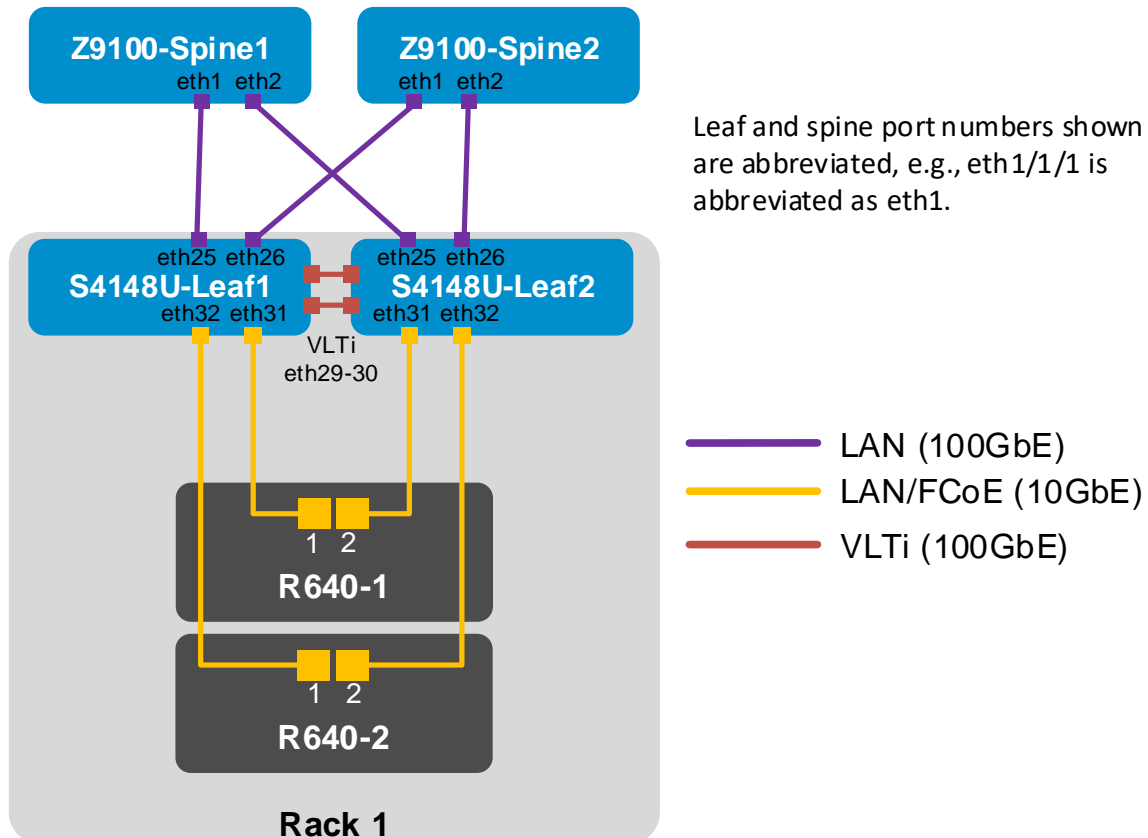


Figure 9 LAN topology

Each S4148U-ON leaf has one 100GbE connection to each Z9100-ON spine upstream. The downstream connections to the server CNAs carry converged FCoE and LAN traffic as covered in the previous section.

The S4148U-ON leaf switches are configured as a Virtual Link Trunking (VLT) pair, with the VLT interconnect (VLTi) also using 100GbE connections.

Note: VLT is used for LAN traffic only, and FCoE traffic is not forwarded across the VLTi. When the FCoE VLAN is added to the vfabri, it is automatically removed from the VLTi. This can be confirmed with the `show vlan` command.

3.3 OOB management network

The out-of-band (OOB) management network is an isolated network for management traffic only. It is used by administrators to remotely configure and manage servers, switches, and storage devices. Production traffic initiated by the network end users does not traverse the management network.

An S3048-ON switch is installed at the top of each rack for OOB management connections as shown.

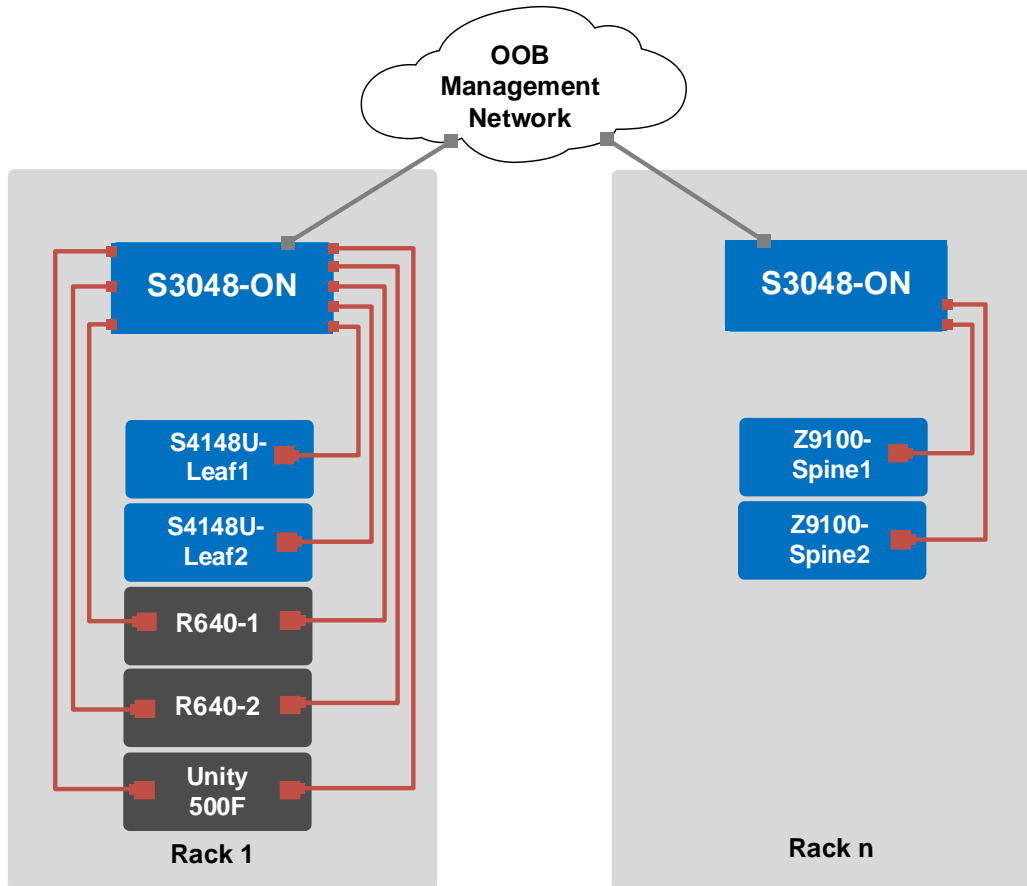


Figure 10 OOB management network connections

Four 10GbE SFP+ ports are available on each S3048-ON switch for use as uplinks to the OOB management network core. Downstream connections to servers, switches, and storage devices are 1GbE BASE-T.

The dedicated OOB management port of each leaf and spine switch is used for these connections. Each PowerEdge R640 server has two connections to the S3048-ON: the server's iDRAC port and a 1GbE NIC that is used for ESXi host management.

Note: For this deployment example, each server's ESXi management connection is made using one of the two QLogic 57800 1GbE BASE-T ports on the rNDC.

The Unity 500F storage array has two dedicated management ports: one for SP A and one for SP B. Both are connected to the S3048-ON.

4 Preparation

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

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 CNA port.
6. From the **Main Configuration Page**, click the **Default** button followed by **Yes** to load the default settings. Click **OK**.
7. Click **Finish**. Notice if the message indicates a reboot is required for changes to take effect.
8. Click **Yes** to save changes. Click **OK**.
9. Repeat for each CNA port listed on the **Device Settings** page.
10. If required per step 7, reboot the system and return to System Setup to configure NIC partitioning.

4.2 Configure NIC partitioning on CNAs

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

Note: This is only done on CNA ports that carry converged traffic. In this example, these are the two 10GbE QLogic 578xx CNA ports on each server that are connected to the leaf switches. It is not done on the 1GbE port connected to the OOB management network.

If the system is 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 the first CNA port.
6. Select **Device Level Configuration**. Set **Virtualization Mode** to **NPar** (if not already set) 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 **iSCSI Offload Mode** to **Disabled**.
- c. Set **FCoE Mode** to **Enabled**, shown in Figure 11, and click **Back**.

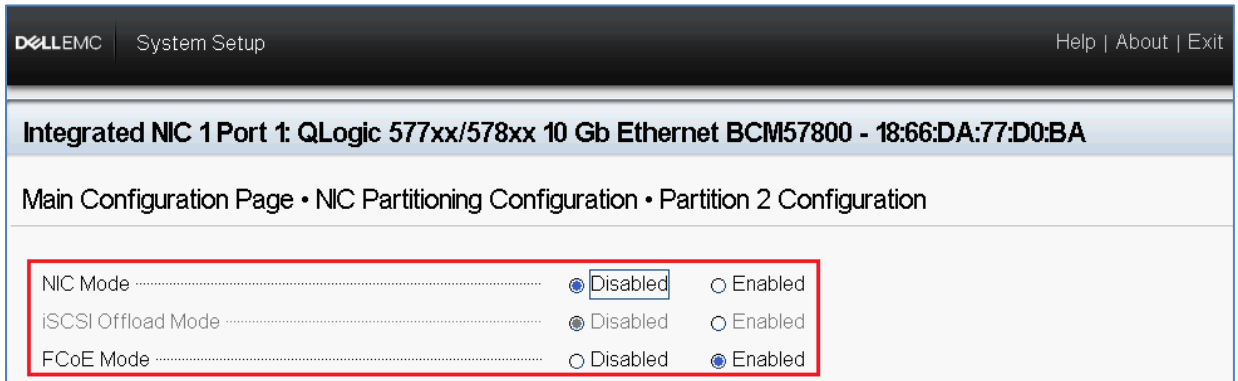


Figure 11 CNA partition 2 configuration

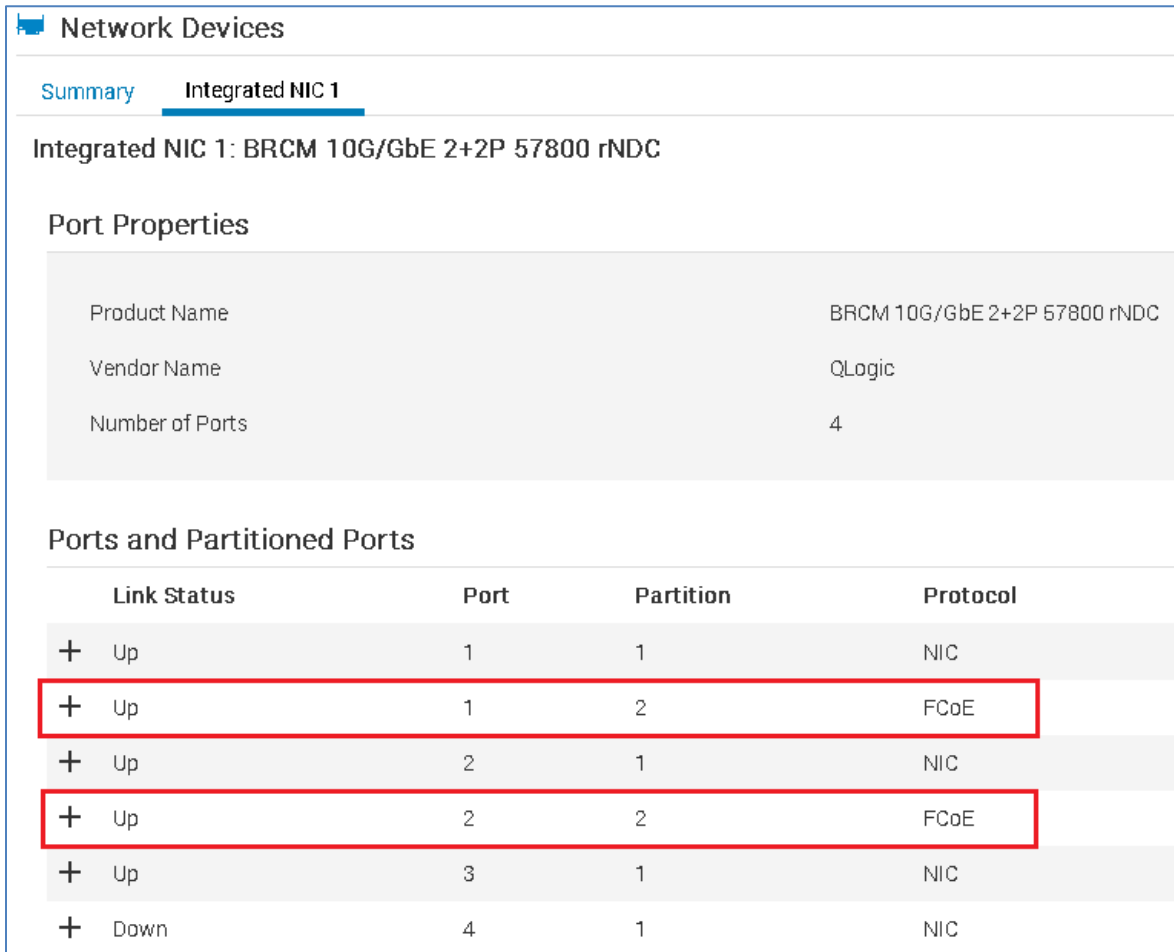
10. If present, select **Partition 3 Configuration**. Set NIC, iSCSI and FCoE modes to **Disabled** and click **Back**.
11. If present, select **Partition 4 Configuration**. Set NIC, iSCSI and FCoE modes to **Disabled** and click **Back**.
12. Click **Back > Finish**.
13. When prompted, answer **Yes**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

Repeat on remaining servers as needed.

4.3 Determine CNA FCoE port WWPNs

The PowerEdge R640 server's FCoE adapter World Wide Port Names (WWPNs) are used for FC zone configuration. Adapter WWPNs are determined as follows:

1. Connect to the first server's iDRAC in a web browser and log in.
2. Select **System > Network Devices**.
3. Click on the CNA. In this example, it is **Integrated NIC 1**. Under **Ports and Partitioned Ports**, the FCoE partition for each port is displayed as shown in Figure 12:



Network Devices

Summary **Integrated NIC 1**

Integrated NIC 1: BCM 10G/GbE 2+2P 57800 rNDC

Port Properties

Product Name	BCM 10G/GbE 2+2P 57800 rNDC
Vendor Name	QLogic
Number of Ports	4

Ports and Partitioned Ports

Link Status	Port	Partition	Protocol
+ Up	1	1	NIC
+ Up	1	2	FCoE
+ Up	2	1	NIC
+ Up	2	2	FCoE
+ Up	3	1	NIC
+ Down	4	1	NIC

Figure 12 FCoE partitions in iDRAC

4. The first FCoE partition is **Port 1, Partition 2**. Click the **+** icon to view the MAC Addresses as shown in Figure 13.

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

Figure 13 MAC address and FCoE WWPN for CNA port 1

- Record the **MAC Address** and **WWPN**, outlined in red above. A convenient method is to copy and paste these values into a text file.

Note: While the WWPN is used in the actual switch configuration for FC zoning, the MAC address is recorded to identify the corresponding vmnic number in VMware.

- Repeat steps 4 and 5 for the FCoE partition on port 2.
- Repeat steps 1-6 for remaining servers.


The FCoE WWPNs and MAC addresses used in this deployment example are shown in Table 1. The **Switch** column has been added for reference per the cable connections in the SAN topology diagram (Figure 8).

Table 1 Server CNA FCoE port WWPNs and MACs

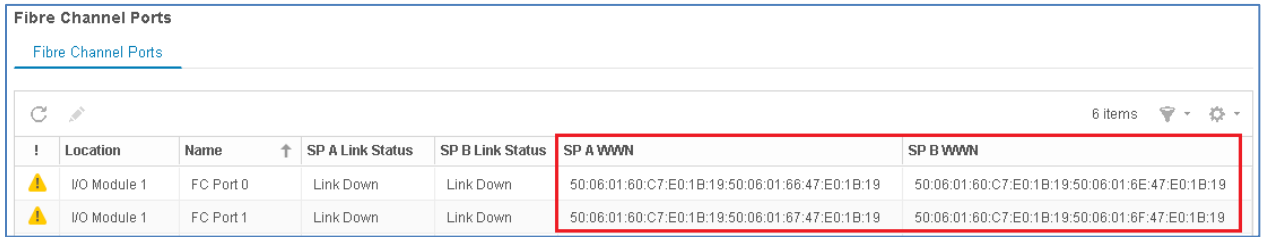
Server	Port	Switch	WWPN	MAC
R640-1	1	S4148U-Leaf1	20:01:18:66:DA:71:50:AD	18:66:DA:71:50:AC
R640-1	2	S4148U-Leaf2	20:01:18:66:DA:71:50:AF	18:66:DA:71:50:AE
R640-2	1	S4148U-Leaf1	20:01:18:66:DA:77:D0:C3	18:66:DA:77:D0:C2
R640-2	2	S4148U-Leaf2	20:01:18:66:DA:77:D0:C5	18:66:DA:77:D0:C4

4.4 Determine Unity 500F storage array FC WWPNs

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

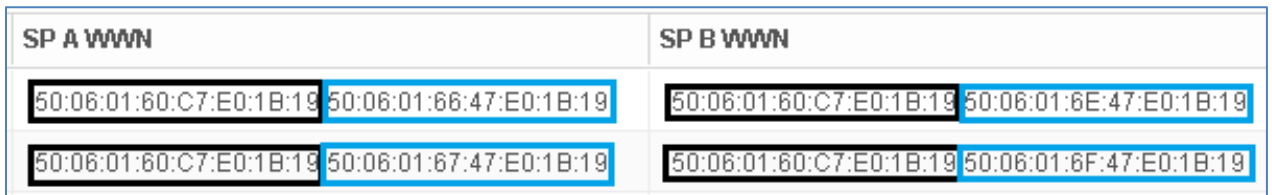
1. Connect to the Unisphere GUI in a web browser and log in. Click the **Settings** icon  near the top right corner of the page.
2. In the left pane of the **Settings** window, select **Access > Fibre Channel**.

The **Fibre Channel Ports** page is displayed as shown in Figure 14. A zoomed-in view of the area inside the red box is shown in Figure 15.



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

Figure 14 Unisphere Fibre Channel Ports page



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

Figure 15 Zoomed-in view of SP A and SP B WWPNs on Fibre Channel Ports page

3. Two WWPNs are listed for each port. The World Wide Node Name (WWNN), outlined in black, identifies this Unity storage array (the node). It is not used in zone configuration. The WWPNs, outlined in blue, identify the individual ports and are used for FC zoning. Record the WWPNs as shown in Table 2. The **Switch** column has been added based on the physical cable connections shown in Figure 8.

Table 2 Storage array FC adapter WWPNs

Service processor	Port	Switch	WWPN
SP A	0	S4148U-Leaf1	50:06:01:66:47:E0:1B:19
SP A	1	S4148U-Leaf2	50:06:01:67:47:E0:1B:19
SP B	0	S4148U-Leaf1	50:06:01:6E:47:E0:1B:19
SP B	1	S4148U-Leaf2	50:06:01:6F:47:E0:1B:19

4.5 VMware preparation

4.5.1 VMware ESXi download and installation

Install VMware ESXi 6.5 U1 or later on each PowerEdge server. Dell EMC recommends using the latest Dell EMC customized ESXi .iso image available on support.dell.com. The correct drivers for the PowerEdge server 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.

Each ESXi host has one 1GbE BASE-T interface connected to the OOB management network for communication with the vCenter Server.

4.5.2 Install and configure VMware vCenter Server 6.5 U1

For information on the installation and configuration of vCenter Server, refer to [vSphere Installation and Setup](#). In this deployment, the vCenter Server communicates with ESXi hosts via the OOB management network, and it may be located on any host accessible over this network.

4.5.3 Add ESXi hosts to vCenter Server

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 two ESXi hosts are added to it.

A cluster named **Rack1** is created and added to the datacenter object, and the two R640 hosts are added to the cluster. When complete, the vSphere Web Client **Hosts and Clusters** tab in the **Navigator** pane appears as shown in Figure 16.








Figure 16 Datacenter and cluster created with ESXi hosts

4.5.4 Add VMkernel adapters for FCoE

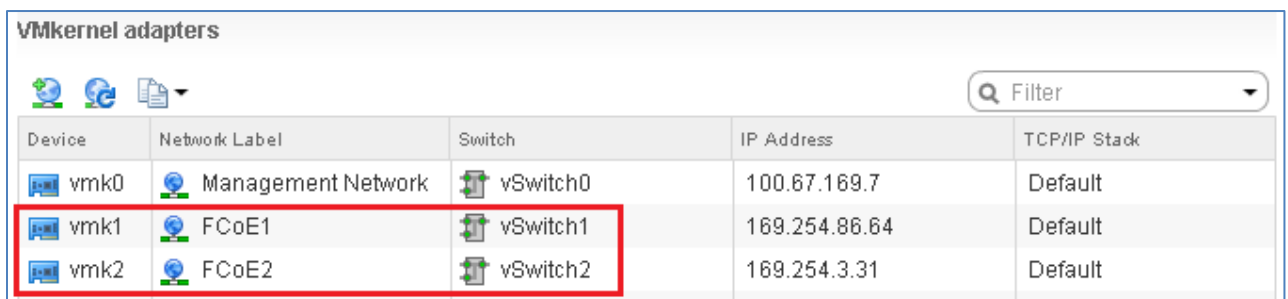
Note: Before starting this section, be sure the vmnic-to-physical adapter mapping for each host's FCoE ports is known. Vmnics and their MAC addresses are visible in the Web Client by selecting the host in the **Navigator** pane. In the center pane, go to **Configure > Networking > Physical adapters**. MAC addresses for this deployment were recorded earlier in Table 1. In this example, the two FCoE adapters on each host are vmnic4 and vmnic5. 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, e.g., **r640-1.dell.local**. In the center pane, select **Configure > Networking > Virtual switches**.
3. For the first CNA port, 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., vmnic4, 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**, and the **VLAN ID**, e.g., **1002**. (VLAN 1002 is used for FCoE traffic on both S4148U switches in this guide). 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 CNA port, 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., vmnic5. 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 using the same VLAN ID.

When complete, the **Configure > Networking > VMkernel adapters** page for each host appear should appear similar to Figure 17.



Device	Network Label	Switch	IP Address	TCP/IP Stack
vmk0	Management Network	vSwitch0	100.67.169.7	Default
vmk1	FCoE1	vSwitch1	169.254.86.64	Default
vmk2	FCoE2	vSwitch2	169.254.3.31	Default


Figure 17 Host VMkernel adapters configured

Note: The IP addresses shown next to the two FCoE adapters are automatically assigned private addresses and are not used for FCoE.

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.

Use the following steps to 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 18.

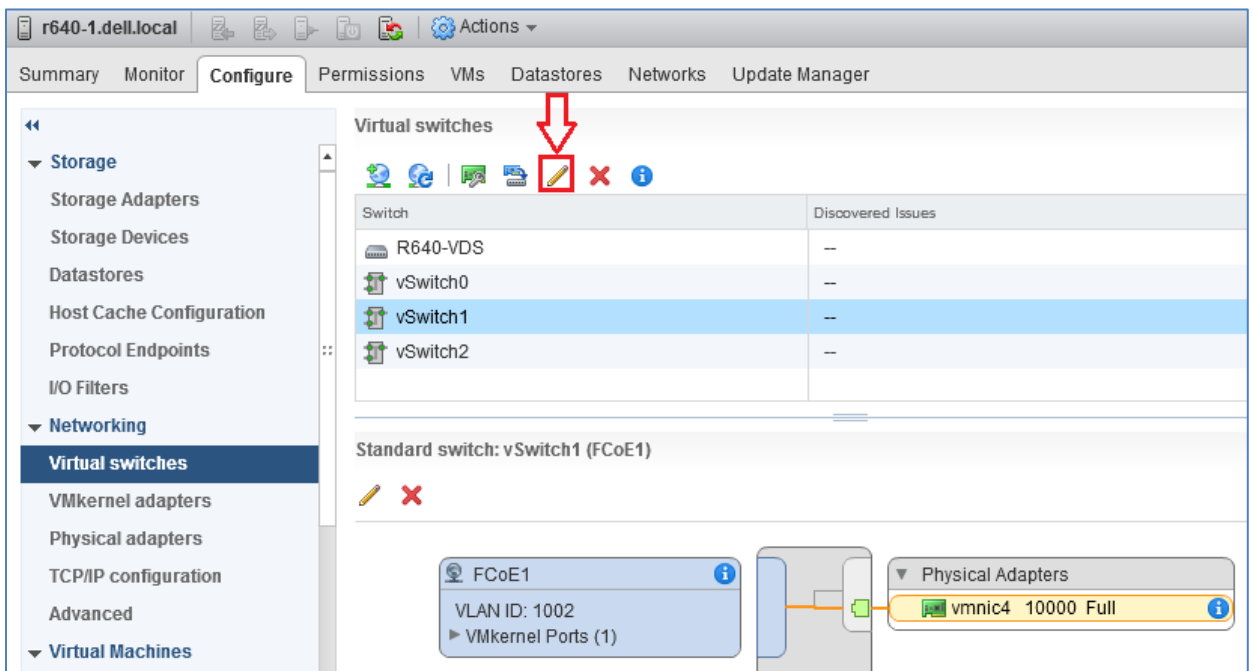


Figure 18 Selecting the correct vSwitch edit icon

4. In the **Edit Settings** dialog box, change the value of **MTU (Bytes)** to 2500 as shown in Figure 19.

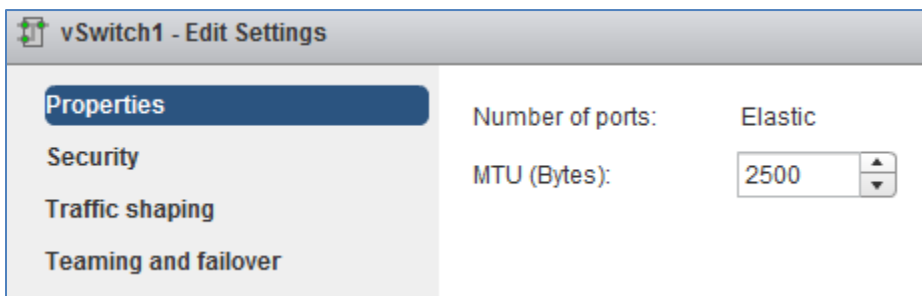



Figure 19 vSwitch MTU set to 2500 bytes



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.

Use the following steps to 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. Select 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. Leave the settings at their default values and click **OK** to create the FCoE adapter.
5. Click the  icon to refresh the storage adapter view. The newly added FCoE adapter appears under **QLogic 57800 10 Gigabit Ethernet Multi Function Adapter**.
6. Repeat steps 3-5 to add the second FCoE adapter.

When complete, two FCoE storage adapters, named vmhba32 and vmhba33 in this example, are listed:

Storage Adapters						
Adapter	Type	Status	Identifier	Targets	Devices	Paths
Lewisburg SATA AHCI Controller						
vmhba1	Block SCSI	Unknown		0	0	0
vmhba2	Block SCSI	Unknown		0	0	0
PERC H330 Mini						
vmhba0	SAS	Unknown	5d09466020bede00	1	1	1
QLogic 57800 10 Gigabit Ethernet Multi Function Adapter						
vmhba32	Fibre Channel over Ethernet	Offline	20:00:18:66:da:71:50:ad 20:01:18:66:da:71:50:ad	0	0	0
vmhba33	Fibre Channel over Ethernet	Offline	20:00:18:66:da:71:50:af 20:01:18:66:da:71:50:af	0	0	0

Figure 20 FCoE storage adapters created on ESXi host

Repeat the steps above for the remaining ESXi hosts.

5 S4148U-ON switch configuration

This section covers steps to configure the S4148U-ON leaf switches running OS10.

5.1 Prepare switches

5.1.1 Factory default configuration

The configuration commands in the sections that follow begin with S4148U-ON switches at their factory default settings. Dell EMC Networking switches running OS10 can be reset to their default configuration as follows:

```
OS10# delete startup-configuration  
Proceed to delete startup-configuration [confirm yes/no(default)]:y
```

```
OS10# reload  
System configuration has been modified. Save? [yes/no]:n  
Proceed to reboot the system? [confirm yes/no]:y
```

The switch reboots to its factory default configuration. Repeat on the second S4148U-ON switch.

Note: Dell EMC Networking OS10 switches at their default settings have Telnet disabled, SSH enabled, and the OOB management interface configured to get its IP address via DHCP. The default username and password are both `admin`. Dell EMC recommends changing the admin password to a complex password during the first login.

5.1.2 Set switch port profile

A switch port profile determines the enabled front-panel ports and supported breakout modes on Ethernet and unified ports. Changing the profile removes the existing switch configuration, so the switch port profile is configured first. Switch port profile-3 is used in this topology to enable four 100GbE ports on each leaf switch. These ports are used for spine switch and VLTi connections.

Note: See [OS10 Enterprise Edition User Guide Release 10.4.0E\(R3\)](#) for switch-port profile options and details.

To change the switch-port profile from its current setting:

```
OS10# configure terminal  
OS10(config)# switch-port-profile 1/1 profile-3  
Warning: Switch port profile will be applied only after a save and reload.  
All management port configurations will be retained but all other  
configurations will be wiped out after the reload.  
OS10(config)# exit  
OS10# write memory  
OS10# reload
```

```
Proceed to reboot the system? [confirm yes/no]:y
```

5.2 Configure switches

After both S4148U-ONs are set to port profile-3, the commands in the tables that follow are run to complete the configuration of both switches. The port numbers used correspond to the topology diagrams shown in Figure 8 and Figure 9.

Note: The commands in the tables below should be entered in the order shown. Complete switch running-configuration files are provided as attachments named **S4148U-Leaf1.txt** and **S4148U-Leaf2.txt**.

5.2.1 Initial configuration settings

Configure the hostname, OOB management IP address, and OOB management default gateway. Set the Spanning Tree Protocol mode to RSTP. Leaf1 is configured as the primary RSTP root bridge with the `priority 0` command option, and Leaf2 is configured as the secondary RSTP root bridge with the `priority 4096` option.

S4148U-Leaf1	S4148U-Leaf2
<pre>configure terminal hostname S4148U-Leaf1 interface mgmt 1/1/1 no ip address dhcp ip address 100.67.169.31/24 no shutdown management route 0.0.0.0/0 100.67.169.254 spanning-tree mode rstp spanning-tree rstp priority 0</pre>	<pre>configure terminal hostname S4148U-Leaf2 interface mgmt 1/1/1 no ip address dhcp ip address 100.67.169.30/24 no shutdown management route 0.0.0.0/0 100.67.169.254 spanning-tree mode rstp spanning-tree rstp priority 4096</pre>

5.2.2 Fibre Channel configuration

Globally enable DCBX and F_Port mode on each switch. The `feature fc domain-id domain-id` command puts the switch in F_port mode.

S4148U-Leaf1	S4148U-Leaf2
<pre>dcbx enable feature fc domain-id 100</pre>	<pre>dcbx enable feature fc domain-id 100</pre>

Configure FC aliases for CNA and storage WWPNs. Using aliases is optional, but makes the configuration more user-friendly.

S4148U-Leaf1	S4148U-Leaf2
<pre>fc alias r640-1p1 member wwn 20:01:18:66:DA:71:50:AD fc alias r640-2p1 member wwn 20:01:18:66:DA:77:D0:C3 fc alias SpB-0 member wwn 50:06:01:6E:47:E0:1B:19 fc alias SpA-0 member wwn 50:06:01:66:47:E0:1B:19</pre>	<pre>fc alias r640-1p2 member wwn 20:01:18:66:DA:71:50:AF fc alias r640-2p2 member wwn 20:01:18:66:DA:77:D0:C5 fc alias SpB-1 member wwn 50:06:01:6F:47:E0:1B:19 fc alias SpA-1 member wwn 50:06:01:67:47:E0:1B:19</pre>

CNA and storage adapter WWPNs are combined into zones to allow communication between devices in the same zone. Dell EMC recommends single-initiator zoning. In other words, no more than one server CNA port per zone. For high availability, each CNA port should be zoned to at least one port from SP A and one port from SP B.

In this deployment, one zone is created for each CNA port. The zone contains the CNA port and the two storage processor ports connected to the same S4148U-ON.

S4148U-Leaf1	S4148U-Leaf2
<pre>fc zone r640-1p1zone member alias-name r640-1p1 member alias-name SpB-0 member alias-name SpA-0 fc zone r640-2p1zone member alias-name r640-2p1 member alias-name SpB-0 member alias-name SpA-0</pre>	<pre>fc zone r640-1p2zone member alias-name r640-1p2 member alias-name SpB-1 member alias-name SpA-1 fc zone r640-2p2zone member alias-name r640-2p2 member alias-name SpB-1 member alias-name SpA-1</pre>

A zone set is a collection of zones. A zone set named zoneset1 is created on each switch, and the zones are added to it.

Note: Multiple zone sets may be configured, but only one zone set can be active at any time.

S4148U-Leaf1	S4148U-Leaf2
<pre>fc zoneset zoneset1 member r640-1p1zone member r640-2p1zone exit</pre>	<pre>fc zoneset zoneset1 member r640-1p2zone member r640-2p2zone exit</pre>

Create the FCoE VLAN, vfabric, FCoE map, and activate the zone set. In this deployment, the FCoE VLAN is set to 1002, the vfabric ID is set to 100 (valid range is 1-255), and the fcmmap is set to 0xEFC64 (valid range is 0xEFC00-0xEFCFF).

S4148U-Leaf1	S4148U-Leaf2
<pre>interface vlan 1002 vfabric 100 vlan 1002 fcoe fcmmap 0xEFC64 zoneset activate zoneset1</pre>	<pre>interface vlan 1002 vfabric 100 vlan 1002 fcoe fcmmap 0xEFC64 zoneset activate zoneset1</pre>

Note: If changes are subsequently made to aliases, zones, or zone sets, the zone set must be reactivated for the changes to take effect. This is done with the following command sequence:

```
S4148U-Leaf1# configure terminal
S4148U-Leaf1(config)# vfabric vfabric_number
S4148U-Leaf1(conf-vfabric-100)# zoneset activate zoneset_name
```

Put a sufficient number of unified ports in 16Gb FC mode for connections to the FC storage array. In this deployment example, two FC ports per S4148U-ON switch are used. In switch port profile-3, unified port group 1 (entered as 1/1/1) provides two 16Gb FC interfaces, 1/1/1 and 1/1/3. Ports 1/1/2 and 1/1/4 are disabled.

Note: For more information on mapping ports to switch profiles and port groups, see [OS10 Enterprise Edition User Guide Release 10.4.0E\(R3\)](#), p.83-85.

Add the Fibre Channel and Ethernet interfaces to vfabric 100. This command adds the Ethernet interfaces as tagged ports in the FCoE VLAN. The Ethernet interfaces are also added as access (untagged) ports on VLAN 1 for the initial Fibre Channel Initialization Protocol (FIP) VLAN request. Using trunk mode enables multiple tagged VLANs and the access VLAN to traverse the same link. Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS) are enabled for lossless storage traffic and bandwidth management, respectively.

S4148U-Leaf1	S4148U-Leaf2
<pre>port-group 1/1/1 mode fc 16g-2x interface fibrechannel 1/1/1 description "To SPB" no shutdown vfabric 100 interface fibrechannel 1/1/3 description "To SPA" no shutdown</pre>	<pre>port-group 1/1/1 mode fc 16g-2x interface fibrechannel 1/1/1 description "To SPB" no shutdown vfabric 100 interface fibrechannel 1/1/3 description "To SPA" no shutdown</pre>

S4148U-Leaf1	S4148U-Leaf2
<pre> vfabric 100 interface ethernet 1/1/31 description "To Server 1" switchport access vlan 1 switchport mode trunk spanning-tree port type edge priority-flow-control mode on ets mode on vfabric 100 no shutdown interface ethernet 1/1/32 description "To Server 2" switchport access vlan 1 switchport mode trunk spanning-tree port type edge priority-flow-control mode on ets mode on vfabric 100 no shutdown </pre>	<pre> vfabric 100 interface ethernet 1/1/31 description "To Server 1" switchport access vlan 1 switchport mode trunk spanning-tree port type edge priority-flow-control mode on ets mode on vfabric 100 no shutdown interface ethernet 1/1/32 description "To Server 2" switchport access vlan 1 switchport mode trunk spanning-tree port type edge priority-flow-control mode on ets mode on vfabric 100 no shutdown </pre>

Note: On Ethernet ports, the MTU is automatically set to 2500 bytes when the port is added to the vfabric. This is verified by running the command `show interface ethernet interface_number` after configuration.

5.2.3 QoS Configuration

Quality of Service (QoS) configuration is a 3-step process:

1. Create class maps to classify traffic.
2. Create QoS and policy maps for the classified traffic.
3. Apply the QoS and policy maps.

FCoE traffic is assigned dot1p priority value 3 by default. In the following tables, dot1p priority value 3 is mapped to QoS group 3. The remaining dot1p priority values, 0-2 and 4-7, are mapped to QoS group 0. QoS group 3 is mapped to queue 3, and QoS group 0 is mapped to queue 0. In this example, bandwidth is allocated 50% to queue 3 for FCoE traffic and 50% to queue 0 for all other traffic using the `bandwidth percent` command. (Bandwidth allocations should total 100% and may be adjusted accordingly to optimize traffic).

Note: For more information, see the *Quality of service* and *Converged data center services* sections of [OS10 Enterprise Edition User Guide Release 10.4.0E\(R3\)](#).

The following commands are run to create the class maps:

S4148U-Leaf1	S4148U-Leaf2
<pre>class-map type network-qos class_Dot1p_3 match qos-group 3 class-map type queuing map_ETSQueue_0 match queue 0 class-map type queuing map_ETSQueue_3 match queue 3</pre>	<pre>class-map type network-qos class_Dot1p_3 match qos-group 3 class-map type queuing map_ETSQueue_0 match queue 0 class-map type queuing map_ETSQueue_3 match queue 3</pre>

QoS and policy maps are configured as follows:

S4148U-Leaf1	S4148U-Leaf2
<pre>trust dot1p-map map_Dot1pToGroups qos-group 0 dot1p 0-2,4-7 qos-group 3 dot1p 3 qos-map traffic-class map_GroupsToQueues queue 0 qos-group 0 queue 3 qos-group 3 policy-map policy_Input_TrustDot1p class class-trust trust dot1p policy-map type network-qos policy_Input_PFC class class_Dot1p_3 pause pfc-cos 3 policy-map type queuing policy_Output_BandwidthPercent class map_ETSQueue_0 bandwidth percent 50 class map_ETSQueue_3 bandwidth percent 50</pre>	<pre>trust dot1p-map map_Dot1pToGroups qos-group 0 dot1p 0-2,4-7 qos-group 3 dot1p 3 qos-map traffic-class map_GroupsToQueues queue 0 qos-group 0 queue 3 qos-group 3 policy-map policy_Input_TrustDot1p class class-trust trust dot1p policy-map type network-qos policy_Input_PFC class class_Dot1p_3 pause pfc-cos 3 policy-map type queuing policy_Output_BandwidthPercent class map_ETSQueue_0 bandwidth percent 50 class map_ETSQueue_3 bandwidth percent 50</pre>

The QoS and policy maps defined above are applied using the `system qos` command as follows:

S4148U-Leaf1	S4148U-Leaf2
<pre>system qos trust-map dot1p map_Dot1pToGroups qos-map traffic-class map_GroupsToQueues service-policy input type qos policy_Input_TrustDot1p service-policy input type network-qos policy_Input_PFC service-policy output type queuing policy_Output_BandwidthPercent</pre>	<pre>system qos trust-map dot1p map_Dot1pToGroups qos-map traffic-class map_GroupsToQueues service-policy input type qos policy_Input_TrustDot1p service-policy input type network-qos policy_Input_PFC service-policy output type queuing policy_Output_BandwidthPercent</pre>

5.2.4 LAN configuration for IP traffic

The final configuration steps are for the LAN side of the topology which includes connections to the spine switches.

Note: For more information on the configuration settings in this section, see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#). Spine switch configuration examples are also provided.

Configure the S4148U leaf switches as VLT peers for LAN traffic. 100GbE interfaces 1/1/29-1/1/30 are used for the VLTi in this deployment. The backup destination is the management IP address of the VLT peer switch. Enable peer routing.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface range ethernet 1/1/29-1/1/30 description VLTi no switchport vlt-domain 127 backup destination 100.67.169.30 discovery-interface ethernet 1/1/29 discovery-interface ethernet 1/1/30 peer-routing</pre>	<pre>interface range ethernet 1/1/29-1/1/30 description VLTi no switchport vlt-domain 127 backup destination 100.67.169.31 discovery-interface ethernet 1/1/29 discovery-interface ethernet 1/1/30 peer-routing</pre>

Create a server-facing VLAN interface for IP traffic. Use the same VLAN ID, VLAN 50 in this example, on both leaf switches. Assign an IP address to the VLAN interface. The address must be unique but on the same network on both leaf switches. Configure VRRP to use VRRP version 3. Create a VRRP group and specify the group's virtual IP address.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface Vlan 50 ip address 172.16.1.1/24 no shutdown exit vrrp version 3 interface vlan 50 vrrp-group 50 virtual-address 172.16.1.254</pre>	<pre>interface Vlan 50 ip address 172.16.1.2/24 no shutdown exit vrrp version 3 interface vlan 50 vrrp-group 50 virtual-address 172.16.1.254</pre>

Add VLAN 50 to each downstream server-facing interface.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface ethernet 1/1/31 switchport trunk allowed vlan 50 interface ethernet 1/1/32 switchport trunk allowed vlan 50</pre>	<pre>interface ethernet 1/1/31 switchport trunk allowed vlan 50 interface ethernet 1/1/32 switchport trunk allowed vlan 50</pre>

Note: Additional configuration commands for the downstream Ethernet interfaces above are provided as part of Fibre Channel configuration in Section 5.2.2.

The two upstream layer 3 interfaces connected to the spine switches are configured. The IP addresses match those in [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#). Configure a loopback interface to be used as the router ID for BGP.

Note: If multiple loopback interfaces exist on a system, the interface with the highest numbered IP address is used as the router ID. This configuration only uses one loopback interface.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface ethernet 1/1/25 description "Z9100-Spine1 eth 1/1/1" no switchport ip address 192.168.1.1/31 no shutdown interface ethernet 1/1/26 description "Z9100-Spine2 eth 1/1/1"</pre>	<pre>interface ethernet 1/1/25 description "Z9100-Spine1 eth 1/1/2" no switchport ip address 192.168.1.3/31 no shutdown interface ethernet 1/1/26 description "Z9100-Spine2 eth 1/1/2"</pre>

S4148U-Leaf1	S4148U-Leaf2
<pre>no switchport ip address 192.168.2.1/31 no shutdown interface loopback 0 description "Router ID" ip address 10.0.2.1/32 no shutdown</pre>	<pre>no switchport ip address 192.168.2.3/31 no shutdown interface loopback 0 description "Router ID" ip address 10.0.2.2/32 no shutdown</pre>

Configure a route map and an IP prefix list to redistribute all loopback addresses and leaf networks via BGP.

S4148U-Leaf1	S4148U-Leaf2
<pre>route-map spine-leaf permit 10 match ip address prefix-list spine-leaf ip prefix-list spine-leaf seq 10 permit 10.0.0.0/8 ge 24 ip prefix-list spine-leaf seq 20 permit 172.16.0.0/16 ge 24</pre>	<pre>route-map spine-leaf permit 10 match ip address prefix-list spine-leaf ip prefix-list spine-leaf seq 10 permit 10.0.0.0/8 ge 24 ip prefix-list spine-leaf seq 20 permit 172.16.0.0/16 ge 24</pre>

Configure Uplink Failure Detection (UFD). This shuts down the downstream interfaces if all uplinks to the spine switches fail. The hosts attached to the leaf switch use their connection to the other leaf switch to continue sending traffic across the fabric.

S4148U-Leaf1	S4148U-Leaf2
<pre>uplink-state-group 1 name Shut_downstr_if_all_upstr_fail downstream ethernet1/1/31 downstream ethernet1/1/32 upstream ethernet1/1/25 upstream ethernet1/1/26 enable</pre>	<pre>uplink-state-group 1 name Shut_downstr_if_all_upstr_fail downstream ethernet1/1/31 downstream ethernet1/1/32 upstream ethernet1/1/25 upstream ethernet1/1/26 enable</pre>

Enable external Border Gateway Protocol (eBGP) with the `router bgp ASN` command. The `bestpath as-path multipath-relax` command enables equal-cost multi-path routing (ECMP). The `maximum-paths ebgp 2` command specifies the maximum number of parallel paths to a destination to add to the routing table.

Graceful restart enables the data plane to continue forwarding traffic for a time if the BGP process fails or quits. Neighbor fall-over is enabled, and BGP neighbors are configured.

Finally, exit configuration mode and save the configuration with the `end` and `write memory` commands.

S4148U-Leaf1	S4148U-Leaf2
<pre>router bgp 64701 bestpath as-path multipath-relax maximum-paths ebgp 2 address-family ipv4 unicast redistribute connected route-map spine-leaf exit graceful-restart role receiver-only template spine-leaf fall-over advertisement-interval 1 timers 3 9 exit neighbor 192.168.1.0 remote-as 64601 inherit template spine-leaf no shutdown exit neighbor 192.168.2.0 remote-as 64602 inherit template spine-leaf no shutdown end write memory</pre>	<pre>router bgp 64702 bestpath as-path multipath-relax maximum-paths ebgp 2 address-family ipv4 unicast redistribute connected route-map spine-leaf exit graceful-restart role receiver-only template spine-leaf fall-over advertisement-interval 1 timers 3 9 exit neighbor 192.168.1.2 remote-as 64601 inherit template spine-leaf no shutdown exit neighbor 192.168.2.2 remote-as 64602 inherit template spine-leaf no shutdown end write memory</pre>

6 S4148U-ON validation

After connected devices are configured, many commands are available to validate the network configuration. This section provides a list of the most common commands and their output for this topology.

Note: The commands and output shown below are for S4148U-Leaf1. The output for S4148U-Leaf2 is similar. For additional commands and output related to the leaf-spine portion of the topology, such as BGP, etc., see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).

6.1 General commands

6.1.1 show interface status

The **show interface status | grep up** command is used to verify all required interfaces are up and links are established at their appropriate speeds.

```
S4148U-Leaf1# show interface status | grep up
Port          Description      Status  Speed  Duplex  Mode  Vlan  Tagged-Vlans
Eth 1/1/25    Z9100-Spine1 .. up    100G   full   -
Eth 1/1/26    Z9100-Spine2 .. up    100G   full   -
Eth 1/1/29    VLTi             up     100G   full   -
Eth 1/1/30    VLTi             up     100G   full   -
Eth 1/1/31    To Server 1     up     10G    full   T    1    50,1002
Eth 1/1/32    To Server 2     up     10G    full   T    1    50,1002
Fc 1/1/1      To SPB          up     16G    auto   -
Fc 1/1/3      To SPA          up     16G    auto   -
```

6.1.2 show vlan

Port VLAN assignment is verified with the **show vlan** command. Po1000 is the VLTi, and it is automatically created when the VLTi links are selected during the configuration of VLT. Po1000 is a member of all active VLANs except for VLAN 1002. When VLAN 1002 is specified as the FCoE VLAN during vfabric configuration, it is automatically removed from the VLTi link. VLAN 4094 is reserved and is the internal control VLAN for the VLT domain.

```
S4148U-Leaf1# show vlan
Codes: * - Default VLAN, M - Management VLAN, R - Remote Port Mirroring VLANs
Q: A - Access (Untagged), T - Tagged
  NUM    Status  Description                                Q Ports
*   1     Active  Po1000                                     T Po1000
    50    Active  Eth1/1/5-1/1/24,1/1/31-1/1/54            A Eth1/1/5-1/1/24,1/1/31-1/1/54
    1002  Active  Eth1/1/31-1/1/32                          T Eth1/1/31-1/1/32
    4094  Active  Po1000                                     T Po1000
```

6.1.3 show vlt domain_id

VLT configuration is verified by running the `show vlt domain_id` command on each of the S4148U-ON switches. The VLTi is shown as port channel 1000, and the link status is up. The role of one switch in the VLT pair is primary and its peer switch (not shown) is assigned to the secondary role.

```
S4148U-Leaf1# show vlt 127
Domain ID           : 127
Unit ID            : 1
Role                : primary
Version            : 1.0
Local System MAC address : e4:f0:04:6b:04:42
VLT MAC address     : e4:f0:04:6b:04:42
IP address          : fda5:74c8:b79e:1::1
Delay-Restore timer : 90 seconds
Peer-Routing        : Enabled
Peer-Routing-Timeout timer : 0 seconds
VLTi Link Status
  port-channel1000 : up
```

VLT Peer	Unit ID	System MAC Address	Status	IP Address	Version
2		e4:f0:04:6b:01:42	up	fda5:74c8:b79e:1::2	1.0

6.1.4 show vlt domain_id backup link

The `show vlt domain_id backup link` command verifies the VLT backup link is functional on the OOB management network.

```
S4148U-Leaf1# show vlt 127 backup-link
VLT Backup Link
-----
Destination           : 100.67.169.30
Peer Heartbeat status : Up
```

6.2 FC commands

6.2.1 show fc switch

The `show fc switch` command verifies the switch is in F_Port mode for FC traffic.

```
S4148U-Leaf1# show fc switch
Switch Mode : FPORT
Switch WWN  : 10:00:e4:f0:04:6b:04:42
```

6.2.2 show fc ns switch

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

```
S4148U-Leaf1# show fc ns switch
```

```
Total number of devices = 4
```

```
Switch Name          10:00:e4:f0:04:6b:04:42
Domain Id            100
Switch Port          fibrechannel1/1/1
FC-Id                64:04:00
Port Name            50:06:01:6e:47:e0:1b:19
Node Name            50:06:01:60:c7:e0:1b:19
Class of Service     8
Symbolic Port Name   UNITY::::SPB10::FC:::::
Symbolic Node Name   UNITY::::SPA::FC:::::
Port Type            N_PORT
Registered with NameServer Yes
Registered for SCN   Yes
```

```
Switch Name          10:00:e4:f0:04:6b:04:42
Domain Id            100
Switch Port          fibrechannel1/1/3
FC-Id                64:0c:00
Port Name            50:06:01:66:47:e0:1b:19
Node Name            50:06:01:60:c7:e0:1b:19
Class of Service     8
Symbolic Port Name   UNITY::::SPA10::FC:::::
Symbolic Node Name   UNITY::::SPA::FC:::::
Port Type            N_PORT
Registered with NameServer Yes
Registered for SCN   Yes
```

```
Switch Name          10:00:e4:f0:04:6b:04:42
Domain Id            100
Switch Port          ethernet1/1/31
FC-Id                64:7c:00
Port Name            20:01:18:66:da:71:50:ad
Node Name            20:00:18:66:da:71:50:ad
Class of Service     8
Symbolic Port Name   QLogic 57800 bnx2fc v1.713.30.v60.6 over vmnic4
Symbolic Node Name   QLogic 57800 bnx2fc v1.713.30.v60.6 over vmnic4
Port Type            N_PORT
Registered with NameServer Yes
Registered for SCN   Yes
```

```

Switch Name          10:00:e4:f0:04:6b:04:42
Domain Id            100
Switch Port          ethernet1/1/32
FC-Id                64:80:00
Port Name            20:01:18:66:da:77:d0:c3
Node Name            20:00:18:66:da:77:d0:c3
Class of Service     8
Symbolic Port Name   QLogic 57800 bnx2fc v1.713.30.v60.6 over vmn4
Symbolic Node Name   QLogic 57800 bnx2fc v1.713.30.v60.6 over vmn4
Port Type            N_PORT
Registered with NameServer Yes
Registered for SCN   Yes

```

6.2.3 show fc zoneset active

The **show fc zoneset active** command shows the zones and zone members in the active zone set. Members logged into the fabric are shown with an asterisk (*).

```

S4148U-Leaf1# show fc zoneset active

vFabric id: 100
Active Zoneset: zoneset1
ZoneName          ZoneMember
=====
r640-1plzone      *20:01:18:66:da:71:50:ad
                  *50:06:01:66:47:e0:1b:19
                  *50:06:01:6e:47:e0:1b:19

r640-2plzone      *20:01:18:66:da:77:d0:c3
                  *50:06:01:66:47:e0:1b:19
                  *50:06:01:6e:47:e0:1b:19

```

6.2.4 show fcoe sessions

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

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

```

S4148U-Leaf1# show fcoe sessions
Enode MAC          Enode Interface  FCF MAC          FCF interface  VLAN
FCoE MAC           FC-ID           PORT WWPN        PORT WWNN
-----
-----

```

```

18:66:da:71:50:ad Eth 1/1/31 e4:f0:04:6b:05:41 ~ 1002
0e:fc:64:64:7c:00 64:7c:00 20:01:18:66:da:71:50:ad 20:00:18:66:da:71:50:ad

18:66:da:77:d0:c3 Eth 1/1/32 e4:f0:04:6b:05:41 ~ 1002
0e:fc:64:64:80:00 64:80:00 20:01:18:66:da:77:d0:c3 20:00:18:66:da:77:d0:c3

```

6.2.5 show vfabric

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

```

S4148U-Leaf1# show vfabric
Fabric Name
Fabric Type          FPORT
Fabric Id            100
Vlan Id              1002
FC-MAP               0xEFC64
Vlan priority        3
FCF Priority          128
FKA-Adv-Period       Enabled,8
Config-State         ACTIVE
Oper-State           UP
=====
Switch Config Parameters
=====
Domain ID             100
=====
Switch Zoning Parameters
=====
Default Zone Mode:   Deny
Active ZoneSet:      zoneset1
=====
Members
fibrechannel1/1/1
fibrechannel1/1/3
ethernet1/1/31
ethernet1/1/32
=====

```

Note: By default, the Default Zone Mode is set to Deny. This setting denies access between FC nodes except those in the same zone of the active zone set. To allow access between FC nodes in the absence of an active zone set, change the Default Zone Mode to Allow with the following command sequence:

```

S4148U-Leaf1# configure terminal
S4148U-Leaf1(config)# vfabric vfabric_number
S4148U-Leaf1(conf-vfabric-100)# zone default-zone permit

```

6.3 QoS commands

6.3.1 show qos maps

The **show qos maps** command shows the configured traffic-class-to-queue and dot1p-to-traffic-class maps. The command output also displays default traffic-class maps (not used in this deployment and removed from the output below to save space).

```
S4148U-Leaf1# show qos maps
Traffic-Class to Queue Map: map_GroupsToQueues
Queue          Traffic-Class
-----
0              0
3              3
DOT1P Priority to Traffic-Class Map : map_Dot1pToGroups
Traffic-Class   DOT1P Priority
-----
0              0-2,4-7
3              3
```

(output truncated)

6.3.2 show policy-map

The **show policy-map** command shows the configured policy maps.

Note: The `policy-iscsi` application policy map is configured by default and is not used in this deployment.

```
S4148U-Leaf1# show policy-map
Service-policy (application) input: policy-iscsi
Service-policy (qos) input: policy_Input_TrustDot1p
  Class-map (qos): class-trust
    trust dot1p
Service-policy (network-qos) input: policy_Input_PFC
  Class-map (network-qos): class_Dot1p_3
    pause
    pfc-cos 3
Service-policy (queuing) output: policy_Output_BandwidthPercent
  Class-map (queuing): map_ETSQueue_0
    bandwidth percent 50
  Class-map (queuing): map_ETSQueue_3
    bandwidth percent 50
```

6.3.3 show class-map

The `show class-map` command shows the configured class maps.

Note: The `class-iscsi` application class map is configured by default and is not used in this deployment.

```
S4148U-Leaf1# show class-map
Class-map (application): class-iscsi
Class-map (qos): class-trust
Class-map (network-qos): class_Dot1p_3
  Match: qos-group 3
Class-map (queuing): map_ETSQueue_0
  Match: queue 0
Class-map (queuing): map_ETSQueue_3
  Match: queue 3
```

6.3.4 show qos interface ethernet interface_number

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

```
S4148U-Leaf1# show qos interface ethernet 1/1/31
Interface
unknown-unicast-storm-control : Disabled
multicast-storm-control : Disabled
broadcast-storm-control : Disabled
flow-control-rx : Disabled
flow-control-tx : Disabled
ets mode : Enabled
pfc mode : Enabled
```

6.4 DCBX commands

6.4.1 show lldp dcbx interface ethernet interface_number

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

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

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

```

S4148U-Leaf1# show lldp dcbx interface ethernet 1/1/31
E-ETS Configuration TLV enabled           e-ETS Configuration TLV disabled
R-ETS Recommendation TLV enabled          r-ETS Recommendation TLV disabled
P-PFC Configuration TLV enabled           p-PFC Configuration TLV disabled
F-Application priority for FCOE enabled    f-Application Priority for FCOE
disabled
I-Application priority for iSCSI enabled    i-Application Priority for iSCSI
disabled

```

```

Interface ethernet1/1/31
  Port Role is Manual
  DCBX Operational Status is Enabled
  Is Configuration Source? FALSE
  Local DCBX Compatibility mode is CEE
  Local DCBX Configured mode is AUTO
  Peer Operating version is CEE
  Local DCBX TLVs Transmitted: ErPFI

```

Local DCBX Status

```

-----
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 1
Acknowledgment Number: 1
Protocol State: In-Sync

```

Peer DCBX Status

```

-----
DCBX Operational Version is 0
DCBX Max Version Supported is 0
Sequence Number: 1
Acknowledgment Number: 1
  32 Input PFC TLV pkts, 40 Output PFC TLV pkts, 0 Error PFC pkts
  32 Input PG TLV Pkts, 42 Output PG TLV Pkts, 0 Error PG TLV Pkts
  32 Input Appln Priority TLV pkts, 33 Output Appln Priority TLV pkts, 0 Error
Appln Priority TLV Pkts

```

```

Total DCBX Frames transmitted 47
Total DCBX Frames received 32
Total DCBX Frame errors 0
Total DCBX Frames unrecognized 0

```


6.4.2 show lldp dcbx interface ethernet *interface_number* pfc detail

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

```
S4148U-Leaf1# show lldp dcbx interface ethernet 1/1/31 pfc detail
Interface ethernet1/1/31
  Admin mode is on
  Admin is enabled, Priority list is 3
  Remote is enabled, Priority list is 3
  Remote Willing Status is enabled
  Local is enabled, Priority list is 3
  Oper status is init
  PFC DCBX Oper status is Up
  State Machine Type is Feature
  PFC TLV Tx Status is enabled
  Application Priority TLV Parameters :
  -----
  FCOE TLV Tx Status is enabled
  ISCSI TLV Tx Status is disabled
  Local FCOE PriorityMap is 0x08

  14 Input TLV pkts, 14 Output TLV pkts, 0 Error pkts
  14 Input Appln Priority TLV pkts, 6 Output Appln Priority TLV pkts, 0 Error
  Appln Priority TLV Pkts
```

6.4.3 show lldp dcbx interface ethernet *interface_number* ets detail

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

```
S4148U-Leaf1# show lldp dcbx interface ethernet 1/1/31 ets detail
Interface ethernet1/1/31
Max Supported PG is 8
Number of Traffic Classes is 8
Admin mode is on

Admin Parameters :
-----
Admin is enabled

PG-grp      Priority#      Bandwidth      TSA
-----
0           0,1,2,4,5,6,7  50%           ETS
1           1                0%            -
```

2		0%	-
3	3	50%	ETS
4		0%	-
5		0%	-
6		0%	-
7		0%	-
15		0%	SP

Remote Parameters :

Remote is enabled

PG-grp	Priority#	Bandwidth	TSA

0	0,1,2,4,5,6,7	50%	ETS
1		0%	-
2		0%	-
3	3	50%	ETS
4		0%	-
5		0%	-
6		0%	-
7		0%	-
15		0%	SP

Remote Willing Status is enabled

Local Parameters :

Local is enabled

PG-grp	Priority#	Bandwidth	TSA

0	0,1,2,4,5,6,7	50%	ETS
1		0%	-
2		0%	-
3	3	50%	ETS
4		0%	-
5		0%	-
6		0%	-
7		0%	-
15		0%	SP

Oper status is init

ETS DCBX Oper status is Up

State Machine Type is Feature

Conf TLV Tx Status is enabled

Reco TLV Tx Status is disabled

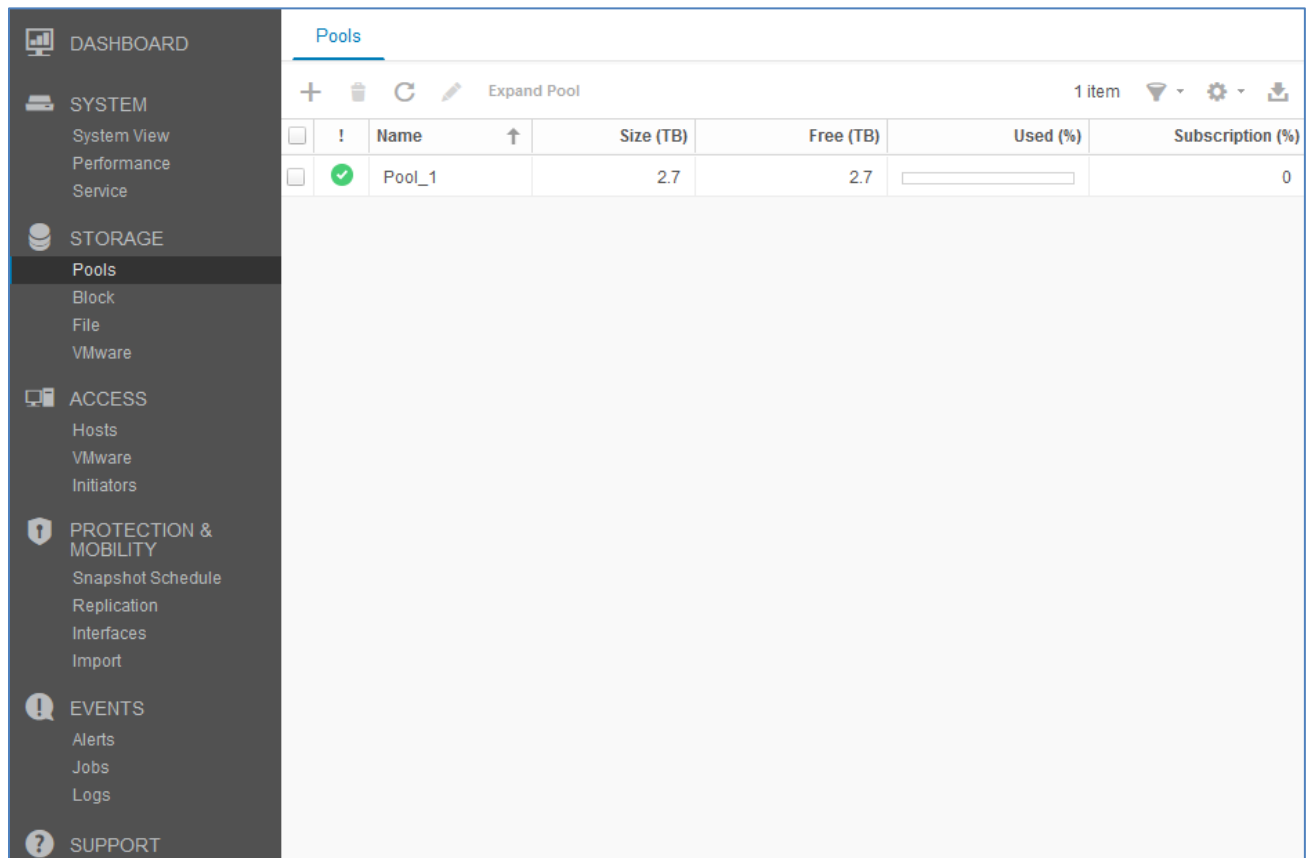
14 Input Conf TLV Pkts, 16 Output Conf TLV Pkts, 0 Error Conf TLV Pkts

7 Configure Unity FC storage

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

7.1 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 and RAID configuration for the pool. Click **Next**.
5. Under **Select Amount of Storage**, select the number of drives section. The total number of drives and the total capacity will be displayed next to **Totals**. Click **Next**.
6. The **Capability Profile Name** section is optional. Click **Next**.
7. Review selections on the **Summary** page and click **Finish** to create the pool. Once the **Overall status** shows 100%, click **Close**.
8. The pool is displayed on the **STORAGE > Pools** page as shown in Figure 21.



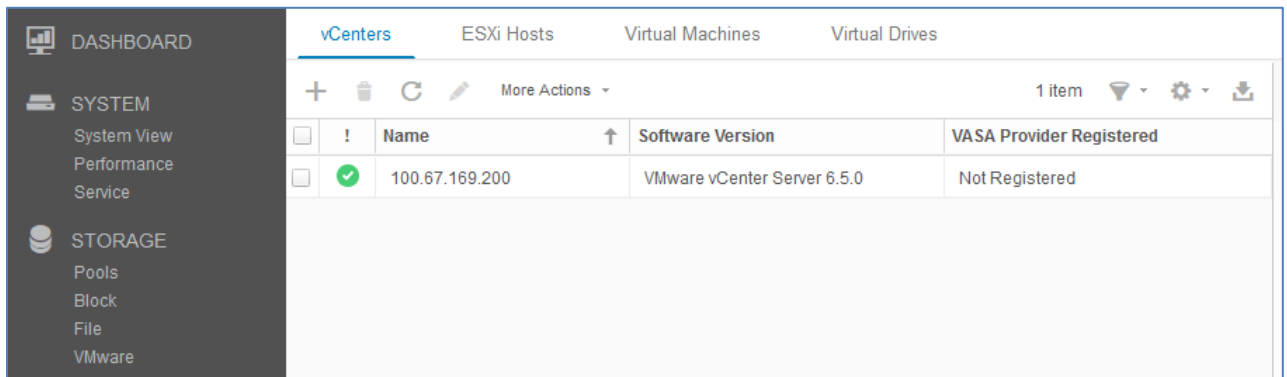
The screenshot displays the 'Pools' page in the Unisphere GUI. The left sidebar shows the navigation menu with 'STORAGE > Pools' selected. The main content area shows a table with one row representing the created pool 'Pool_1'. The table columns are: Name, Size (TB), Free (TB), Used (%), and Subscription (%). The 'Used (%)' column contains a progress bar. Above the table, there are icons for adding, deleting, refreshing, and expanding a pool, along with the text 'Expand Pool' and '1 item'.

	!	Name	↑	Size (TB)	Free (TB)	Used (%)	Subscription (%)
<input type="checkbox"/>	<input checked="" type="checkbox"/>	Pool_1		2.7	2.7	<div style="width: 0%;"></div>	0

Figure 21 Storage pool created

7.2 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 discovered ESXi hosts is displayed. Select the applicable hosts and click **Next**.
5. A VMware API for Storage Awareness (**VASA**) **Provider** is not used in this example. Click **Next**.
6. On the **Summary** page, review the ESXi Hosts to be added. Click **Finish**.
7. When the **Overall status** shows **100% Completed**, click **Close**.
8. The vCenter server is displayed as shown in Figure 22.

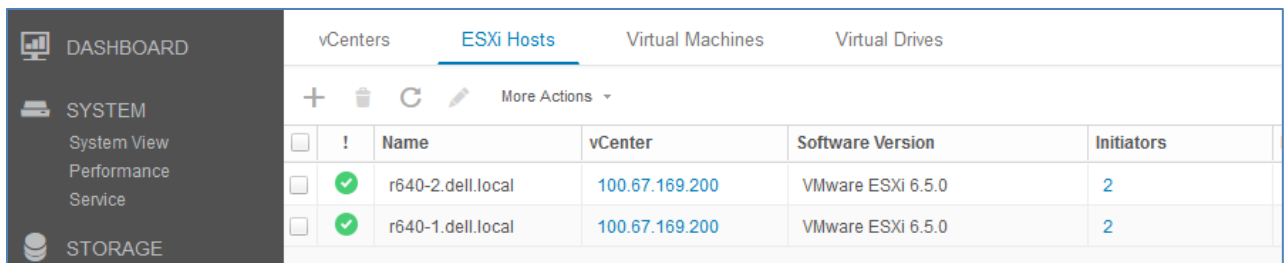


The screenshot shows the Unisphere interface with the 'vCenters' tab selected. The left sidebar contains 'DASHBOARD', 'SYSTEM' (System View, Performance, Service), and 'STORAGE' (Pools, Block, File, VMware). The main area displays a table with one row for a discovered vCenter server.

		Name	Software Version	VASA Provider Registered
<input type="checkbox"/>	!	100.67.169.200	VMware vCenter Server 6.5.0	Not Registered

Figure 22 vCenter server added to Unisphere

9. The list of added ESXi hosts is displayed on the **ESXi Hosts** tab, as shown in Figure 23.



The screenshot shows the Unisphere interface with the 'ESXi Hosts' tab selected. The left sidebar is the same as in Figure 22. The main area displays a table with two rows for discovered ESXi hosts.

		Name	vCenter	Software Version	Initiators
<input type="checkbox"/>	!	r640-2.dell.local	100.67.169.200	VMware ESXi 6.5.0	2
<input type="checkbox"/>	✓	r640-1.dell.local	100.67.169.200	VMware ESXi 6.5.0	2

Figure 23 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.

7.3 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 the **+** icon and select host(s) to be granted access to the LUN. Click **OK > Next**.


5. On the **Snapshot** page, leave settings at their defaults and click **Next**.
6. On the **Replication** page, leave settings at their defaults and click **Next**.
7. On the **Summary** page, review the details and click **Finish** to create the LUN.
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 24. In this example, a LUN named FC-80GB that is 80GB in size has been created.

		LUNs		Consistency Groups		iSCSI Interfaces	
				More Actions ▾		1 item	
<input type="checkbox"/>	!	Name ↓	Size (GB)	Allocated (%)	Pool	Thin Clone Base	
<input type="checkbox"/>	✓	FC-80GB	80.0	<div style="width: 0%;"></div>	Pool_1	--	

Figure 24 LUN Created


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

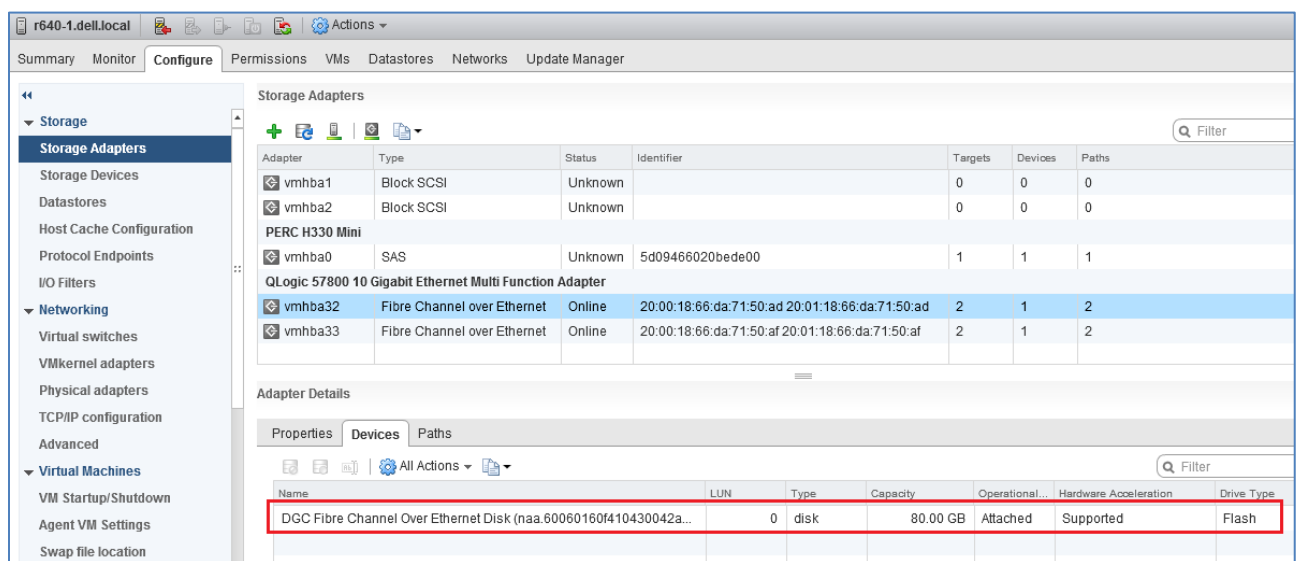
Note: To modify host access at any time, check the box next to the LUN to select it. Click the  icon, and select the **Host Access** tab.

8 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, hosts, and VMs.

8.1 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, vmhba32 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 25.



The screenshot shows the vSphere Web Client interface for host r640-1.dell.local. The left sidebar shows the navigation tree with 'Storage Adapters' selected. The main pane displays a table of storage adapters. The 'vmhba32' adapter is selected, and the 'Adapter Details' pane is open to the 'Devices' tab. A table below shows the discovered LUNs, with one LUN highlighted in red.

Adapter	Type	Status	Identifier	Targets	Devices	Paths
vmhba1	Block SCSI	Unknown		0	0	0
vmhba2	Block SCSI	Unknown		0	0	0
PERC H330 Mini						
vmhba0	SAS	Unknown	5d09466020bede00	1	1	1
QLogic 57800 10 Gigabit Ethernet Multi Function Adapter						
vmhba32	Fibre Channel over Ethernet	Online	20:00:18:66:da:71:50:ad 20:01:18:66:da:71:50:ad	2	1	2
vmhba33	Fibre Channel over Ethernet	Online	20:00:18:66:da:71:50:af 20:01:18:66:da:71:50:af	2	1	2

Name	LUN	Type	Capacity	Operational...	Hardware Acceleration	Drive Type
DGC Fibre Channel Over Ethernet Disk (naa.60060160f410430042a...	0	disk	80.00 GB	Attached	Supported	Flash

Figure 25 LUN visible to ESXi host

6. Repeat for host's second adapter, vmhba33 in this example. The LUN information on the **Adapter Details > Devices** tab is identical to the first adapter.
7. Select the first storage adapter, e.g., vmhba32, then select the **Adapter Details > Paths** tab as shown in Figure 26. The target, LUN number (e.g., LUN 0) and path status are shown. The target field includes the two active storage WWPNs connected to vmhba32. The status field is marked either **Active** or **Active (I/O)** for each path.

Adapter Details			
Properties Devices Paths			
Enable		Disable	
Runtime Name	Target	LUN	Status
vmhba32:C0:T0:L0	50:06:01:60:c7:e0:1b:19 50:06:01:66:47:e0:1b:19	0	◆ Active
vmhba32:C0:T1:L0	50:06:01:60:c7:e0:1b:19 50:06:01:6e:47:e0:1b:19	0	◆ Active (I/O)

Figure 26 Adapter Details - Paths tab

The **Paths** tab includes similar information for the host's second storage adapter.

8.2 Create a datastore

In this section, a datastore that uses the Unity LUN is created on the 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 the 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 27. In this example, the 80GB LUN mapped to this host appears in the list of devices.

New Datastore

1 Type
2 Name and device selection
 3 VMFS version
 4 Partition configuration
 5 Ready to complete

Name and device selection
 Select a name and a disk/LUN for provisioning the datastore.

Datastore name:

Filter

Name	LUN	Capacity	Hardware Accel...	Drive Type	Sector format
Local ATA Disk (t10.ATA_____DELLBOSS_VD_____...)	0	111.73 GB	Not supported	HDD	512e
DGC Fibre Channel Over Ethernet Disk (naa.60060...)	0	80.00 GB	Supported	Flash	512n

Figure 27 Name and device selection page

5. Provide a **Datastore name**, e.g., **Unity 80GB 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 is now accessible by selecting the host in the **Navigator** pane. Select the **Configure** tab > **Storage > Datastores** as shown in Figure 28.

Name	Status	Type	Datastore Cluster	Capacity	Free
host1localdatastore	Normal	VMFS 6		104.25 GB	84.85 GB
Unity 80GB LUN	Normal	VMFS 5		79.75 GB	78.8 GB

Figure 28 Datastore configured


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

8.3 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. VMs should be deployed before proceeding with this section.

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 using 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 80GB LUN, and click **OK**. The screen will look similar to Figure 29.

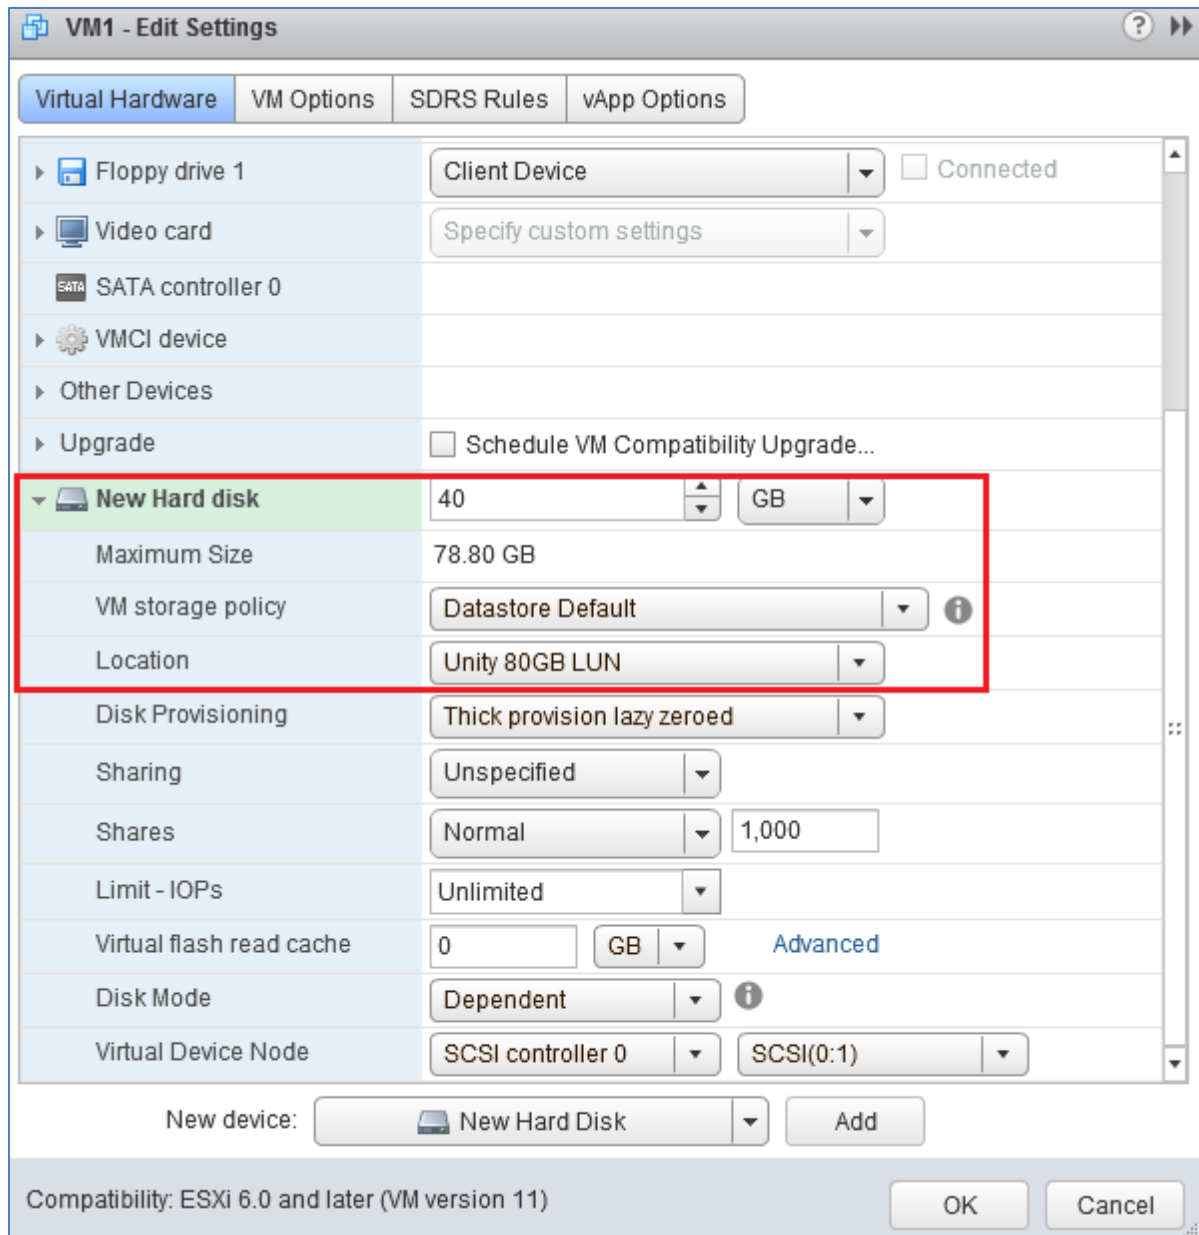


Figure 29 New hard disk configuration options

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.

8.4 Configure the virtual disk

Note: The following example is applicable for VMs running Windows Server 2008, 2012, or 2016. See the 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**.

Note: If an **Initialize Disk** window appears, select **OK** to initialize now, or **Cancel** to initialize in step 5.

3. If the new disk is not present in the list, right click on **Disk Management** and select **Rescan Disks**.
4. The new hard disk appears in the list (**Disk 1** in Figure 30).

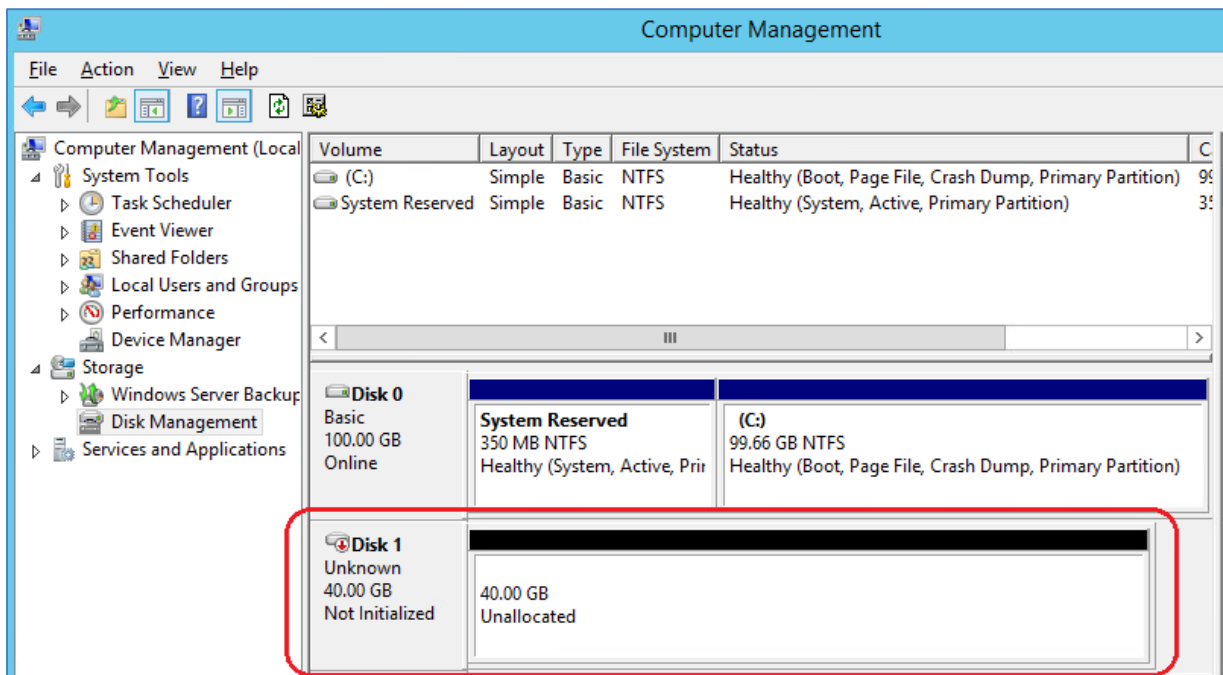


Figure 30 Windows Server Disk Management utility

5. If the disk is not initialized or online, 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.
6. Continue using the Disk Management application to format the disk.

9 Configure ESXi hosts for LAN traffic

In this section, ESXi hosts and VMs are configured for TCP/IP access to the production network. (Refer to the LAN topology shown in Figure 9.)

9.1 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 vSphere Web Client, and the configuration is populated across all hosts associated with the switch. A VDS is 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, connect the same vmnic on every host to the same uplink port on the distributed switch.

9.2 Create a VDS

In this section, a VDS named **R640-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, e.g., **R640-VDS**. Click **Next**.
4. On the **Select version** page, select **Distributed switch: 6.5.0** and click **Next**.
5. On the **Edit settings** page:
 - a. Since two CNA ports are 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 31.

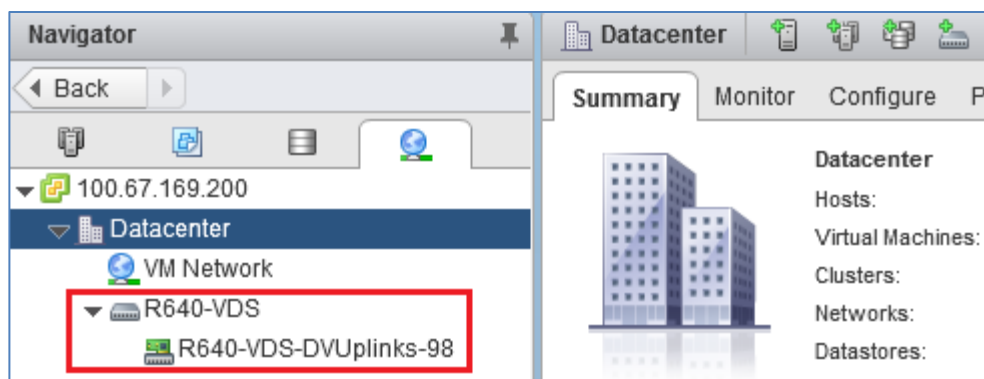


Figure 31 R640-VDS created with uplinks port group

9.3 Add a distributed port group

In this section, a distributed port group is created on the distributed switch.

To create the port group:

1. On the Web Client **Home** screen, select **Networking**.
2. Right-click on the distributed switch, e.g., **R640-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, e.g., **Production**. Click **Next**.
4. On the **Configure settings** page, next to **VLAN type**: select **VLAN**. Next to **VLAN ID**, specify the VLAN configured earlier for LAN traffic on the leaf switches. In this deployment, **VLAN 50** is used.
5. Leave the remaining values at their defaults and click **Next > Finish**.

When complete, the **Navigator** pane appears as shown:

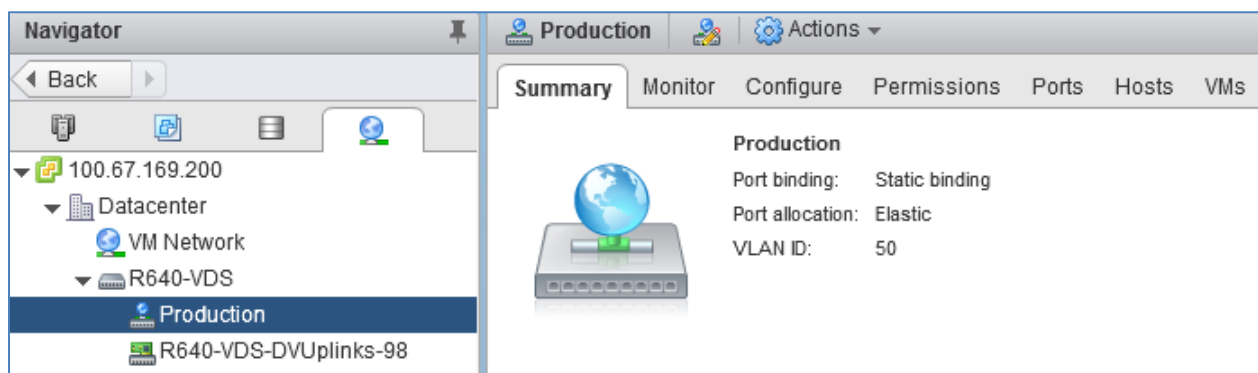


Figure 32 Distributed port group created for production LAN traffic

9.4 Configure load balancing

Note: It is a best practice to use **Route based on Physical NIC load** as the load balancing algorithm on distributed port groups. For more information, see [VMware Validated Design Documentation](#), release 4.2.

To configure load balancing on the distributed port group:

1. In the **Navigator** pane, right-click on the port group, e.g., **Production**, and select **Edit Settings**.
2. In the **Edit Settings** window, select **Teaming and Failover** in the left pane. In the right pane, change the **Load Balancing** method to **Route based on Physical NIC load**. Other settings should remain at their defaults as shown:

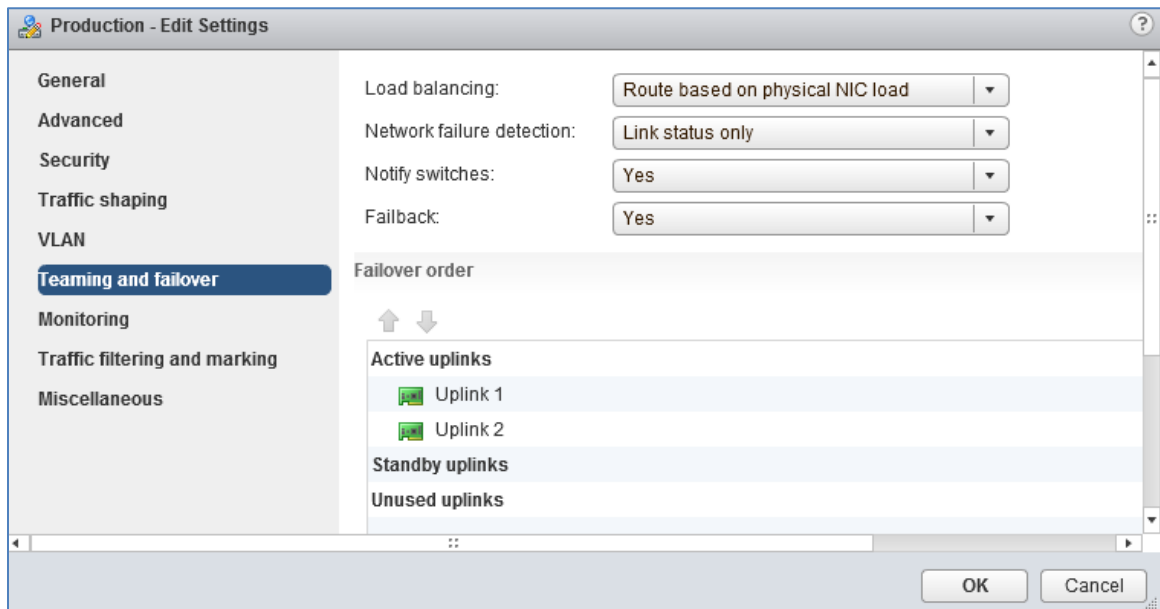


Figure 33 Port group teaming and failover settings

3. Click **OK**.

9.5 Add hosts to the VDS

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

Note: Before starting this section, be sure the vmnic-to-physical adapter mapping for each host's ports for production IP traffic is known. This is determined by selecting the host in the **Navigator** pane, then going to **Configure > Networking > Physical adapters** in the center pane. Vmnics and their MAC addresses are shown. Running the command `show lldp neighbors` on the S4148U-ON switches displays the MAC addresses of the connected QLogic 57800 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 R640-VDS:

1. On the Web Client **Home** screen, select **Networking**.
2. Right click on **R640-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 **+ New hosts...** icon. Select the checkbox by each host to add. 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 **R640-VDS** will look similar to Figure 34.

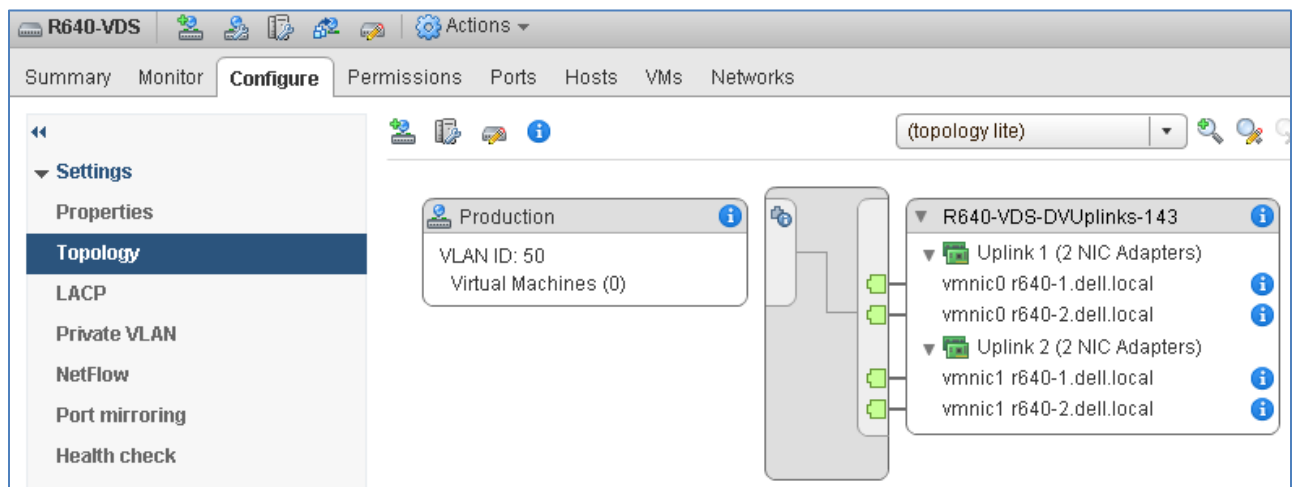



Figure 34 Uplinks configured on R640-VDS

9.6 Add a virtual network adapter to VMs

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

1. In the vSphere Web Client, go to **Home > Hosts and Clusters**.
2. Under the **Rack1** cluster, right click on a VM and click **Edit Settings**.
3. Next to **New Device**, select **Network**. Click **Add**.
4. Next to **New Network**, 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 to select the **Production** port group. If only the VDS that is shown in the drop-down menu is selected, e.g., R640-VDS, a Port ID error will be displayed.

5. On the **Select Network** page, select the **Production** port group on **R640-VDS** and click **OK**.
6. Click the  icon next to **New Network** to view configuration options.
 - a. Next to **Status**, ensure the **Connect at Power On** box is checked.
 - b. 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, with key network settings outlined in red.

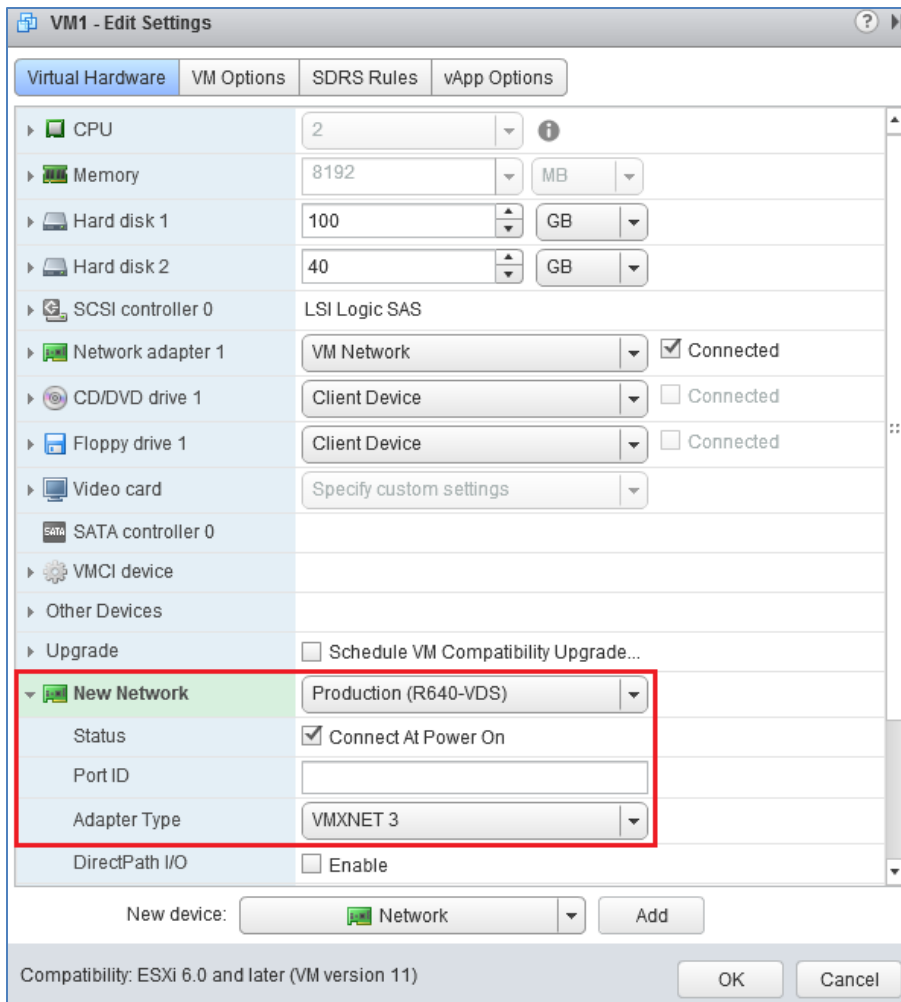


Figure 35 Virtual network adapter configured

Note: In the top half of Figure 35, an additional adapter named **Network adapter 1** on **VM Network** is present. VM Network is the OOB Management network in this deployment. If present, the adapter connected to VM Network may be retained or deleted as needed.

7. Click **OK** to add the **New Network** adapter.

Repeat the steps above for remaining VMs that will access the leaf-spine production network.

9.7 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 and default gateway on the newly added vNIC.

VMs in this deployment have IP addresses on the 172.16.1.0/24 network. The default gateway is the VRRP IP address configured on the S4148U-ON leaf switches, 172.16.1.254.

Test connectivity by pinging the default gateway and other configured VMs in the rack. If spines, leaf switches, and hosts in other racks are configured per [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#), pings will also succeed 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 Validated components

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

A.1 Switches

Table 3 Switches and firmware versions

Qty	Item	OS Version
2	Dell EMC Networking S4148U-ON leaf switches	10.4.0E(R3P2)
2	Dell EMC Networking Z9100-ON spine switches	10.4.0E(R3P2)
1	Dell EMC Networking S3048-ON management switch	10.4.0E(R3P2)

A.2 PowerEdge R640 servers

Table 4 R640 server components

Qty per server	Item	Firmware version
1	Intel Xeon Gold 6126 CPU @ 2.60GHz, 12 cores	-
6	8GB DDR4 DIMMs (48GB total)	-
1	120GB SATA SSD	-
1	PERC H330 Mini Storage Controller	25.5.3.0005
1	QLogic 57800 rNDC - 2x10GbE SFP+ (CNA) ports plus 2x1GbE BASE-T ports	14.02.12
-	BIOS	1.3.7
-	iDRAC with Lifecycle Controller	3.00.00.00

A.3 Storage

Table 5 Storage firmware version

Qty	Item	Firmware Version
1	Unity 500F	4.3.0.1522077968

A.4 VMware software

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

Table 6 VMware software versions

Item	Version
VMware ESXi	6.5 U1 - Dell EMC customized image version A10, build 7967591
VMware vCenter Server Appliance	6.5 Update 1d – build 7312210 (Includes vSphere Web Client)

A.5 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 – 4 CPU sockets
- vCenter 6 Server Standard – 1 instance

B 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.

B.1 Dell EMC product manuals and technical guides

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

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

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

[Manuals and documentation for Dell EMC PowerEdge R640](#)

[OS10 Enterprise Edition User Guide Release 10.4.0E\(R3\)](#)

[Dell EMC Unity Technical Documentation](#)

[Dell EMC TechCenter Networking Guides](#)

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

B.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](#)

C Support and feedback

Contacting Technical Support

Support Contact Information

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

Telephone: USA: 1-800-945-3355

Feedback for this document

We encourage readers to provide feedback on the quality and usefulness of this publication by sending an email to Dell_Networking_Solutions@Dell.com.