

Dell EMC Networking FCoE Deployment with S4148U-ON in NPG Mode

Connecting server FCoE CNAs to Fibre Channel storage using two Dell EMC PowerSwitch S4148U-ON switches and Brocade Fibre Channel switches

Abstract

This document provides the deployment steps for configuring S4148U-ON switches in NPIV Proxy Gateway (NPG) mode connected to FC switches and storage.

April 2019

Revisions

Date	Description
April 2019	Initial release

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

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

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

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

Table of contents

Revisions.....	2
1 Introduction.....	5
1.1 Typographical conventions.....	6
1.2 Attachments.....	6
2 Hardware Overview.....	7
2.1 Dell EMC PowerSwitch models.....	7
2.1.1 Dell EMC PowerSwitch S4148U-ON.....	7
2.1.2 Dell EMC PowerSwitch S3048-ON.....	7
2.2 Storage arrays, Fibre Channel switches, and servers.....	7
2.2.1 Dell EMC Unity 500F storage array.....	8
2.2.2 Dell EMC PowerEdge R740xd server.....	8
2.2.3 Brocade 6510 Switch.....	8
3 Topology overview.....	9
3.1 FC SAN topology detail.....	9
3.2 OOB management network.....	10
4 Deployment Overview.....	12
4.1 Configuration strategy and sequence.....	12
4.1.1 FC storage array.....	12
4.1.2 Dell EMC PowerSwitch S4148U-ON leaf switches.....	12
4.1.3 FC Switch.....	12
4.1.4 Server HW and Virtualization.....	12
5 S4148U-ON switch configuration.....	13
5.1 Prepare switches.....	13
5.1.1 Factory default configuration.....	13
5.1.2 Set switch port profile.....	13
5.2 Configure switches.....	14
6 S4148U-ON validation.....	21
6.1 show interface status.....	21
6.2 show fc switch.....	21
6.3 show fcoe sessions.....	22
6.4 show fcoe enode.....	22
6.5 show vfabric.....	23
A Validated components.....	24
B PowerEdge server, Unity storage, and VMware setup.....	25
B.1 PowerEdge server configuration.....	25

B.1.1	Reset server CNA interfaces to factory defaults	25
B.1.2	Determine FCoE CNA port WWPNS	25
B.2	Dell EMC Unity 500F storage array configuration	27
B.2.1	Create a storage pool	28
B.2.2	Add ESXi hosts.....	28
B.2.3	Create LUNs and configure host access.....	29
B.3	VMware preparation	30
B.3.1	VMware ESXi download and installation.....	30
B.3.2	Install and configure VMware vCenter Server 6.7 U1	30
B.3.3	Add ESXi hosts to vCenter Server	30
B.3.4	Configure storage on ESXi hosts	31
B.3.5	Rescan storage.....	31
B.3.6	Create a datastore	32
B.3.7	Create a virtual disk.....	33
B.3.8	Configure the virtual disk in Windows Server.....	35
C	Technical resources	36
D	Support and feedback	37

1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company from the edge, to the core, and to the cloud. Dell EMC Networking ensures modernization for today's applications and for the emerging cloud-native world. Dell EMC is committed to disrupting the fundamental economics of the market with an open strategy that gives you the freedom of choice for networking operating systems and top-tier merchant silicon. The Dell EMC strategy enables business transformations that maximize the benefits of collaborative software and standards-based hardware, including lowered costs, flexibility, freedom, and security. Dell EMC provides further customer enablement through validated deployment guides which demonstrate these benefits while maintaining a high standard of quality, consistency, and support.

In a leaf-spine data center environment, two Dell EMC PowerSwitch S4148U-ON switches running OS10 Enterprise Edition (OS10EE) may be used as a leaf pair with Ethernet ports connected to spine switches for production TCP/IP traffic. The same leaf pair is utilized in NPG mode for connections to Fibre Channel (FC) switches servicing FC storage.

This guide covers configuration of the Fibre Channel/Fibre Channel over Ethernet (FC/FCoE) portion of the network. The configuration includes two S4148U-ON switches, PowerEdge server FCoE converged network adapters (CNAs), VMware ESXi hosts, and a Unity 500F storage array. The goal of this guide is to enable a network administrator or deployment services engineer to deploy a pair of S4148U-ON switches in an FC 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
<i>Italic Monospace Text</i>	Variables in CLI examples
Bold Monospace Text	Commands entered at the CLI prompt, or to highlight information in CLI output
Bold text	UI elements and information entered in the GUI

1.2 Attachments

This document in .pdf format includes one or more 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 hardware that is used to validate the deployment examples in this document. Appendix A contains a complete listing of hardware and software validated for this guide.

Note: While the steps in this document were validated using the specified Dell EMC PowerSwitch models and operating system(s), they may be leveraged for other Dell EMC PowerSwitch models utilizing the same networking OS version or later assuming the switch has the available port numbers, speeds, and types. At the date of publication, the S4148U-ON switch is the only OS10 Enterprise Edition (OS10EE) based switch supporting Fibre Channel services.

2.1 Dell EMC PowerSwitch models

This section details the Dell EMC PowerSwitch models in the networking topology shown in Figure 5.

2.1.1 Dell EMC PowerSwitch S4148U-ON

The S4148U-ON enables converging LAN and SAN traffic in a single rack unit (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 Fibre Channel connectivity to the SAN and Ethernet connectivity to the leaf-spine network.



Figure 1 Dell EMC PowerSwitch S4148U-ON

2.1.2 Dell EMC PowerSwitch S3048-ON

The Dell EMC PowerSwitch S3048-ON is a 1-RU switch with forty-eight 1GbE BASE-T ports and four 10GbE SFP+ ports. In this document, one S3048-ON supports out-of-band (OOB) management traffic for all examples.



Figure 2 Dell EMC PowerSwitch S3048-ON

2.2 Storage arrays, Fibre Channel switches, and servers

This section details the supplemental hardware used to validate a complete storage solution. Comparable hardware models can be substituted for the hardware listed in this section to operate on the network topology described in this document.

2.2.1 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 3 Dell EMC Unity 500F front view

2.2.2 Dell EMC PowerEdge R740xd server

The Dell EMC PowerEdge R740xd server is a dual socket 2-RU platform that brings scalable storage performance and data set processing to adapt to a variety of applications. The R740xd server features Intel Xeon processors, with up to 3 terabytes of expandable memory and network interface technologies to cover a variety of high-performance NIC and rNDC network interface options. The server adds extraordinary storage capacity options, making it well-suited for data-intensive applications that require greater storage, without sacrificing I/O performance.



Figure 4 Dell EMC PowerEdge R740xd server

Note: The servers used in supporting the validation of the network are Dell EMC PowerEdge R740xd models. The server model is not the focus of this document and are referenced using the generic term node.

2.2.3 Brocade 6510 Switch

The Brocade 6510 Fibre Channel switch is a 48 port SAN switch used for dedicated FC storage networks supporting 2, 4, 8, 10, or 16 Gbps speeds.

3 Topology overview

Each S4148U-ON provides unified ports, configured as 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 also act as a leaf in a leaf-spine network for production TCP/IP traffic.

Note: FC SAN traffic does not traverse the leaf-spine network, or the VLT interconnect (VLTi).

A combined FC SAN and leaf-spine topology are shown in the diagram below.

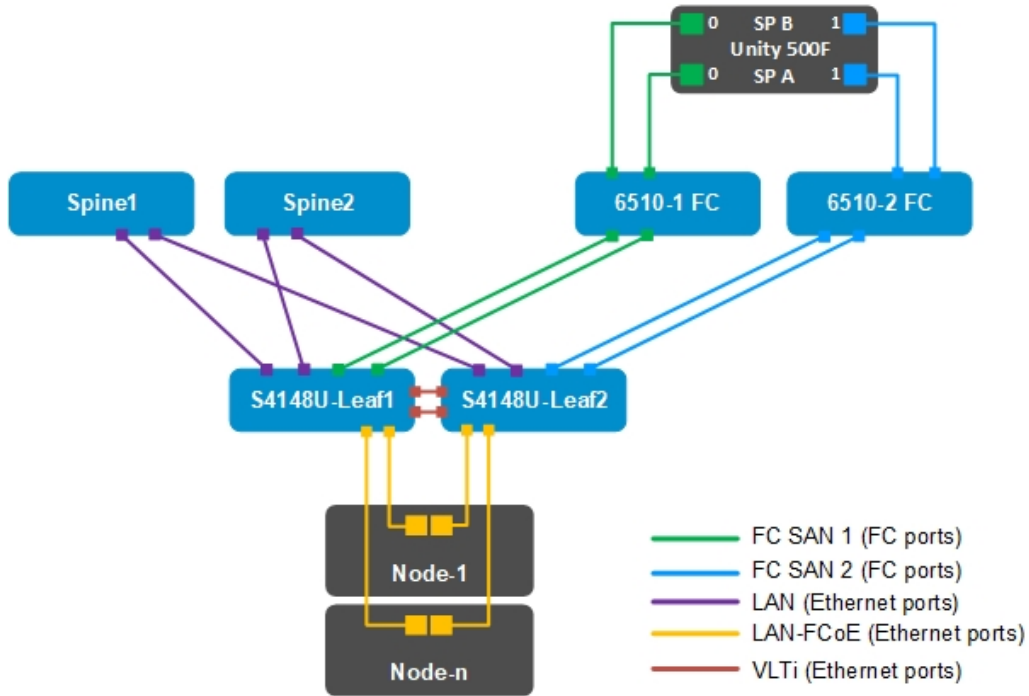


Figure 5 Combined FC SAN and leaf-spine topology

Note: Using a leaf-spine network in the data center is considered a best practice. This document covers configuration of the FC SAN portion of the topology above. For the leaf-spine Ethernet portion, refer to [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

3.1 FC SAN topology detail

Each S4148U-ON switch is placed in NPG mode and provides FCoE to FC bridging. FC services are provided by the dedicated FC network. In this example, the FC switches are Brocade 6510 switches.

The FC SAN topology used in this guide is shown in Figure 6.

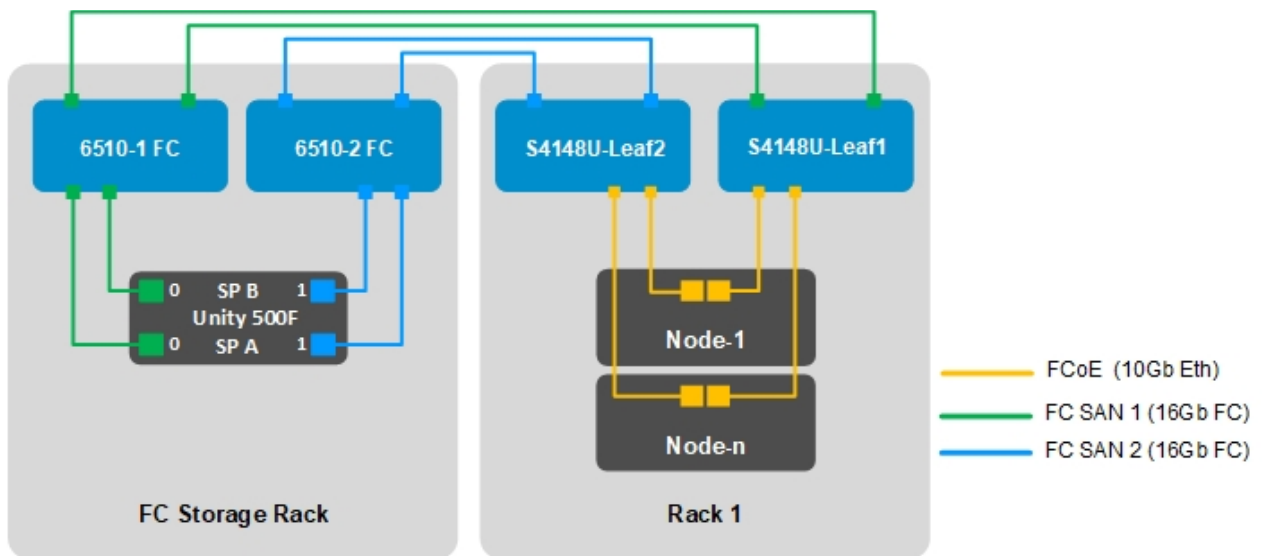


Figure 6 FC SAN topology

3.2 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 management switch is installed at the top of each rack for OOB management connections as shown.

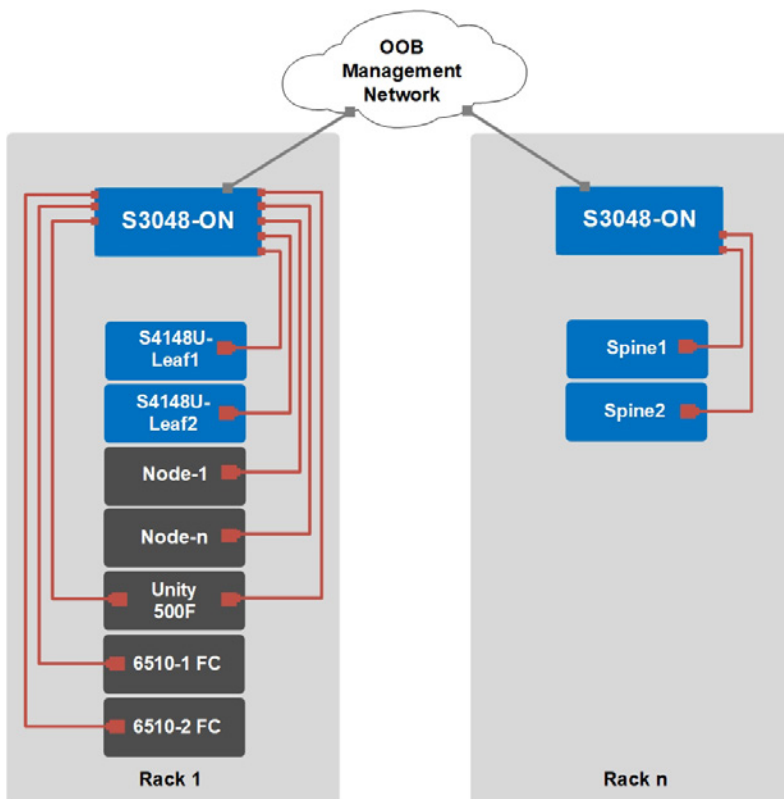


Figure 7 OOB management network connections

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 R740xd server has a connection to the S3048-ON for the server's iDRAC port.

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

4 Deployment Overview

This section provides high-level guidance for deploying the total solution including FC storage, networking, server resources, and virtualization.

4.1 Configuration strategy and sequence

This document provides specific configuration examples for the S4148U-ON leaf pair.

Note: A Dell EMC Unity 500F storage array, Brocade 6510 FC switches, and Dell EMC PowerEdge R740xd servers were used to validate the complete solution. For details related to setup and their components, see Appendix A & B.

The following sequence can be used as a guide for how the complete solution can be set up from start to finish. Basic connectivity and related SW installation are required before starting.

4.1.1 FC storage array

1. Create storage pools
2. Add hosts
3. Create LUNs
4. Configure host access

4.1.2 Dell EMC PowerSwitch S4148U-ON leaf switches

1. Enable NPG
2. Enable Data Center Bridging Exchange (DCBx)
3. Configure FCoE VLAN
4. Configure vFabric
5. Disable link-level flow control (LLFC) on FCoE interfaces
6. Configure Enhanced Transmission Selection/ Priority-based Flow Control (ETS/PFC)
7. Configure Quality of Service (QoS)
8. Configure host-facing interfaces
9. Configure FC switch-facing interfaces

4.1.3 FC Switch

Configure zoning

4.1.4 Server HW and Virtualization

1. Enable FCoE on CNAs in System Setup
2. Configure virtual machines

5 S4148U-ON switch configuration

This section details steps to configure the S4148U-ON leaf switches running OS10EE.

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 PowerSwitch models running OS10EE 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: OS10EE at its default settings has 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 deployment to enable four 100GbE ports on each S4148U-ON for the leaf-spine portion of the topology. This profile also provides twelve 16Gb FC ports for storage connections.

Note: See [OS10 Enterprise Edition User Guide Release 10.4.2.0](#) 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 setting the switch port profile on both S4148U-ONs, the commands in the tables that follow are run to complete the FC configuration on both switches.

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

Configure global switch settings

Configure the hostname, OOB management IP address, and OOB management default gateway. Configure the first port group in FC mode for 16GFC.

These FC interfaces are used to connect to the Brocade 6510 FC switches.

S4148U-Leaf1	S4148U-Leaf2
<pre>configure terminal hostname S4148U-Leaf1 interface mgmt 1/1/1 no ip address dhcp ip address 100.67.166.219/24 no shutdown management route 0.0.0.0/0 100.67.166.254 port-group 1/1/1 mode fc 16g-2x</pre>	<pre>configure terminal hostname S4148U-Leaf2 interface mgmt 1/1/1 no ip address dhcp ip address 100.67.166.220/24 no shutdown management route 0.0.0.0/0 100.67.166.254 port-group 1/1/1 mode fc 16g-2x</pre>

Note: For more information on mapping ports to switch profiles and port groups, see [OS10 Enterprise Edition User Guide Release 10.4.2.0](#), Unified port groups, page 84.

Configure VLT between the two S4148U-ON switches

VLT synchronizes Layer 2 and ARP table information between the switches and enables a single port channel to be connected to the switch pair as if they are a single switch.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface range ethernet 1/1/29-1/1/30 no switchport vlt-domain 127 backup destination 100.67.166.220 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 no switchport vlt-domain 127 backup destination 100.67.166.219 discovery-interface ethernet 1/1/29 discovery-interface ethernet 1/1/30 peer-routing</pre>

Configure the FC VLAN and virtual fabrics

For each switch define the VLANs and virtual fabrics. Apply to appropriate FC and server interfaces. The global command `feature fc npg` puts the switch in NPG mode.

S4148U-Leaf1	S4148U-Leaf2
<pre>feature fc npg interface vlan 1001 vfabric 101 vlan 1001 fcoe fcmmap 0xEFC64 interface fibrechannel 1/1/1 no shutdown vfabric 101 interface fibrechannel 1/1/3 no shutdown vfabric 101 interface ethernet 1/1/31 vfabric 101 interface ethernet 1/1/32 vfabric 101 interface ethernet 1/1/33 vfabric 101 interface ethernet 1/1/34 vfabric 101</pre>	<pre>feature fc npg interface vlan 1002 vfabric 102 vlan 1002 fcoe fcmmap 0xEFC65 interface fibrechannel 1/1/1 no shutdown vfabric 102 interface fibrechannel 1/1/3 no shutdown vfabric 102 interface ethernet 1/1/31 vfabric 102 interface ethernet 1/1/32 vfabric 102 interface ethernet 1/1/33 vfabric 102 interface ethernet 1/1/34 vfabric 102</pre>

Note: The FCoE VLAN and FCoE map are used internally by the switch to forward FC traffic, even if FCoE is not used by nodes connected to the switch.

The bulk of the FC portion of the configuration is now complete.

The following steps complete the configuration example by configuring common features found in a layer 3 leaf-spine network. For more information on leaf-spine networks see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

Configure the compute VLAN interfaces

The VLANs shown in this example represent a generic converged or hyper-converged deployment and should be changed to apply to the deployment specifications.

S4148U-Leaf1	S4148U-Leaf2
<pre>interface Vlan 1612 ip address 172.16.12.1/24 description "vMotion" no shutdown interface Vlan 1613 ip address 172.16.13.1/24 description "vSAN" no shutdown interface Vlan 1614 ip address 172.16.14.1/24 description "web" no shutdown interface Vlan 1615 ip address 172.16.15.1/24 description "app" no shutdown interface Vlan 1616 ip address 172.16.16.1/24 description "db" no shutdown</pre>	<pre>interface Vlan 1612 ip address 172.16.12.2/24 description "vMotion" no shutdown interface Vlan 1613 ip address 172.16.13.2/24 description "vSAN" no shutdown interface Vlan 1614 ip address 172.16.14.2/24 description "web" no shutdown interface Vlan 1615 ip address 172.16.15.2/24 description "app" no shutdown interface Vlan 1616 ip address 172.16.16.2/24 description "db" no shutdown</pre>

Configure VRRP

VRRP is an active/standby first-hop redundancy protocol. When used among VLT peers, it becomes active/active. Both VLT peers have the VRRP virtual MAC address in their forwarding table as a local destination address. This allows the backup VRRP router to forward intercepted frames whose destination MAC address matches the VRRP virtual MAC address.

S4148U-Leaf1	S4148U-Leaf2
<pre>vrrp version 3 interface vlan 1612 vrrp-group 12 virtual-address 172.16.12.254 interface vlan 1613 vrrp-group 13 virtual-address 172.16.13.254 interface vlan 1614 vrrp-group 14 virtual-address 172.16.14.254</pre>	<pre>vrrp version 3 interface vlan 1612 vrrp-group 12 virtual-address 172.16.12.254 interface vlan 1613 vrrp-group 13 virtual-address 172.16.13.254 interface vlan 1614 vrrp-group 14 virtual-address 172.16.14.254</pre>

S4148U-Leaf1	S4148U-Leaf2
<pre>interface vlan 1615 vrrp-group 15 virtual-address 172.16.15.254 interface vlan 1616 vrrp-group 16 virtual-address 172.16.16.254</pre>	<pre>interface vlan 1615 vrrp-group 15 virtual-address 172.16.15.254 interface vlan 1616 vrrp-group 16 virtual-address 172.16.16.254</pre>

Configure server-facing interfaces

S4148U-Leaf1	S4148U-Leaf2
<pre>interface ethernet 1/1/31 description "Server 1" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/32 description "Server 2" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/33 description "Server 3" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/34 description "Server 4" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown</pre>	<pre>interface ethernet 1/1/31 description "Server 1" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/32 description "Server 2" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/33 description "Server 3" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown interface ethernet 1/1/34 description "Server 4" switchport mode trunk switchport trunk allowed vlan 1612- 1616 spanning-tree port type edge mtu 9216 no shutdown</pre>

Configure upstream interfaces to spine switches

S4148U-Leaf1	S4148U-Leaf2
<pre>interface ethernet 1/1/25 description "Z9264F-Spine1 eth 1/1/1" no switchport ip address 192.168.1.1/31 no shutdown interface ethernet 1/1/26 description "Z9264F-Spine2 eth 1/1/1" 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>interface ethernet 1/1/25 description "Z9264F-Spine1 eth 1/1/2" no switchport ip address 192.168.1.3/31 no shutdown interface ethernet 1/1/26 description "Z9264F-Spine2 eth 1/1/2" 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>

Note: Configuration of the spine switches are not detailed in this document. For information on the leaf-spine architecture and configuration, see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

Configure BGP routing

The example configuration use of BGP is for the application traffic associated with the layer 3 leaf-spine network. Other routing protocols can be utilized and will not affect the FC storage configuration.

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 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</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 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</pre>

S4148U-Leaf1	S4148U-Leaf2
<pre>remote-as 64601 inherit template spine-leaf no shutdown exit neighbor 192.168.2.0 remote-as 64602 inherit template spine-leaf no shutdown</pre>	<pre>remote-as 64601 inherit template spine-leaf no shutdown exit neighbor 192.168.2.2 remote-as 64602 inherit template spine-leaf no shutdown</pre>

Configure uplink failure detection (UFD)

UFD is recommend on all server-facing interfaces.

S4148U-Leaf1	S4148U-Leaf2
<pre>uplink-state-group 1 name "UFD_Group_1" downstream ethernet1/1/31-1/1/34 upstream ethernet1/1/25-1/1/26 enable</pre>	<pre>uplink-state-group 1 name "UFD_Group_1" downstream ethernet1/1/31-1/1/34 upstream ethernet1/1/25-1/1/26 enable</pre>

Configure QoS

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).

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

S4148U-Leaf1	S4148U-Leaf2
<pre> dcbx enable 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 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 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 interface range ethernet 1/1/31- 1/1/34,1/1/41 flowcontrol receive off flowcontrol transmit off priority-flow-control mode on ets mode on service-policy input type network-qos policy_Input_PFC service-policy output type queuing policy_Output_BandwidthPercent system qos trust-map dot1p map_Dot1pToGroups qos-map traffic-class map_GroupsToQueues end write memory </pre>	<pre> dcbx enable 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 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 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 interface range ethernet 1/1/31- 1/1/34,1/1/41 flowcontrol receive off flowcontrol transmit off priority-flow-control mode on ets mode on service-policy input type network-qos policy_Input_PFC service-policy output type queuing policy_Output_BandwidthPercent system qos trust-map dot1p map_Dot1pToGroups qos-map traffic-class map_GroupsToQueues end write memory </pre>

6 S4148U-ON validation

After configuring connected devices, 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 VLT, BGP, etc., see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#).

6.1 show interface status

The `show interface status | grep up` command is used to verify 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
Fc 1/1/1      To SPB           up       16G     auto     -
Fc 1/1/3      To SPA           up       16G     auto     -
Eth 1/1/25    Z9264F-Spine1.. up       100G    full     -
Eth 1/1/26    Z9264F-Spine2.. up       100G    full     -
Eth 1/1/29    VLTi             up       100G    full     -
Eth 1/1/30    VLTi             up       100G    full     -
Eth 1/1/31    Server 1         up       10G     full     T    1    1001,1612-1616
Eth 1/1/32    Server 2         up       10G     full     T    1    1001,1612-1616
Eth 1/1/33    Server 3         up       10G     full     T    1    1001,1612-1616
Eth 1/1/34    Server 4         up       10G     full     T    1    1001,1612-1616
Eth 1/1/41                    up       10G     full     T    1    1001
```

6.2 show fc switch

The `show fc switch` command verifies the switch is in NPG mode.

```
S4148U-Leaf1# show fc switch
Switch Mode : NPG
Switch WWN   : 10:00:e4:f0:04:81:46:cc
```

6.3 show fcoe sessions

The `show fc fcoe sessions` command shows all currently active FCoE sessions on the switch. In this example, four FCoE sessions are active on each switch.

```
S4148U-Leaf1# show fcoe sessions
Enode MAC          Enode Interface    FCF MAC           FCF interface
VLAN              FCoE MAC          FC-ID            PORT WWPN         PORT WNNN
-----
-----
f4:e9:d4:61:c6:6a  Eth 1/1/31        e4:f0:04:81:46:cd Fc 1/1/1
1001  0e:fc:64:01:08:01 01:08:01 20:01:f4:e9:d4:61:c6:6a
20:00:f4:e9:d4:61:c6:6a
f4:e9:d4:62:4a:d0  Eth 1/1/32        e4:f0:04:81:46:cd Fc 1/1/1
1001  0e:fc:64:01:08:02 01:08:02 20:01:f4:e9:d4:62:4a:d0
20:00:f4:e9:d4:62:4a:d0
f4:e9:d4:62:4b:72  Eth 1/1/34        e4:f0:04:81:46:cf Fc 1/1/3
1001  0e:fc:64:01:09:01 01:09:01 20:01:f4:e9:d4:62:4b:72
20:00:f4:e9:d4:62:4b:72
f4:e9:d4:62:4b:ba  Eth 1/1/33        e4:f0:04:81:46:cf Fc 1/1/3
1001  0e:fc:64:01:09:02 01:09:02 20:01:f4:e9:d4:62:4b:ba
20:00:f4:e9:d4:62:4b:ba
```

6.4 show fcoe enode

The `show fcoe enode` command shows the enodes that are currently logged in through the switch. The following example shows the output for the leaf switch #1.

```
S4148U-Leaf1# show fcoe enode
Enode MAC          Enode Interface    VLAN    FCFs    Sessions
-----
-----
f4:e9:d4:61:c6:6a  Eth 1/1/31        1001    1        1
f4:e9:d4:62:4a:d0  Eth 1/1/32        1001    1        1
f4:e9:d4:62:4b:72  Eth 1/1/34        1001    1        1
f4:e9:d4:62:4b:ba  Eth 1/1/33        1001    1        1
```

6.5 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          NPG
Fabric Id            101
Vlan Id              1001
FC-MAP               0xEFC64
Vlan priority        3
FCF Priority          128
FKA-Adv-Period       Enabled,8
Config-State         ACTIVE
Oper-State           UP
=====
Members
fibrechannel1/1/1
fibrechannel1/1/3
ethernet1/1/31
ethernet1/1/32
ethernet1/1/33
ethernet1/1/34
=====
```

A Validated components

Table 1 Leaf switches

Qty	Item	Version
2	Dell EMC S4841U-ON	10.4.2.1

Table 2 Management switches

Qty	Item	Version
2	Dell EMC S3048-ON	10.4.2.1

Table 3 Spine switches

Qty	Item	Version
2	Dell EMC Z9264-ON	10.4.2.1

Table 4 Fibre Channel switches

Qty	Item	Version
2	Brocade 6510	v8.1.0a

Table 5 Storage

Qty	Item	Version
1	Dell EMC Unity 500F	4.3.0.1522077968

Table 6 Servers

Qty	Item	Version
4	Dell EMC PowerEdge R740xd	
	- BIOS	1.6.12
	- iDRAC	3.21.26.22
	- rNDC Intel(R) Gigabit 4P X710/I350 rNDC	18.8.9
	- CNA QLogic 25GE 2P QL41262HxCU-DE Adapter firmware	14.07.50
	- CNA QLogic 25GE 2P QL41262HxCU-DE Adapter ESXi drivers	qedf 1.2.24.6 qedentv 3.7.9.1

B PowerEdge server, Unity storage, and VMware setup

B.1 PowerEdge server configuration

This section details the configuration of the CNAs used to validate the network topology.

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.

B.1.1 Reset server CNA interfaces to factory defaults

Note: Resetting to defaults is only necessary if installed CNA has 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 adapter port.
6. From the **Main Configuration Page**, click the **Default** button followed by **Yes** to load the default settings. Click **OK**.
7. Click **Finish**.
8. Click **Yes** to save changes. Click **OK**.
9. Repeat steps 5-8 for the second FC adapter port.
10. Repeat steps 1-9 for remaining servers.

B.1.2 Determine FCoE CNA port WWPNS

The PowerEdge R740xd server's FCoE CNA World Wide Port Names (WWPNs) are used for FC zone configuration. WWPNs may be determined using the iDRAC as follows:

1. Connect to the first server's iDRAC in a web browser and log in.
2. Select **System > Network Devices**.
3. Under **Network Devices**, click on the CNA. In this example, it is NIC Slot 2. The NIC ports are displayed as shown in Figure 8:

Network Devices

Summary Integrated NIC 1 NIC Slot 1 NIC Slot 4 **NIC Slot 5**

NIC Slot 5: QLogic 25GE 2P QL41262 Adapter

Port Properties

Product Name	QLogic 25GE 2P QL41262 Adapter
Vendor Name	QLogic
Number of Ports	2

Ports and Partitioned Ports

Link Status	Port	Partition	Protocol	Switch Connection ID
+ Up	1	1	NIC, RDMA	3c:2c:30:10:36:00
+ Up	1	2	FCoE	N/A
+ Unknown	1	3	none	N/A
+ Unknown	1	4	none	N/A
+ Up	2	1	NIC, RDMA	3c:2c:30:10:41:00
+ Up	2	2	FCoE	N/A

Figure 8 CNA ports listed in iDRAC

- Under **Ports and Partitioned Ports**, click the icon next to the first port to expand the details as shown:

NIC in Slot 5 Port 1 Partition 2

Link Connection

Link Status	Up
Link Speed	25 Gbps
OS Driver State	Operational
Auto Negotiation	Disabled

MAC Addresses

MAC Addresses	✓ F4:E9:D4:62:4A:D0	Virtual MAC Addresses
FIP	✓ F4:E9:D4:62:4A:D0	Virtual FIP
WWN	✓ 20:00:F4:E9:D4:62:4A:D0	Virtual WWN
WWPN	✓ 20:01:F4:E9:D4:62:4A:D0	Virtual WWPN

Figure 9 WWPN for FCoE CNA port 1

- Record the WWPN, outlined in red in Figure 9. A convenient method is to copy and paste it into a text file. The WWPN is used in the S4148U-ON switch FC zone configuration.
- Repeat steps 4 and 5 for CNA port 2.
- Repeat steps 1-6 for remaining servers.


The FC WWPNs used in this deployment example are shown in Table 7. The Switch column has been added for reference per the cable connections in the SAN topology diagram (Figure 6).

Table 7 Server FCoE CNA port WWPNs

Server	Port	Switch	WWPN
R740xd-1	1	S4148U-Leaf1	20:01:F4:E9:D4:61:C6:6A
R740xd-1	2	S4148U-Leaf2	20:01:F4:E9:D4:61:C6:6B
R740xd-2	1	S4148U-Leaf1	20:01:F4:E9:D4:62:4A:D0
R740xd-2	2	S4148U-Leaf2	20:01:F4:E9:D4:62:4A:D1
R740xd-3	1	S4148U-Leaf1	20:01:F4:E9:D4:62:4B:BA
R740xd-3	2	S4148U-Leaf2	20:01:F4:E9:D4:62:4B:BB
R740xd-4	1	S4148U-Leaf1	20:01:F4:E9:D4:62:4B:72
R740xd-4	2	S4148U-Leaf2	20:01:F4:E9:D4:62:4B:73

B.2 Dell EMC Unity 500F storage array configuration

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

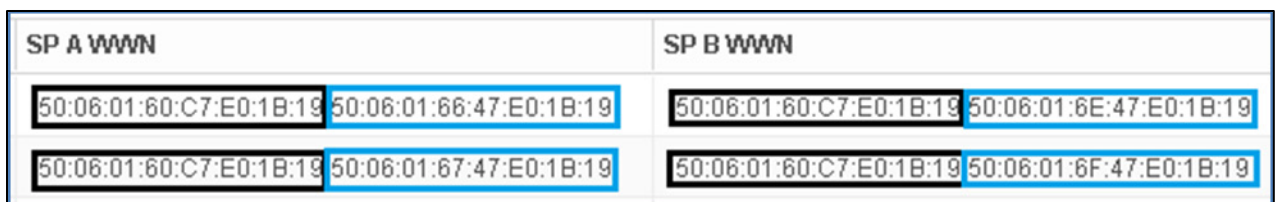
1. Connect to the Unisphere GUI in a web browser and log in. Click the **Settings** icon  near the top right corner of the page.
2. In the left pane of the **Settings** window, select **Access > Fibre Channel**.

The **Fibre Channel Ports** page is displayed as shown in Figure 10. A zoomed-in view of the area inside the red box is shown in Figure 11.



!	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 10 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 11 Zoomed-in view of SP A and SP B WWNs on Fibre Channel Ports page

Two WWNs 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 8. The Switch column has been added based on the physical cable connections shown in Figure 6.

Table 8 Storage array FC adapter WWPNS

Service processor	Port	Switch	WWPN
SP A	0	6510-1 FC	50:06:01:66:47:E0:1B:19
SP A	1	6510-2 FC	50:06:01:67:47:E0:1B:19
SP B	0	6510-1 FC	50:06:01:6E:47:E0:1B:19
SP B	1	6510-2 FC	50:06:01:6F:47:E0:1B:19

B.2.1 Create a storage pool

- Connect to the Unisphere GUI in a web browser and log in.
- In the left pane under **STORAGE**, select **Pools**.
- Click the **+** icon. In the **Create Pool** dialog box, provide a **Name** and click **Next**.
- Select appropriate storage tiers and RAID configuration for the pool. Click **Next**.
- 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**.
- The **Capability Profile Name** section is optional. Click **Next**.
- Review selections on the **Summary** page and click **Finish** to create the pool. Once the **Overall status** shows 100%, click **Close**.
- The pool is displayed on the **STORAGE > Pools** page as shown in Figure 12.

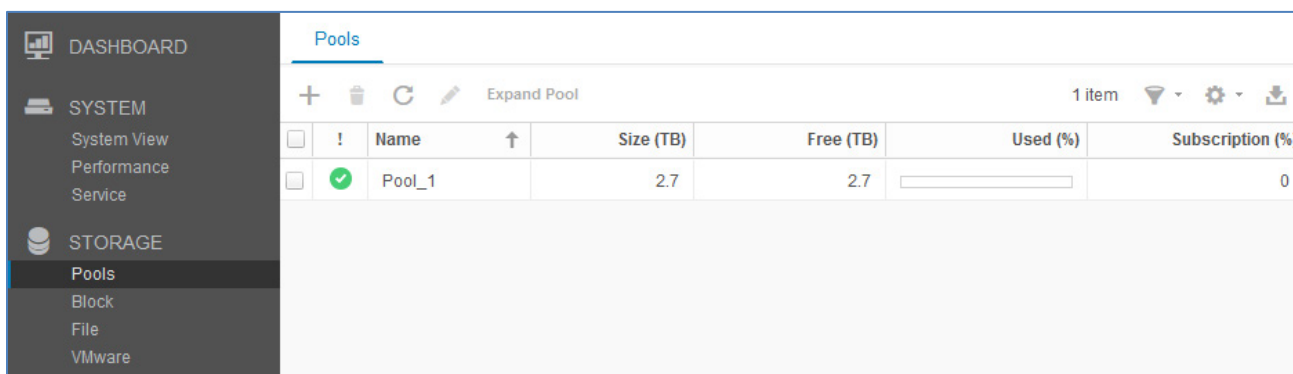


Figure 12 Storage pool created

B.2.2 Add ESXi hosts

- In the Unisphere left pane under **ACCESS**, select **VMware**.
- On the **vCenters** tab, click the **+** icon to open the **Add vCenter** dialog box.
- 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 13.

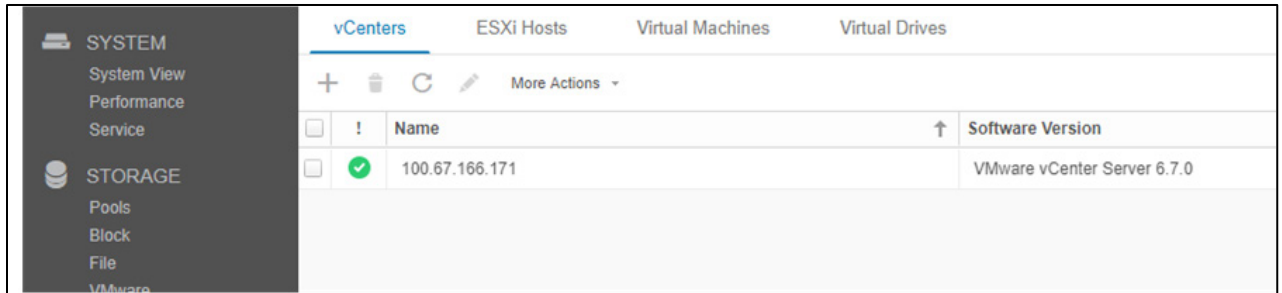


Figure 13 vCenter server added to Unisphere

9. The list of added ESXi hosts is displayed on the **ESXi Hosts** tab, as shown in Figure 14.

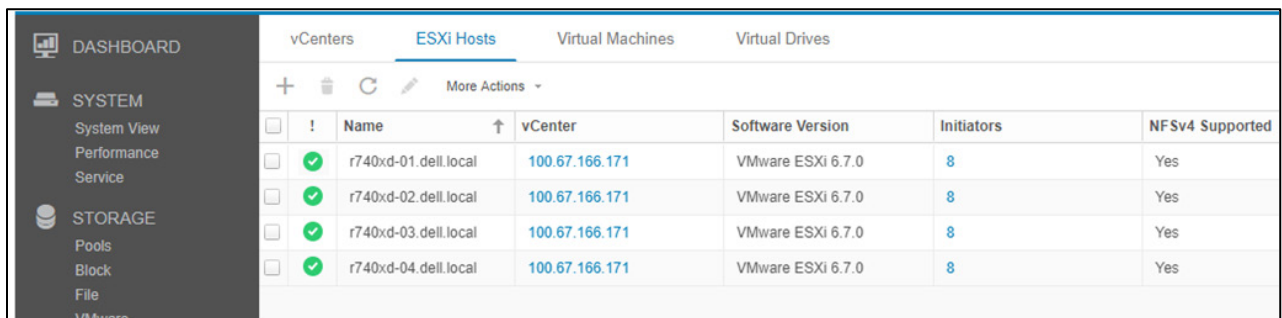


Figure 14 ESXi Hosts added to Unisphere

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.

B.2.3 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 15. In this example, a LUN named FC-80GB that is 80GB in size has been created.

	!	Name	Size (GB)	Allocated (%)	Pool	Thin Clone Base
<input type="checkbox"/>	<input checked="" type="checkbox"/>	FC-80GB	80.0	<input type="text"/>	Pool_1	--

Figure 15 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.

B.3 VMware preparation

B.3.1 VMware ESXi download and installation

Install VMware ESXi 6.7 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 is 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.

B.3.2 Install and configure VMware vCenter Server 6.7 U1

For information on the installation and configuration of vCenter Server, refer to [vSphere Installation and Setup](#). In this deployment, the vCenter Server is located on a separate host within the data center.

B.3.3 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 Management is created and added to the datacenter object, and the four R740xd 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.

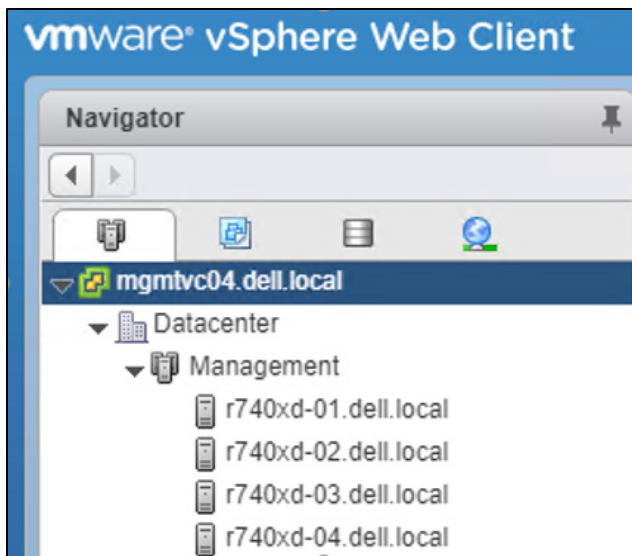



Figure 16 Datacenter and cluster created with ESXi hosts

B.3.4 Configure storage on ESXi hosts

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.

B.3.5 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 FC adapter, vmhba4 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 17.

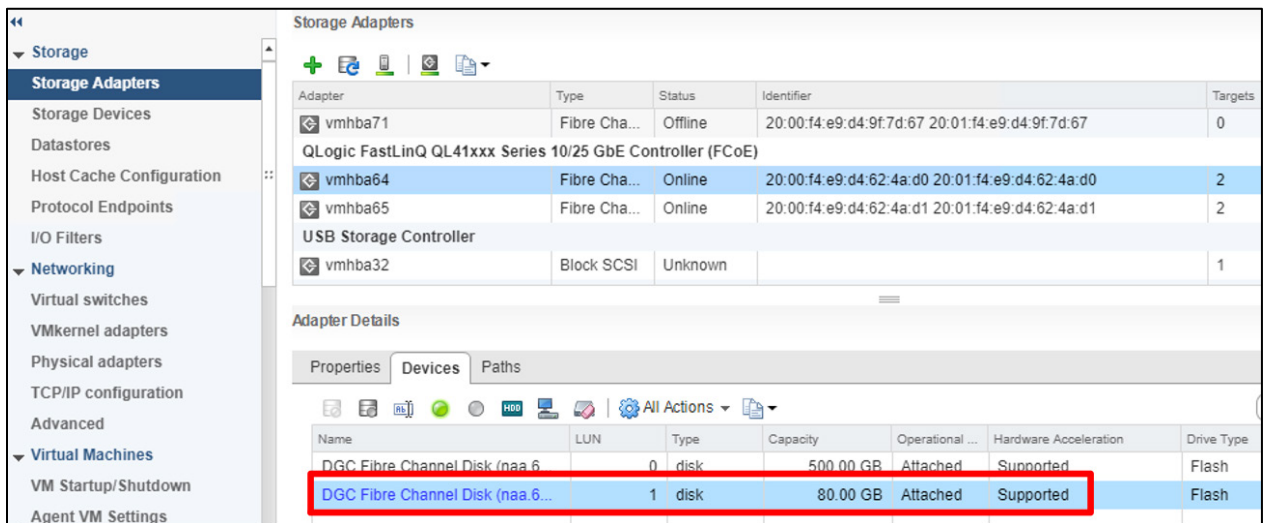


Figure 17 LUN visible to ESXi host

- Repeat for host's second adapter, vmhba5 in this example. The LUN information on the **Adapter Details > Devices** tab is identical to the first adapter.
- Select the first storage adapter, e.g., vmhba4, then select the **Adapter Details > Paths** tab as shown in Figure 18. The target, LUN number (e.g., LUN 0) and path status are shown. The target field includes the two active storage WWPNs connected to vmhba4. The status field is marked either **Active** or **Active (I/O)** for each path.

Runtime Name	Target	LUN	Status
vmhba64:C0:T0:L0	50:06:01:60:c7:e0:1b:19 50:06:0...	0	Active
vmhba64:C0:T1:L0	50:06:01:60:c7:e0:1b:19 50:06:0...	0	Active (I/O)

Figure 18 Adapter Details - Paths tab

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

B.3.6 Create a datastore

A datastore that uses the Unity LUN is created on the ESXi host.

To create the datastore:

- In the vSphere Web Client, go to **Home > Hosts and Clusters**.
- In the **Navigator** pane, right-click on the ESXi host and select **Storage > New Datastore**.
- In the **New Datastore** window, leave the **Type** set to **VMFS** and click **Next**.
- The **Name and device selection** page appears as shown in Figure 19. 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.

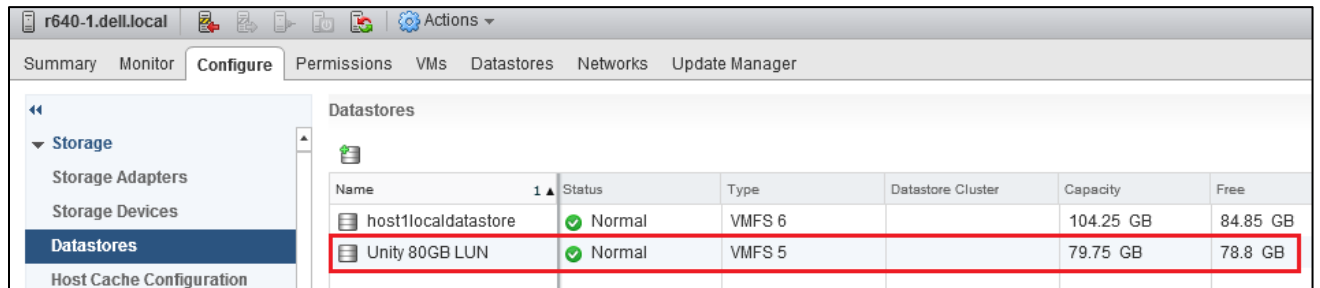
Datastore name:

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 19 Name and device selection page

- Provide a **Datastore name**, e.g., **Unity 80GB LUN**, select the LUN in the list, and click **Next**.
- Select the **VMFS version**. For this guide, it is left at its default setting, **VMFS 5**. Click **Next**.
- 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 20.



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 20 Datastore configured

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

B.3.7 Create a virtual disk

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.

Note: Virtual machine guest operating system deployment steps are not included in this document. For instructions, see the VMware vSphere 6.7 Documentation. Guest operating systems can be any supported by ESXi 6.7. VMs should be deployed before proceeding with this section.

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 looks similar to Figure 21.

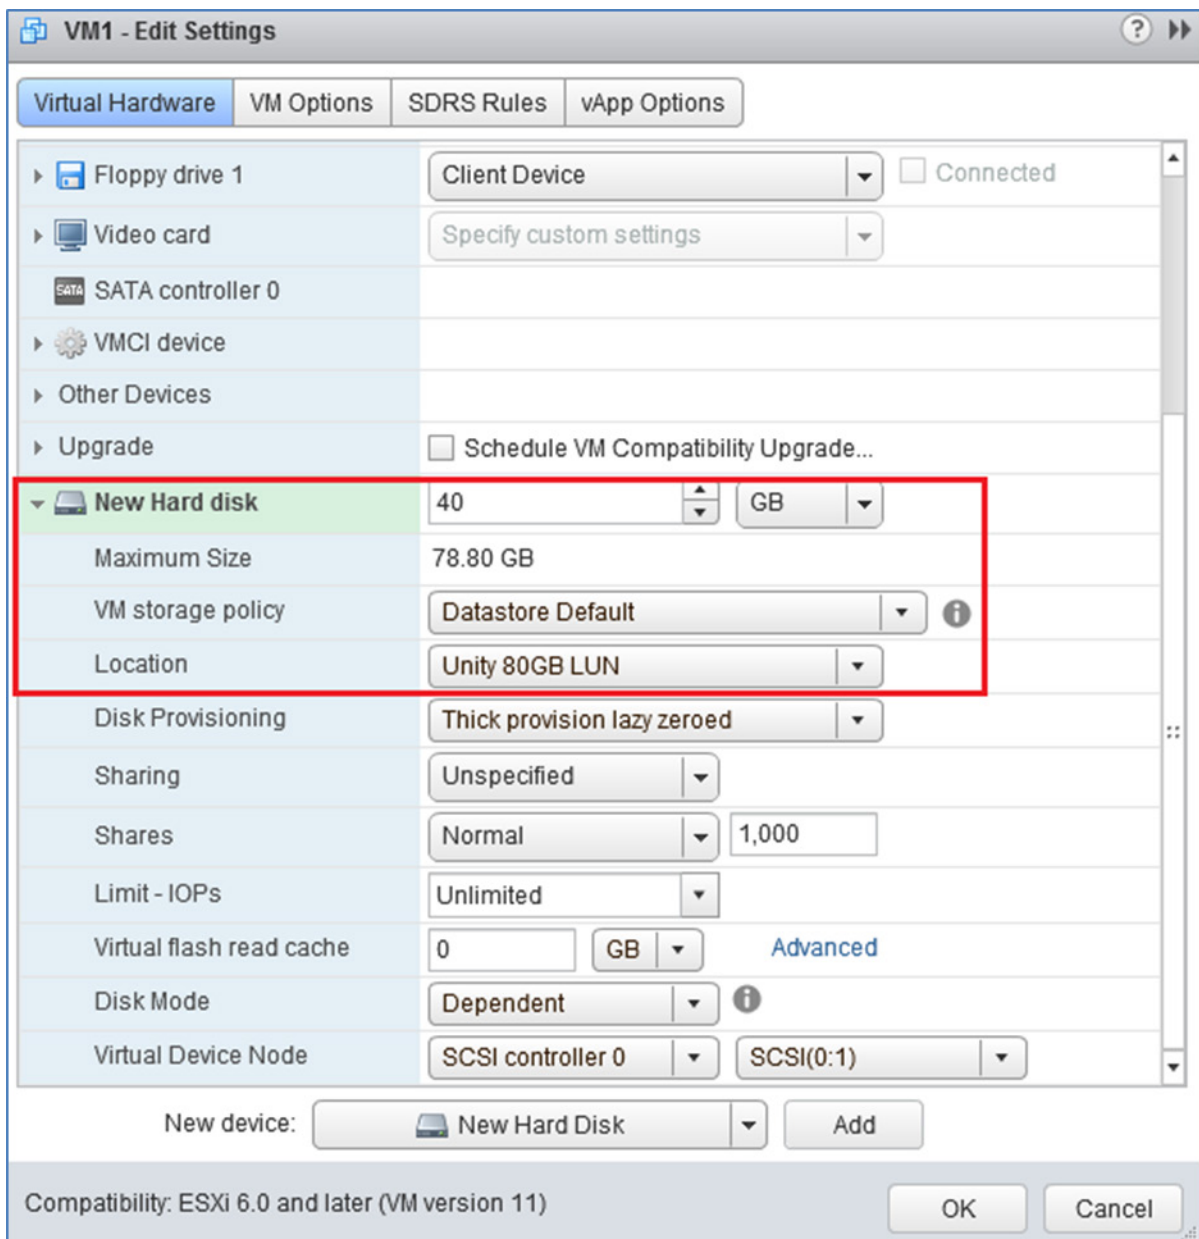


Figure 21 New hard disk configuration options

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.

B.3.8 Configure the virtual disk in Windows Server

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 22).

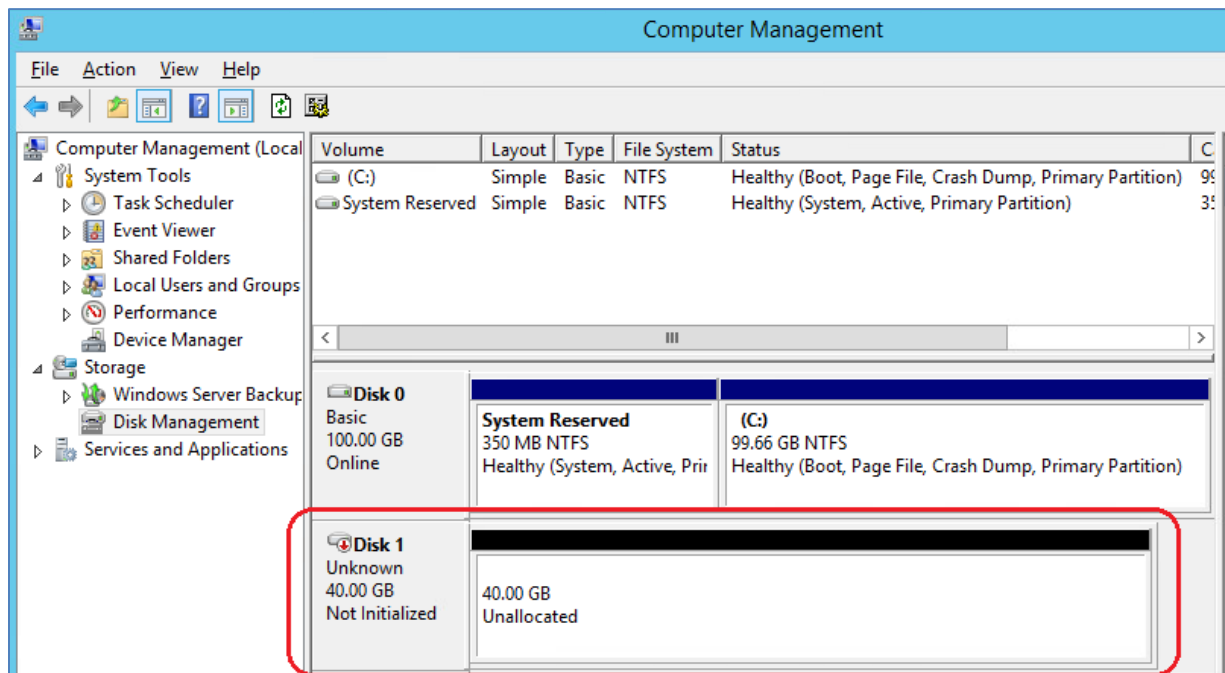


Figure 22 Windows Server Disk Management utility

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.

C Technical resources

[Dell EMC Networking Guides](#)

[Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10EE](#)

[OS10 Enterprise Edition User Guide Release 10.4.2.0](#)

D 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.