

Dell EMC VMware VSAN storage deployment

A step-by-step VMware VSAN storage deployment on Dell EMC FX2 chassis with FC430 compute nodes.

Dell EMC EMEA
February 2017

Revisions

Date	Revision	Description	Authors
February 2017	1.0	Initial Release	Nir Goldshmid, Yory Frenklakh

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2017 Dell Inc. or its subsidiaries. All rights reserved. Dell and the Dell EMC logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

Table of contents

Revisions	2
1 Introduction	6
1.1 Validated System for Virtualization	6
1.1.1 Addressing the need for flexibility	6
1.1.2 Differentiated approach addresses challenges and limitations	7
2 Hardware overview	8
2.1 Dell PowerEdge FX2s enclosure and supported modules	8
2.1.1 PowerEdge FC430 server	9
2.1.2 PowerEdge FD332 storage sled	9
2.1.3 PowerEdge FN410S I/O Module	10
2.2 S4048-ON	10
2.3 N1548	11
3 Network topology	12
3.1 Management network	13
4 Network connections	14
4.1 Production network connections	14
4.1.1 S4048-ON spine switches	14
4.2 Management network connections	15
4.3 IP Address Management	15
4.3.1 VLANs and IP addressing	15
4.4 VRRP	16
4.5 VLT (S4048-ON Spine and FN-410S Leaf switches)	17
4.6 Uplink Failure Detection	17
5 Configure physical switches	18
5.1 Factory default settings	18
5.2 FN410S switch configuration	19
5.3 S4048-ON Spine switch configuration	24
5.4 N1548 management switch configuration	26
5.5 Verify switch configuration	26
5.5.1 S4048-ON Spine Switch	27

5.5.2	FN410S I/O Module	28
6	Prepare Servers	30
6.1	Confirm CPU virtualization is enabled in BIOS	30
6.2	Confirm network adapters are at factory default settings.....	30
6.3	Confirm Firmware versions for the infrastructure.	31
6.4	FD332 Physical Disk Configuration	37
6.5	Confirm storage controllers for VSAN disks are in HBA mode	37
6.6	Install ESXi.....	39
6.7	Configure the ESXi management network connection	39
7	Deploy VMware vCenter Server and add hosts	40
7.1	Deploy VMware vCenter Server.....	40
7.2	Connect to the vSphere web client	42
7.3	Install VMware licenses.....	43
7.4	Create a datacenter object and add hosts	44
7.5	Ensure hosts are configured for NTP.....	45
7.6	Create clusters and add hosts	47
7.7	Information on vSphere standard switches	48
8	Deploy vSphere distributed switches.....	49
8.1	Create a VDS for the vSAN cluster.....	49
8.2	Add distributed port groups.....	50
8.3	Create LACP LAGs	52
8.4	Associate hosts and assign uplinks to LAGs.....	53
8.5	Configure teaming and failover on LAGs.....	56
8.6	Add VMkernel adapters for vMotion and VSAN	57
8.7	Verify VDS configuration	60
8.8	Enable LLDP	61
8.8.1	Enable LLDP on each VDS and view information sent.....	61
8.8.2	View LLDP information received from physical switch	62
9	Configure a VSAN datastore in each cluster	63
9.1	VSAN Overview.....	63
9.2	Configure VSAN	64

9.3	Verify VSAN configuration.....	67
9.4	Check VSAN health and resolve issues.....	68
9.4.1	Failure: Virtual SAN HCL DB up-to-date	68
9.4.2	Warning: Controller Driver / Controller Release Support	69
9.4.3	Warning: Performance Service / Stats DB object.....	69
9.5	Verify IGMP snooping functionality	69
10	Scaling guidance	70
10.1	VSAN sizing	70
10.2	Port count and oversubscription.....	70
A	Dell EMC validated hardware and components.....	71
A.1	Switches.....	71
A.2	PowerEdge FX2s chassis and components	71
B	Dell EMC validated software and required licenses	72
B.1	Software	72
B.2	Licenses	72
C	Technical support and resources	73
C.1	Dell EMC product manuals and technical guides	73
C.2	VMware product manuals and technical guides	73
D	Support and Feedback	74

1 Introduction

This guide covers a VMWare VSAN deployment for the Small data center based on the Dell EMC FX2 System for Virtualization.

The goal of this guide is to enable a system administrator or engineer with traditional networking and ESXi experience to build a scalable VSAN virtual storage using the Dell EMC FX2 System for Virtualization hardware and software outlined in this guide.

This document provides a best practice with configuration steps for all components in the topology. It includes step-by-step configuration of a virtual network and storage. It also includes steps to deploy ESXi on PowerEdge servers, vSAN and deployment of a vSphere vCenter Server Appliance.

Note: See the appendices for product versions validated.

1.1 Validated System for Virtualization

The Dell EMC Validated System for Virtualization is the industry's most flexible converged system to date, with choice in building blocks of compute, storage and networking tested and validated to integrate and operate together in support of a virtualized environment. The system incorporates a wide range of form factors, technology choices and deployment options, right-sized to fit each customer's needs. A fully-validated system can be configured, quoted and ordered in minutes, while automated lifecycle management tools allow customers to easily deploy, scale, and update the system.

1.1.1 Addressing the need for flexibility

With increasing business demands and decreasing IT budgets, customers face unprecedented pressures to improve efficiency and lower costs. The current operational model of delivering IT services, which involves procuring technology from best of breed technology providers and managing them in silos, proves to not only be time consuming but problematic. In this approach, customers are typically burdened to manually make design decisions, validate various components, set up and configure components and manage the environment in an ongoing fashion by engaging multiple vendors for assistance. Across the end-to-end infrastructure lifecycle, these elements increase complexity and cost for customers.

Existing integrated solutions that aim to solve these challenges are either pre-integrated and prepackaged offers that optimize time to production and simplify ongoing operations, with customers making a tradeoff on flexibility and choice, or traditional reference architectures that provide some degree of flexibility but do not offer manageability or scalability benefits.

The Dell EMC Validated System for Virtualization bridges this gap by offering a tested and validated integrated system that is highly flexible, scalable, and driven using end-to-end automation throughout the infrastructure lifecycle.

1.1.2 Differentiated approach addresses challenges and limitations

To provide IT services faster, while lowering costs and streamlining operations, Dell EMC engineered the Validated System for Virtualization. This groundbreaking system enables you to achieve greater operational efficiencies and savings, and unparalleled management simplicity, by giving you more power than ever to define and design it.

The system can be deployed with options ranging from “do-it-yourself” using a deployment guide, a system integrated on-site by Dell EMC or by using your own integration vendor.

The Dell EMC Validated System for Virtualization is:

- Built on our best-of-breed products that are designed for virtualization across the ecosystem. By offering various design choices and guidance on choosing the right components, the Dell EMC Validated System for Virtualization takes the guesswork out of solution design and reduces the enormous time it takes to procure, validate, and integrate components. They can be designed to start small, based on the customer’s initial requirements, and grow, based on customer’s ongoing requirements, which reduces the initial investment required for infrastructure deployment.
- Tested, validated, and fully integrated, yet flexible enough to be tailored for your organization, removing risk and accelerating your time to value.
- Delivered with Dell EMC’s global reach, exceptional execution and delivery, providing consistent deployment, management, and maintenance in every region of the world.
- Delivered with a single point of support for the complete system including hardware and software through Dell ProSupport Plus. ProSupport Plus resolves issues faster when they occur and reduces the risk of severe issues and outages.

More information about the Dell EMC Validated System for Virtualization is available [here](#).

1.2 Typographical conventions

This document uses the following typographical conventions:

Monospace text

Command Line Interface (CLI) examples

Bold monospace text

Commands entered at the CLI prompt

Italic monospace text Variables in CLI examples

2 Hardware overview

While the Dell EMC Validated System for Virtualization has flexibility and choice across servers, storage and networking, this guide is focused on a single instance of the system. 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 PowerEdge FX2s enclosure and supported modules

The PowerEdge FX2s enclosure is a 2-rack unit (RU) computing platform. It has capacity for two FC830 full-width servers, four FC630 half-width servers or eight FC430 quarter-width servers. The enclosure is also available with a combination of servers and storage sleds. The FX2s enclosure used in this guide contains four FC430 servers (Section 2.1.1) and two FD332 storage sleds (Section 2.1.2).



Figure 1 Dell PowerEdge FX2s (front) with four PowerEdge FC430 servers and two FD332 storage sleds

The back of the FX2s enclosure includes two I/O networking modules (IOMs) and eight PCIe expansion slots.



Figure 2 Dell PowerEdge FX2s (back) with two PowerEdge FN410S IOMs installed

2.1.1 PowerEdge FC430 server

The PowerEdge FC430 server is a quarter-width, 2 socket server. Four FC430 servers in the top half of the FX2s enclosure combine with the FD332 storage sleds to form the compute cluster for this deployment.



Figure 3 PowerEdge FC430

2.1.2 PowerEdge FD332 storage sled

The PowerEdge FD332 is a half-width, direct-attached storage sled with up to 16 drives. It combines with FC-series servers to build flexible storage solutions. This deployment includes two FD332 storage sleds installed in the bottom half of the FX2s enclosure.



Figure 4 PowerEdge FD332

2.1.3 PowerEdge FN410S I/O Module

The PowerEdge FN410S IOM is a multilayer switch with eight internal, server-facing ports and four external, 10GbE SFP+ ports. Two FN410S IOMs installed in the FX2s enclosure provide fault tolerance.



Figure 5 PowerEdge FN410S

2.2 S4048-ON

The S4048-ON is a 1RU, layer 2/3 switch with forty-eight 10GbE SFP+ ports and six 40GbE QSFP+ ports. Two S4048-ON switches are used as spine switches in the leaf-spine topology covered in this guide.



Figure 6 Dell Networking S4048-ON

2.3 N1548

The N1548 is a 1RU, layer 2+ switch with forty-eight 1GbE ports and four 10GbE SFP+ ports. One N1548 switch used as OOB Management network switch for deployments covered in this guide.



3 Network topology

This section provides an overview of the network topology used in this deployment.

On the production network, a spine and leaf topology is used for performance and scalability. Each Dell FX2 chassis include two FN-410S IO-Modules for network connectivity. Two spine switches (S4048-ONs) are used for solution flexibility, redundancy and increased performance. Dell Virtual-Link Trunking (VLT) connects each pair of spine or leaf switches.

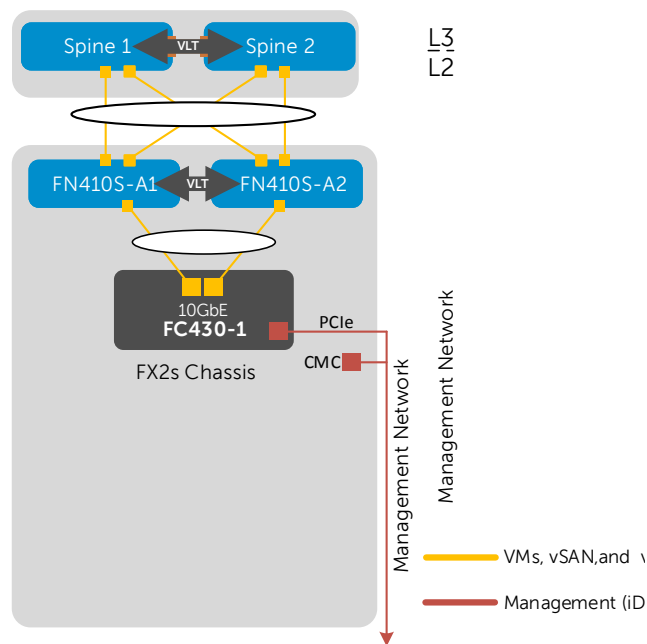


Figure 7 Network topology

3.1 Management network

This guide uses a single management traffic network that is isolated from the production network. An N1548 switch installed in the rack provides connectivity to the management network.

Each FX2s chassis has four 1GbE add-in PCIe network adapters (each connected internally to an FC430 server) for ESXi host management and a built-in CMC for OOB management.

These devices, in addition to the S4048-ON switch management ports, are all connected to the management network as shown in Figure 8.

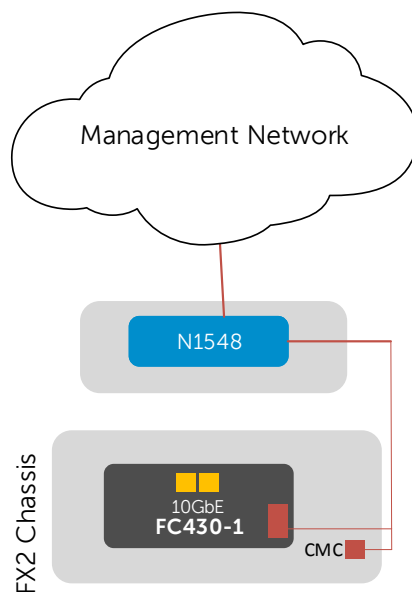


Figure 8 Physical layout of iDRAC, CMC and ESXi management interfaces

4 Network connections

This section details the physical network connections.

4.1 Production network connections

4.1.1 S4048-ON spine switches

Figure 9 shows the network connections from the FN410S switches in the FX2s chassis to Spine 1 and Spine 2. The Spine switches are VLT peers and two FN410S ports connect to each leaf switch. The FN410S switches are also VLT peers. The two FN410S ports functions as the VLTi (VLT interconnect) between the switches.

Inside the FX2s chassis (not shown), four PowerEdge FC430 servers connect via QLogic 57810 dual-port network adapters to FN410S-A1 and A2. For each server, one link connects internally to FN410S-A1 and the other connects to FN410S-A2.




Figure 9 Production network connections (S4048-ON Spine switches)

4.2 Management network connections

These connections are used for non-production, management traffic.

For management traffic, four FC430 servers are connected to an N1548 switch via four Intel I350-T add-in adapters in the FX2s chassis. The FX2s CMC is also connected for OOB management.

For scalability, N1548 ports 1–8 are allocated for FX2s chassis CMC ports. Ports 17–48 are available for Intel I350-t management interfaces (four for each FX2s). This allows up to eight FX2s chassis per N1548 switch.




Figure 10 Management network connections

4.3 IP Address Management

Proper IP Address Management (IPAM) is critical before deploying a spine and leaf topology. This section covers the IP addressing used on the physical network in this guide.

4.3.1 VLANs and IP addressing

Table 1 outlines the VLAN IDs, network and gateway addresses used. The "x" in each network address is replaced by the customer-provided number. The gateway address is the Virtual Router Redundancy Protocol (VRRP) group address, described in the next section. The VLANs and networks are advertised through customer preferred dynamic routing protocol in case of S4048-on spine switches deployment.

Table 1 VLAN and network examples

VLAN ID	Network	Gateway	Used For
22	10.22.x.0/24	10.22.x.1	vMotion
44	10.44.x.0/24	10.44.x.1	VSAN
10	10.10.x.0/24	10.10.x.1	VMs

4.4 VRRP

VRRP is designed to eliminate a single point of failure in a routed network. VRRP is used to create a virtual router which is an abstraction of the two physical spine switches. The virtual router is assigned an IP address that is used as the gateway address by the compute nodes. In the event that one of the spine switches fails, the remaining leaf acts as the gateway until the failed unit recovers.

As illustrated in Figure 11, Node 1 is participating in VLAN 10. The node has an IP address of 10.10.1.1. The node's gateway address is set to 10.10.1.1. This is the Virtual IP (VIP) provided by the VRRP instance running between spine switches 1 and 2.




Figure 11 VRRP configuration example – VLAN 10

A VRRP instance is created for each VLAN in each pair of leaf switches at the top of each rack.

0 shows the VRRP IP addressing scheme for VLANs as an example.

Table 2 VRRP interface configuration for VLANs

VLAN	First Leaf VLAN IP	Second Leaf VLAN IP	Virtual IP
10	10.10.1.2/24	10.10.1.3/24	10.10.1.1
22	10.22.1.2/24	10.22.1.3/24	10.22.1.1
44	10.44.1.3/24	10.44.1.3/24	10.44.1.1

4.5 VLT (S4048-ON Spine and FN-410S Leaf switches)

A pair of spine switches at the top of rack and pair of a FN-IOM provides redundancy. These switches' configurations include the Dell Networking Virtual Link Trunking (VLT) feature.

VLT reduces the role of spanning tree protocols (STPs) by allowing link aggregation group (LAG) terminations on two separate switches and supporting a loop-free topology. VLT provides Layer 2 multipathing and load-balances traffic where alternative paths exist. Virtual Link Trunking offers the following additional benefits:

- Allows a single device to use a LAG across two upstream devices
- Eliminates STP-blocked ports
- Uses all available uplink bandwidth
- Provides fast convergence if either the link or a device fails
- Provides link-level resiliency
- Assures high availability

4.6 Uplink Failure Detection

If a leaf switch loses connectivity to the spine layer, the attached hosts continue to send traffic without a direct path to the destination. The VLTi link to the peer leaf switch handles traffic during such a network outage, but this is not considered a best practice.

Dell EMC recommends enabling Uplink Failure Detection (UFD), which detects the loss of upstream connectivity. An uplink-state group is configured on each leaf switch, which creates an association between the spine uplinks and the downlink interfaces. An uplink-state group is configured on each FN410S.

In the event of an uplink failure, UFD automatically shuts down the corresponding downstream interfaces. This propagates down to the hosts attached to the FN410S switch. The host then uses its remaining Link Aggregation Control Protocol (LACP) port member to continue sending traffic across the leaf-spine network.

5 Configure physical switches

Note:

do not configure the networking on the FN switches before the servers have been fully firmware updated.

This section contains switch configuration details with explanations for one switch in each major role on the production network. This chapter details the following switches:

- FN410S-A1
- S4048-ON: Spine 1

The remaining switches use configurations very similar to one of the three configurations above, with the applicable switches specified in each section.

Notes:

MTU - The MTU is set to 9216 bytes on all switches in the production network in this guide. vSAN require setting the MTU to the 9000 bytes on all switches that will handle vSAN traffic on your network.

Port Channel Numbering – LACP port channel numbers may be any number in the range 1-128.

5.1 Factory default settings

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

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

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

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

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

```
% Warning: The bmp process will stop ...
```

```
Dell>
```

The switch is now ready for configuration.

5.2 FN410S switch configuration

The solution includes a PowerEdge FX2s chassis with four FC430 servers and two FN410S switches.

Each FC430 server has an LACP-enabled port channel connected to internal interfaces of each FN410S. For clarity, only port channel 1 (for server FC430-1) is shown in Figure 12. The remaining port channels are numbered 2-4.

The two FN410S switches are configured as VLT peers. Three of the four FN410S external interfaces, tengigabitethernet 0/9-10, are configured in port channel 10 which is connected to spine switches 1 and 2. The tengigabitethernet 0/11-12 external interfaces, is used as the VLT interconnect between FN410S-A1 and FN410S-A2.

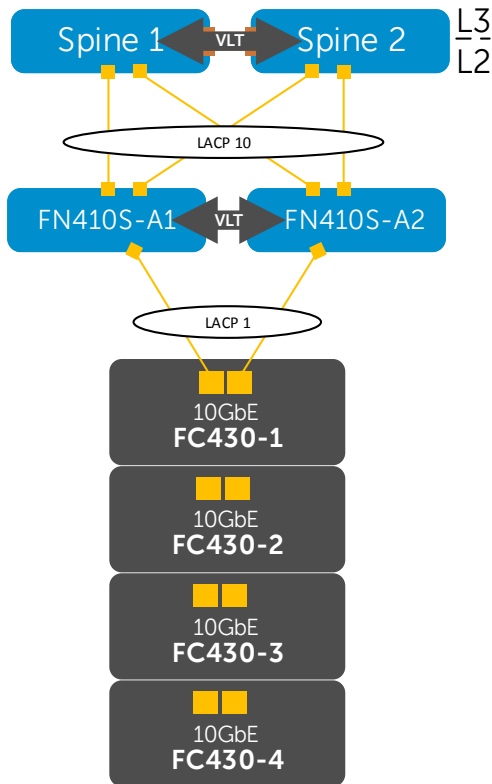


Figure 12 FN410S network connections (internal LACP connections to FC430-2 through 4 not shown)

The following section outlines the configuration commands issued to the FN410S switches. The switches start at their factory default settings per Section 5.1.

After FN410S switches boot to their default settings, place them in full-switch mode as follows:

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

% You are about to configure the Full Switch Mode.
Please reload to effect the changes
```

```
Dell(conf) #do reload
```

```
System configuration has been modified. Save? [yes/no]: yes  
Proceed with reload [confirm yes/no]: yes
```

After FN410S switches boot to full-switch mode, enter the following commands to configure the FN410S-A1.

Note: Ensure FN410S switches have been placed in full-switch mode before proceeding. The following configuration details are specific to switch FN410S-A1. The configuration for FN410S-A2 is similar. See the FN410S-A1.txt and FN410S-A2.txt attachments.

Initial configuration involves setting the hostname, enabling Link Layer Discovery Protocol (LLDP) and disabling Data Center Bridging (DCB). LLDP is useful for troubleshooting (see Section 8.8). DCB is enabled by default on FN410S but is not used in this environment.

Enable Internet Group Messaging Protocol (IGMP) snooping for VSAN traffic. Finally, configure the management interface and default gateway.

Note: IGMP snooping is enabled, but IGMP snooping querier is not enabled on the FN410S switches. The querier role is enabled on the S4048-ON or N4032F spine switches upstream.

```
enable  
configure  
  
hostname FN410S-A1  
protocol lldp  
advertise management-tlv management-address system-description system-  
name  
advertise interface-port-desc  
no dcb enable  
ip igmp snooping enable  
  
interface ManagementEthernet 0/0  
ip address 10.0.0.15/24  
no shutdown  
  
management route 0.0.0.0/0 10.0.0.138
```

Next, the VLT interface between the two switches is configured. In this configuration, interface tengigabitethernet 0/11-12 is used for the VLTi interface. It is added to static port-channel 127. The backup destination is the management IP address of the VLT peer switch, FN410S-A2. The VLT unit-id is set to 0 (and is set to 1 on FN410S-A2).

```

interface port-channel 127
description VLTi
mtu 9216
channel-member tengigabitethernet 0/11
channel-member tengigabitethernet 0/12
no shutdown

interface tengigabitethernet 0/11
description VLTi
no shutdown

interface tengigabitethernet 0/12
description VLTi
no shutdown

vlt domain 127
peer-link port-channel 127
back-up destination 10.0.0.16
unit-id 0

```

The upstream interfaces to the two spine switches are configured in this section. External interfaces tengigabitethernet 0/9-10 are used and placed in LACP-enabled port channel 10. The port channel is configured for VLT and jumbo frames are enabled for vSAN traffic.

```

interface range tengigabitethernet 0/9-10
description To LAN switch te 1/0/9,1/0/11
mtu 9216
port-channel-protocol LACP
port-channel 10 mode active
no shutdown

interface port-channel 10
description To LAN switch te 1/0/9,1/0/11
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 10
no shutdown

```

The downstream interfaces are configured in the next set of commands. Each internal interface is added to a numerically corresponding port channel. The port channels are configured for VLT and jumbo frames are enabled on all interfaces for vSAN traffic.

```

interface tengigabitethernet 0/1
description To FC430-1

```

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

interface port-channel 1
description To FC430-1
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 1
no shutdown

interface tengigabitethernet 0/2
description To FC430-2
mtu 9216
port-channel-protocol LACP
port-channel 2 mode active
no shutdown

interface port-channel 2
description To FC430-2
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 2
no shutdown

interface tengigabitethernet 0/3
description To FC430-3
mtu 9216
port-channel-protocol LACP
port-channel 3 mode active
no shutdown

interface port-channel 3
description To FC430-3
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 3
no shutdown

interface tengigabitethernet 0/4
description To FC430-4
```

```
mtu 9216
port-channel-protocol LACP
port-channel 4 mode active
no shutdown
```

```
interface port-channel 4
description To FC430-4
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 4
```

Finally, the three required VLAN interfaces are created. All downstream and upstream port channels are tagged in each VLAN.

```
interface Vlan 10
description VMs
mtu 9216
tagged Port-channel 1-4,10
no shutdown
```

```
interface Vlan 22
description vMotion
mtu 9216
tagged Port-channel 1-4,10
no shutdown
```

```
interface Vlan 44
description VSAN
mtu 9216
tagged Port-channel 1-4,10
no shutdown
```

UFD is configured. This shuts the downstream interfaces if all uplinks fail. The hosts attached to the switch use the remaining LACP port member to continue sending traffic across the fabric.

```
uplink-state-group 1
description Disable downstream ports in event all uplinks fail
downstream TenGigabitEthernet 0/1-8
upstream TenGigabitEthernet 0/9-10
```

Save the configuration.

```
end
write
```

5.3 S4048-ON Spine switch configuration

Each S4048-ON Spine switch has an LACP-enabled port channel connected to each of the downstream leaf FN410S switches.

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

Note: The following configuration details are specific to Spine 1. The remaining spine switch 2, is similar. Complete configuration details for all spine switches are provided in the attachments named spine1.txt and spine2.txt.

For VSAN traffic, IGMP snooping and IGMP snooping querier are enabled on spine switches. IGMP snooping is enabled globally, and IGMP querier is enabled on VLAN 44.

Initial configuration involves setting the hostname, and enabling LLDP and IGMP snooping. LLDP is useful for troubleshooting (see Section 8.8). IGMP snooping is enabled for VSAN traffic. Finally, the management interface and default gateway are configured.

```
enable
configure

hostname Spine-1
protocol lldp
advertise management-tlv management-address system-description system-
name
advertise interface-port-desc
ip igmp snooping enable

interface ManagementEthernet 1/1
ip address 10.0.0.17/24
no shutdown

management route 0.0.0.0/0 10.0.0.138
```

Next, the VLT interfaces between Spine-1 and Spine-2 are configured. In this configuration, interfaces fortyGigE 1/53-54 are used for the VLT interconnect. They are added to static port-channel 127. The backup destination is the management IP address of the VLT peer switch, Spine2-2.

```
interface port-channel 127
description VLTi
mtu 9216
channel-member fortyGigE 1/53 - 1/54
no shutdown
```



```

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

vlt domain 127
peer-link port-channel 127
back-up destination 10.0.0.18
unit-id 0
exit

```

The downstream interfaces, to the FN430S, are configured in the next set of commands. Each interface is added to a numerically corresponding port channel. The port channels are configured for VLT and jumbo frames are enabled on all interfaces for vSAN traffic.

```

interface range tengigabitethernet 1/41 - 1/43
description To FN410S-A1/A2
mtu 9216
port-channel-protocol LACP
port-channel 10 mode active
no shutdown

interface Port-channel 10
description To FN410S-A1/A2
mtu 9216
portmode hybrid
switchport
vlt-peer-lag port-channel 10
no shutdown

```

The three required VLAN interfaces are created. All downstream port channels are tagged in each VLAN. Each interface is assigned to a VRRP group and a VRRP address is assigned. VRRP priority is set to 254 to make this switch the master. (On the VRRP peer switch, priority is set to 1).

```

interface Vlan 22
description vMotion
ip address 10.22.1.2/24
mtu 9216
tagged Port-channel 10
vrrp-group 22
description vMotion
priority 254
virtual-address 10.22.1.1
no shutdown

interface Vlan 44
description VSAN

```

```
ip address 10.44.1.2/24
mtu 9216
tagged Port-channel 10
ip igmp snooping querier
vrrp-group 44
description VSAN
priority 254
virtual-address 10.44.1.1
no shutdown

interface Vlan 10
description VMs
ip address 10.10.1.2/24
mtu 9216
tagged Port-channel 10
vrrp-group 10
description VMs
priority 254
virtual-address 10.10.1.1
no shutdown
```

Save the configuration.

```
end
write
```

5.4 N1548 management switch configuration

For the N1524 management switches, all ports used are in layer 2 mode and are in the default VLAN. No additional configuration is required.

5.5 Verify switch configuration

The following sections show commands and output to verify switches are configured and connected properly. Except where there are key differences, only output from one spine switch and one FN410S switch is shown to avoid repetition. Output from remaining devices will be similar.

5.5.1 S4048-ON Spine Switch

5.5.1.1 show vlt brief

The Inter-chassis link (ICL) Link Status, Heart Beat Status, and VLT Peer Status must all be up. The role for one switch in the VLT pair will be primary and its peer switch (not shown) will be assigned the secondary role.

```
Spine-1#show vlt brief
```

```
VLT Domain Brief
-----
Domain ID:                127
Role:                     Primary
Role Priority:             32768
ICL Link Status:          Up
HeartBeat Status:         Up
VLT Peer Status:          Up
Local Unit Id:            0
Version:                  6(7)
Local System MAC address:  f4:8e:38:20:37:29
Remote System MAC address: f4:8e:38:20:54:29
Remote system version:    6(7)
Delay-Restore timer:      90 seconds
Delay-Restore Abort Threshold: 60 seconds
Peer-Routing :            Disabled
Peer-Routing-Timeout timer: 0 seconds
Multicast peer-routing timeout: 150 seconds
```

5.5.1.2 show vlt detail

On spine switches 1 and 2 downstream port channel 10 is up because they are connected to properly configured FN410S switches. VLANs 1, 10, 22 and 44 are active.

```
Spine-1#show vlt detail
```

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
10	10	UP	UP	1, 10, 22, 44

5.5.1.3 show vrrp brief

The output from the `show vrrp brief` command should be similar to that shown below. The priority (Pri column) of the master router in the pair is 254 and the backup router (not shown) is assigned priority 1.

```
Spine-1#show vrrp brief
```

Interface	Group	Pri	Pre	State	Master	addr	Virtual	addr(s)	Description
Vl 10	IPv4 10	254	Y	Master	10.55.1.2	10.55.1.1			VMs
Vl 22	IPv4 22	254	Y	Master	10.22.1.2	10.22.1.1			vMotion
Vl 44	IPv4 44	254	Y	Master	10.44.1.2	10.44.1.1			VSAN

5.5.2 FN410S I/O Module

5.5.2.1 show vlt brief

Like the S4048-ON switches above, the ICL Link Status, Heart Beat Status, and VLT Peer Status must all be up. One switch is primary and the peer (not shown) is the secondary.

```
FN410S-A1#show vlt brief
```

```
VLT Domain Brief
```

```
-----
```

```
Domain ID: 127
Role: Primary
Role Priority: 32768
ICL Link Status: Up
HeartBeat Status: Up
VLT Peer Status: Up
Local Unit Id: 0
Version: 6(7)
Local System MAC address: f8:b1:56:6e:fc:5b
Remote System MAC address: f8:b1:56:76:b9:b5
Remote system version: 6(7)
Delay-Restore timer: 90 seconds
Delay-Restore Abort Threshold: 60 seconds
Peer-Routing : Disabled
Peer-Routing-Timeout timer: 0 seconds
Multicast peer-routing timeout: 150 seconds
```

5.5.2.2 show vlt detail

Downstream LAGs (port channels 1-4) are down until LAGs are configured on the directly connected ESXi hosts running on the FC430 servers. This is covered in Section 8.4.

The upstream LAG (port channel 10) is currently up because it is connected to properly configured spine switches (Spine 1 and Spine 2).

VLANs 1, 10, 22 and 44 are active on all LAGs.

```
FN410S-A1#show vlt detail
```

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
1	1	DOWN	DOWN	1, 10, 22, 44
2	2	DOWN	DOWN	1, 10, 22, 44
3	3	DOWN	DOWN	1, 10, 22, 44
4	4	DOWN	DOWN	1, 10, 22, 44
10	10	UP	UP	1, 10, 22, 44

6 Prepare Servers

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

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

6.1 Confirm CPU virtualization is enabled in BIOS

Note: CPU virtualization is typically enabled by default in PowerEdge server BIOS. These steps are provided for reference in case this required feature has been disabled.

1. Connect to the iDRAC in a web browser and launch the virtual console.
2. In the virtual console, from the **Next Boot** menu, select **BIOS Setup**.
3. Reboot the server.
4. From the System Setup Main Menu, select System BIOS, and then select Processor Settings.
5. Verify Virtualization Technology is set to Enabled.
6. To save the settings, click **Back**, **Finish**, and **Yes** if prompted to save changes.
7. If resetting network adapters to defaults, proceed to step 4, **System Setup Main Menu**, in the next section. Otherwise, reboot the server.

6.2 Confirm network adapters are at factory default settings

Note: These steps are only necessary if installed network adapters have been modified from their factory default settings.

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

6.3 Confirm Firmware versions for the infrastructure.

Several ways to update the FX2 infrastructure, update for each server via Lifecycle controller or via the Chassis Management Controller (CMC – using Network shares or update files).

Single server firmware update:

Press F10 and enter to the Lifecycle controller in boot

Go to get the latest firmware.



Figure 13 LifeCycle controller

Choose your update option – *in this example we use FTP Server*

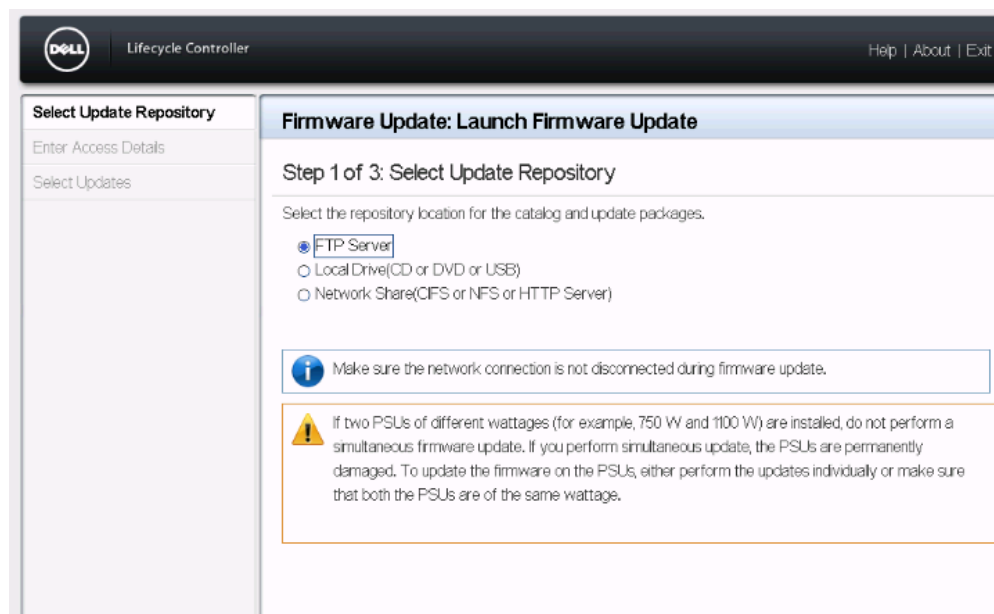


Figure 14 Firmware update

In some cases when the network interfaces are not connected or configured you will need to set an IP address to one of your management NICs, in this case you will click “Yes” to configure the NIC.

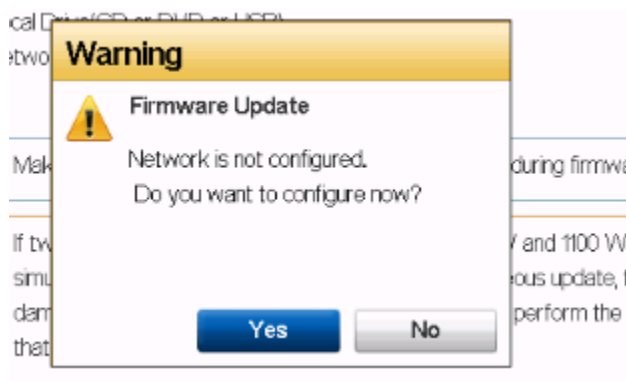


Figure 15 Firmware update warning

Configure the relevant IP address via Static IP or DHCP and click Finish

Dell Lifecycle Controller Help | About | Exit

Settings

Network Settings

Use Network Settings to select and configure the Lifecycle Controller Network Interface Card (NIC).

NIC Card

QLogic 57xx/578xx 10 Gigabit Ethernet (Embedded NIC 1)

IPv4 Network Settings

Select the IP address configuration mode.

IP Address Source Static IP

IP Address 10.0.0.45

Subnet Mask 255.255.255.0

Default Gateway 10.0.0.138

DNS Address 10.0.0.101

IPv6 Network Settings

Select the IP address configuration mode.

IP Address Source No Configuration

IP Address

Prefix Length

PowerEdge FC430
Service Tag: 6JRG2

Cancel Back Finish

Figure 16 Network settings

In FTP mode, add the following address: FTP.DELL.COM and click Next.

For more information, please refer to the following video: <https://www.youtube.com/watch?v=ImWYrS4RGIY>

The screenshot shows the Dell Lifecycle Controller interface for launching a firmware update. The left sidebar contains navigation options: 'Select Update Repository' (checked), 'Enter Access Details' (selected), and 'Select Updates'. The main content area is titled 'Firmware Update: Launch Firmware Update' and 'Step 2 of 3: Enter Access Details'. It is divided into two sections: 'FTP Server Settings' and 'Proxy Settings'. Under 'FTP Server Settings', there are input fields for 'Address' (containing 'ftp.dell.com'), 'User Name', 'Password', and 'File Path or Update Package Path'. Under 'Proxy Settings', there is a checkbox for 'Enable Settings' which is currently unchecked, followed by input fields for 'Server', 'Port', 'User Name', and 'Password', and a dropdown menu for 'Type' set to 'HTTP'. At the bottom of the form is a yellow button labeled 'Test Network Connection'.

Figure 17 Launch Firmware Update

Update all the available system firmware, in this scenario must important for vSAN will be the PERC FD332 controller, BP13G+EXP 0:3 backplane

For more information, please refer to the following video:

<https://www.youtube.com/watch?v=cLj2UVIFBFo>

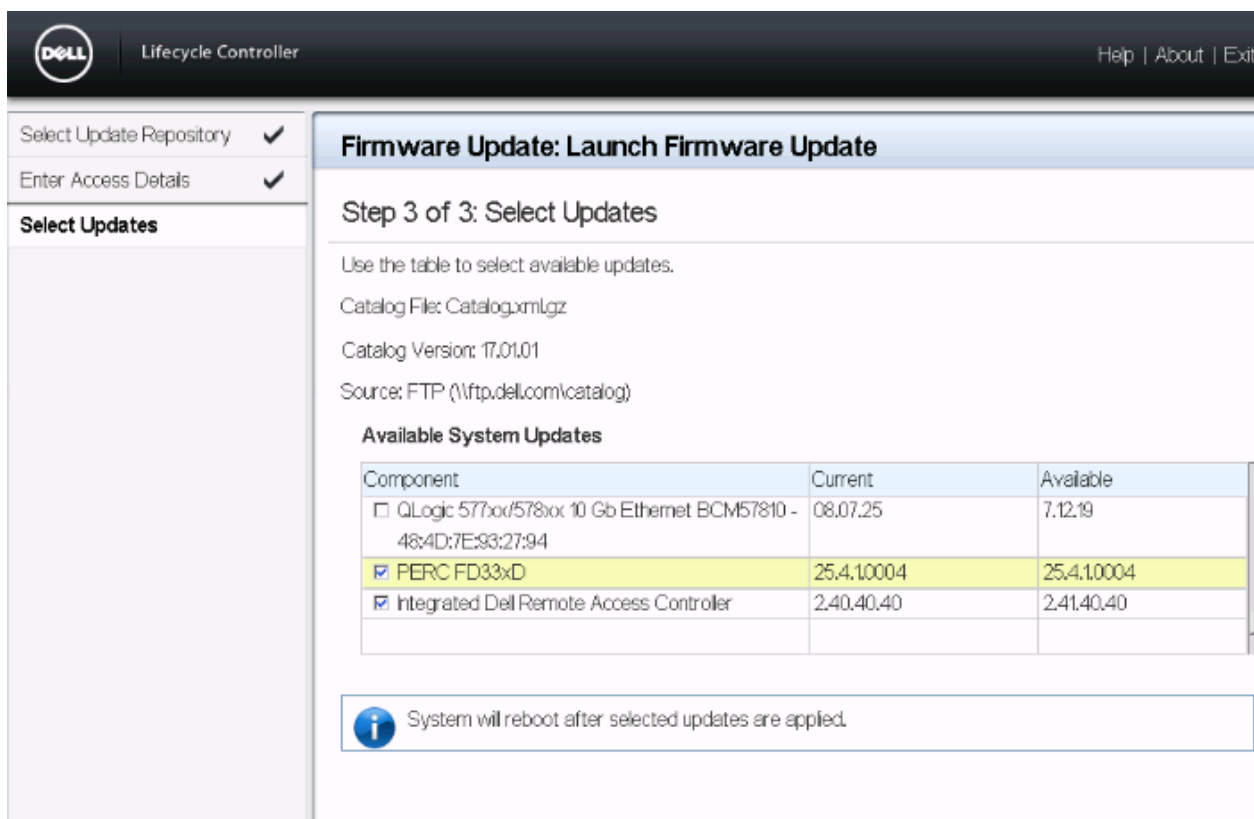


Figure 18 Launch firmware update: Select updates

Please Note, For Vsphere 6.5 you must use PERC Firmware version 25.5.0.0018, in case of automatic update does not update it to correct version please add it manually via the CMC, you can download from the link below:

https://downloads.dell.com/FOLDER03944862M/5/SAS-RAID_Firmware_2H45F_WN32_25.5.0.0018_A08.EXE

Firmware update process via CMC

1. Login to FX2 CMC
2. Navigate to server overview
3. Click update
4. Choose RAID Controller
5. Choose update for the relevant servers

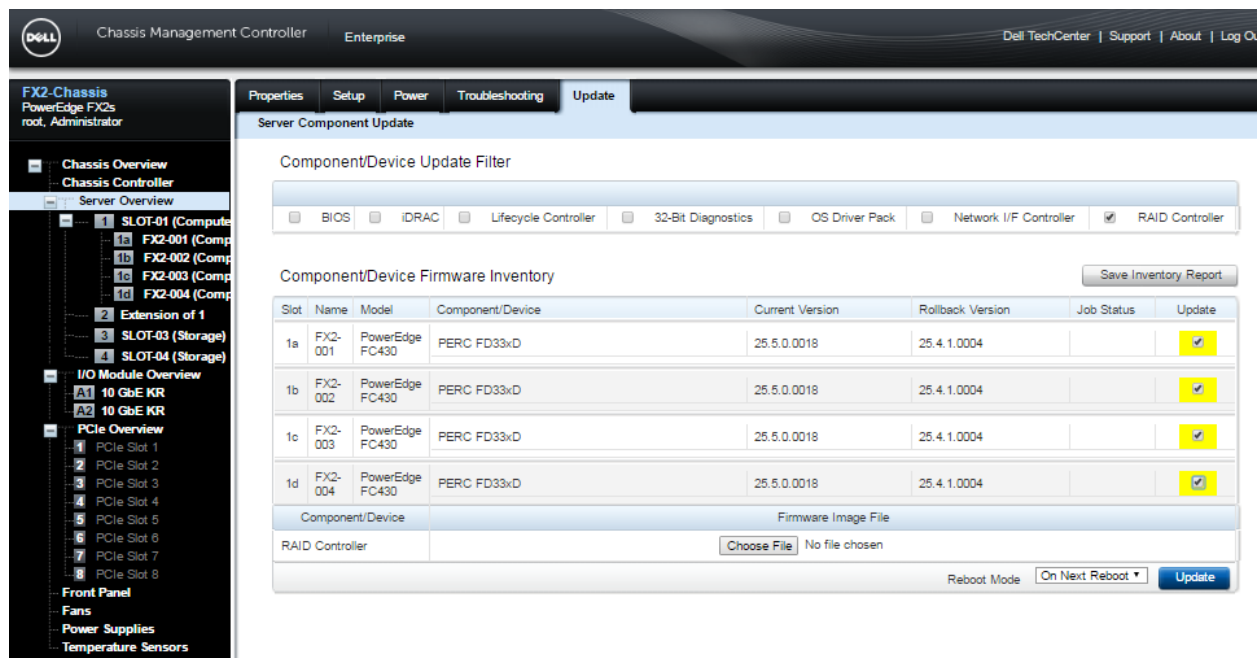


Figure 19 Firmware update process via CMC

6. Choose update file and press update button.




Figure 20 Choose update file and press update button.

7. Reboot the servers once update is finished

** Servers will automatically enter to Lifecycle controller and update the relevant firmware.

Note – Via the CMC you can use repository manager and process the updates via the “update from network Share”

Repository manager is part of [Open Manage Essentials pack](#), you can download Dell Repository manager standalone edition as well from [here](#).

6.4 FD332 Physical Disk Configuration

When working with FD332 in a 3 or more servers scenario attached to the FD332 storage sleds we will use split dual host mode for the controller configuration.

in a vSAN scenario it is very important to set the disks properly so each server has its own disks assigned – 0-7 for Server A / 8-15 for Server B.

Storage Modes:

- **Split Dual Host** — Select this mode if you want to connect the PERC controllers in split mode to two compute sleds. If you select this option, disk drives 0 to 7 are assigned to the PERC controller 1 (primary), and disk drives 8 to 15 are assigned to controller 2.
- **Split Single Host** — Select this mode if you want to connect the PERC controllers in split mode to a single server. If you select this option, disk drives 0 to 7 are assigned to PERC controller 1 (primary), and disk drives 8 to 15 are assigned to PERC controller 2.
- **Joined** — Select this option if you want to assign all disk drives (0 to 15) to the PERC controller (primary) and connect to a compute server.

6.5 Confirm storage controllers for VSAN disks are in HBA mode

For redundancy, VSANs employ software RAID. With the exception of single drive RAID-0 configurations, VSANs do not support hardware RAID.

Note: For more information see the [VMware Virtual SAN Hardware Guidance](#) white paper.

Storage controllers used in VSANs should therefore be set to HBA mode (also referred to as pass-through mode). For the deployment used in this guide, this applies to all PERC FD33xD and PERC H730 controllers in FC430 and R630 servers respectively.

To verify storage controllers are in HBA mode:

1. Connect to the iDRAC in a web browser and launch the virtual console.
2. In the virtual console, from the **Next Boot** menu, select **BIOS Setup**.

3. Reboot the server.
4. From the System Setup Main Menu, select Device Settings.
5. From the list of devices, select the PERC controller. This opens the **Modular RAID Controller Configuration Utility Main Menu**.
6. Select Controller Management. Scroll down to **Controller Mode** and verify it is set to **HBA**. If set to **RAID**, select **Advanced Controller Management > Switch to HBA Mode > OK**.

Note: If unable to switch to HBA mode, configured RAID virtual disks may need to be deleted first. See your system documentation for more information.

7. Click **Back** and **Finish** as needed to exit **System Setup**.

** in some cases the RAID controller is already configured (not factory integrated) so you'll need to go through the following steps:

Start in RAID Mode:

(Can determine raid or hba mode setting before starting in idrac: Storage > Controllers > Properties)

1. Boot to Life Cycle controller > HW Config > Config Wiz > RAID Config.
2. Select Controller > Next. **If get warning that controller is in HBA Mode:**
 - a. Click Back > Back > System Setup > Adv HW Configuration.
 - b. Device Settings > PERC controller > Controller Management.
 - c. Scroll down, Controller Mode shows HBA. Select Advanced Controller Management > Switch to RAID Mode (make sure says operation has been performed successfully) > OK.
 - d. Exit out to go back to Lifecycle controller, Set Next Boot to boot to LC, Exit LC, **Reboot** back to LifeCycle controller (required or will get same error).
 - e. In LC: HW Config > Config Wiz > RAID Config.
 - f. Select Controller > Next.
3. Select all Disks. Use the wizard to configure RAID 0 with all disks.
4. From LifeCycle controller, Go to System Setup > Adv HW Configuration.
5. Device Settings > PERC controller > VD Management.
6. Select the VD > Operation: Delete Virtual Disk > Go > confirm > yes > ok > Back.

Switch to HBA mode:

7. **Controller Mgmt** > Scroll down, Controller Mode shows RAID. Select **Advanced Controller Management > Switch to HBA Mode** (make sure says operation has been performed successfully) > **OK**.
8. Back, Finish, exit and reboot.

Disks should be clean and in HBA mode.

6.6 Install ESXi

Dell EMC recommends using the latest Dell EMC customized ESXi .iso image available on support.dell.com. The correct drivers for your PowerEdge hardware are built into this image.

Install ESXi on all servers that will be part of your deployment. For the example in this guide, ESXi is installed to redundant internal SD cards in the PowerEdge servers. This includes six R630 servers (in the management and edge clusters) and four FC430 servers (in the compute cluster).

A simple way to install ESXi on a PowerEdge server remotely is by using the iDRAC to boot the server directly to the ESXi .iso image. This is done as follows:

1. Connect to the iDRAC in a web browser and launch the virtual console.
2. In the virtual console, select **Virtual Media > Connect Virtual Media**.
3. Select **Virtual Media > Map CD/DVD** > browse to the Dell EMC customized ESXi .iso image > **Open > Map Device**.
4. Select **Next Boot > Virtual CD/DVD/ISO > OK**.
5. Select **Power > Reset System (warm boot)**. Answer **Yes** to reboot the server.
6. The server reboots to the ESXi .iso image and installation starts.
7. Follow the prompts to install ESXi. Select the server's Internal Dual SD Module (IDSDM) when prompted for a location.
8. After installation is complete, click **Virtual Media > Disconnect Virtual Media > Yes**.
9. Reboot the system when prompted.

6.7 Configure the ESXi management network connection

Be sure the host is physically connected to the management network. For this deployment, the Intel I350-T 1GbE add-in PCIe adapter provides this connection for R630 servers and FC430 servers.

1. Log in to the ESXi console and select **Configure Management Network > Network Adapters**.
2. Select the correct vmnic for the management network connection. Follow the prompts on the screen to make the selection.
3. Go to **Configure Management Network > IPv4 Configuration**. If DHCP is not used, specify a static IP address, mask, and default gateway for the management interface.
4. Optionally, configure DNS settings from the **Configure Management Network** menu if DNS is used on your network.
5. Press **Esc** to exit and answer **Y** to apply the changes.
6. From the ESXi main menu, select **Test Management Network**. Verify pings are successful. If there is an error, be sure you have configured the correct vmnic.
7. Optionally, under **Troubleshooting Options**, enable the ESXi shell and SSH to enable remote access to the CLI.
8. Log out of the ESXi console.

7 Deploy VMware vCenter Server and add hosts

7.1 Deploy VMware vCenter Server

VMware vCenter Server is required for managing clusters and vSAN, as well as many other advanced vSphere features. vCenter Server can be installed as a Windows-based application or as a prepackaged SUSE Linux-based VM.

This guide uses the prepackaged VM, called the vCenter Server Appliance (VCSA) and its built in PostgreSQL database. VCSA supports up to 1000 hosts and 10,000 VMs. VCSA is available for download at my.vmware.com.

In this guide, VCSA is installed on a PowerEdge FC430 server running ESXi.

Note: This section provides simplified VCSA installation instructions. Detailed instructions and information are provided in the VMware vCenter Server 6.0 Deployment Guide available at the following location: <https://www.vmware.com/files/pdf/techpaper/vmware-vcenter-server6-deployment-guide.pdf>

1. On a Windows workstation connected to the management network, mount the VCSA .iso image.
2. Install the Client Integration Plugin by running `\vcsa\VMWare-ClientIntegrationPlugin-6.0.0.exe`.
3. Open `\vcsa-setup.html` in a browser and accept the related warning prompts. Click **Install**.
 - a. Accept the license agreement and click **Next**.
 - b. Provide the ESXi host destination IP address, ESXi host username (root) and password. Click **Next**. Click **Yes** to accept the SSL certificate warning if prompted.
 - c. Provide a vCenter **Appliance name** (vctr01 for example), and **password**. Click **Next**.
 - d. Keep the default selection: **Install vCenter Server with an Embedded Platform Services Controller**. Click **Next**.
 - e. Select **Create a new SSO domain > Next**.
 - f. Provide an **SSO Password**, **SSO Domain name** (pct.lab for example), and **SSO Site name** (site for example).
 - g. Select an **Appliance size** depending on your requirements. For this guide **Medium (up to 400 hosts, 4000 VMs)** is selected.
 - h. Select a **datastore**. Optionally, if space is limited, check the **Enable Thin Disk Mode** box. Click **Next**.
 - i. Keep the default selection: **Use an embedded database (PostgreSQL)**. Click **Next**.
 - j. Under **Network Settings**:
 - i. Keep the default network, **VMNetwork**.
 - ii. Select **IPv4** and the network type (**static or DHCP**). A static address is used in this guide.
 - iii. If **static** was selected, provide a **Network address**, **System name** (if not using a fully qualified domain name, retype the Network address), **Subnet mask**, **Network gateway**, and **DNS server**.
 - iv. Under **Configure time sync**, select **Synchronize appliance time with ESXi host**.

Note: If you select Use NTP (Network Time Protocol) servers, a warning appears at the bottom of the screen indicating deployment will fail if the ESXi host clock is not in sync with the NTP server. Since the ESXi

hosts are not yet configured for NTP, select Synchronize appliance time with ESXi host. ESXi hosts are configured for NTP in Section 7.

- v. Checking **Enable SSH** is optional. Click **Next**. Click **OK** if a fully qualified domain name (FQDN) recommendation box is displayed.
- k. Joining the **VMWare Customer Experience Improvement Program** is recommended but optional. Select an option and click **Next**.
- l. Review the summary page and click **Finish** if all settings are correct.

vCenter Server is installed as a virtual machine on the ESXi host. When complete, the message shown in Figure 21 is displayed.

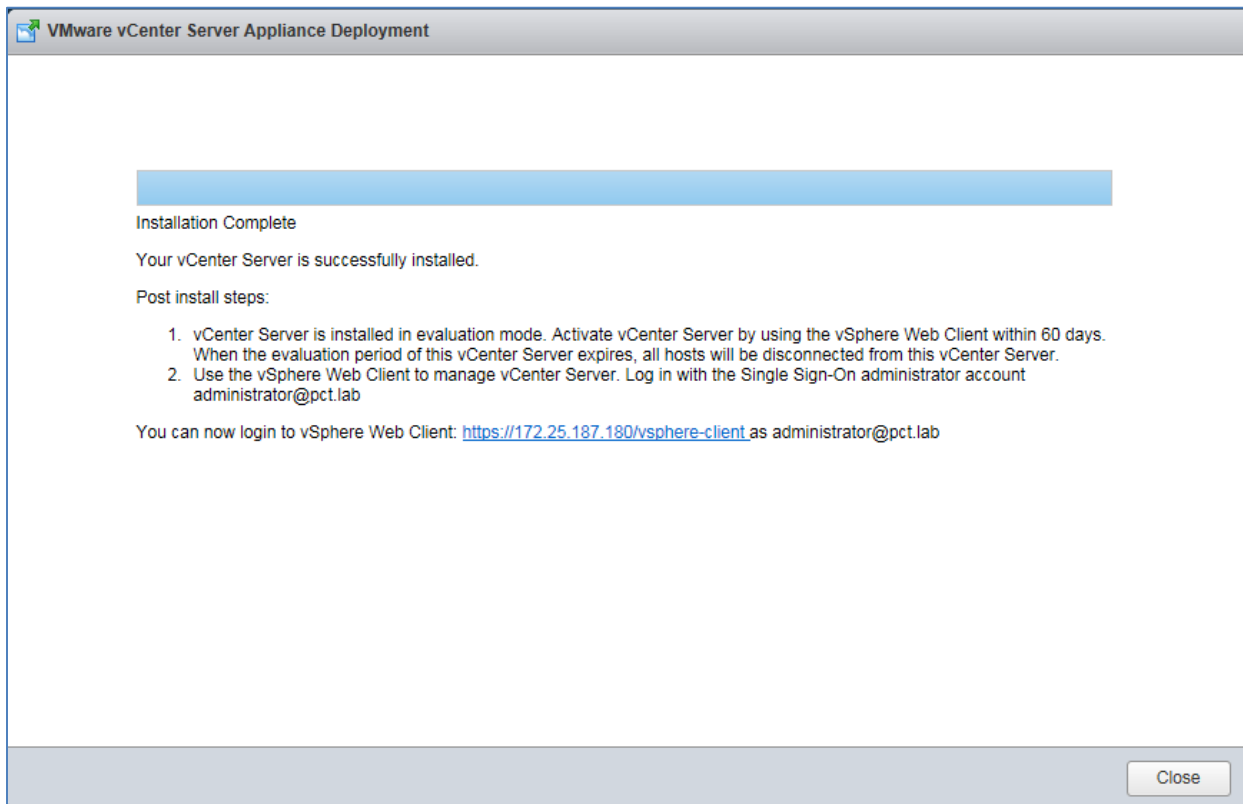


Figure 21 vCenter Server installation complete

7.2 Connect to the vSphere web client

Note: The vSphere Web Client is a service running on vCenter Server.

Connect to the vSphere Web Client in a browser by entering the following in the address bar:

`https://<ip-address-or-hostname-of-vCenter-appliance>/vsphere-client`

Log in with your vCenter credentials. After log in, the web client home page is displayed as shown in Figure 22.

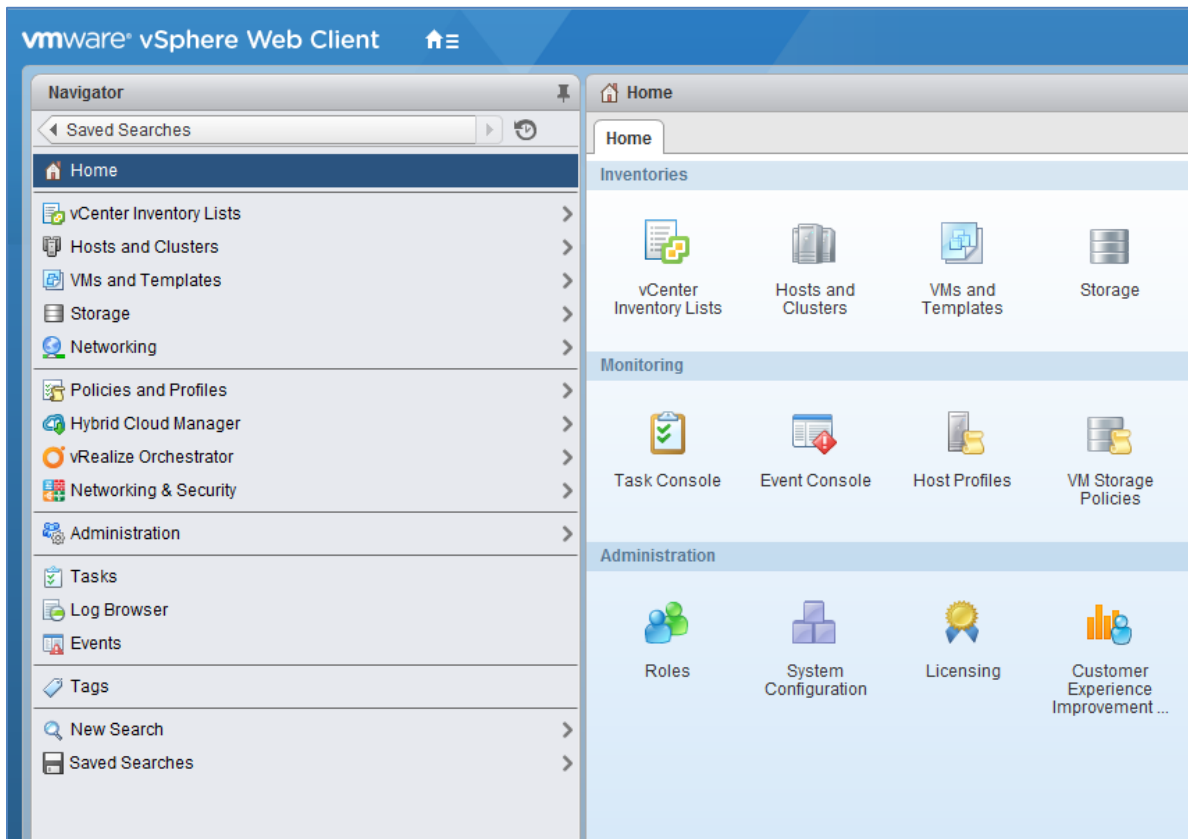


Figure 22 vSphere Web Client home page

The vast majority of management, configuration, and monitoring of your vSphere and vSAN environment is done in the web client.

7.3 Install VMware licenses

The VMware licenses required for this deployment are listed in Appendix B.2. All VMware products used in this guide come with evaluation licenses that can be used for up to 60 days.

To install one or more product licenses:

1. From the web client **Home** page, select **Licensing** in the center pane.
2. Click the **+** icon, and type or paste license keys into the box provided. Click **Next**.
3. Provide **License names** for the keys or use the defaults. Click **Next > Finish**.

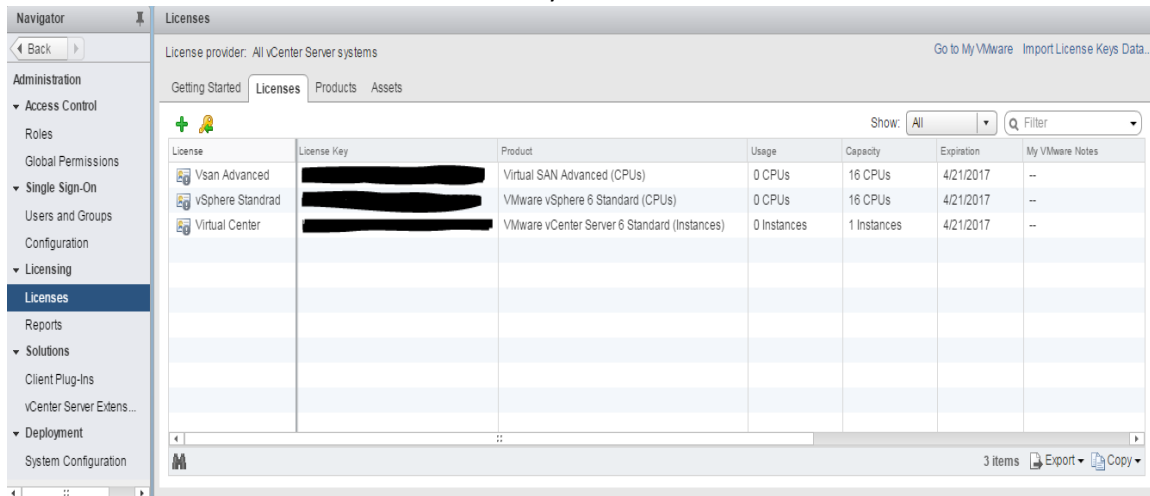


Figure 23 Install VMware licences

7.4 Create a datacenter object and add hosts

A datacenter object needs to be created before hosts can be added. This guide uses a single datacenter object named Datacenter.

To create a datacenter object:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, right click the vCenter Server object and select **New Datacenter**.
3. Provide a name (Datacenter) and click **OK**.

To add ESXi hosts to the datacenter:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, right click on **Datacenter** and select **Add Host**.
3. Specify the **IP address** of an ESXi host (or the **host name** if DNS is configured on your network). Click **Next**.
4. Enter the credentials for the ESXi host and click **Next**. If a security certificate warning box is displayed, click **Yes** to proceed.
5. On the **Host summary** screen, click **Next**.
6. Assign a license or select the evaluation license. This guide uses a VMware vSphere 6 Enterprise Plus license for ESXi hosts. Click **Next**.
7. Select a **Lockdown mode**. This guide uses the default setting, **Disabled**. Click **Next**.
8. For the VM location, select Datacenter and click Next.
9. On the **Ready to complete** screen, select **Finish**.

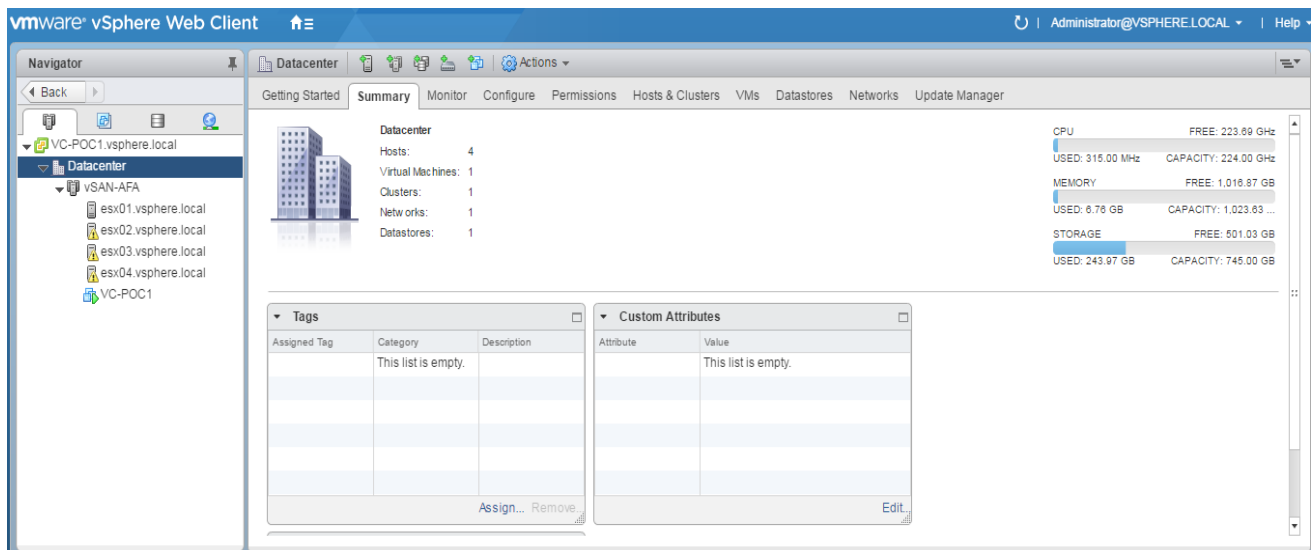



Figure 24 ESXi hosts added to the datacenter

Repeat for all servers running ESXi that will be part of the environment. This deployment example uses four FC430 servers running ESXi. When complete, all ESXi hosts will be added to the datacenter as shown in Figure 24.

Some hosts may have a warning icon () as shown in Figure 24. By selecting the host and going to the **Summary** tab, the warning message can be viewed. In this case the warning indicates the ESXi Shell and SSH have been enabled (as described in Section 6.7). If the behavior is desired, click **Suppress Warning**.

The following warning messages may also appear:

- *No datastores have been configured.* This message will be resolved when either a local datastore or VSAN datastore is configured. VSAN datastore configuration is covered in section 9 or see your ESXi documentation to create a local datastore.
- *System logs on host are stored on non-persistent storage.* This message may appear when ESXi is installed to the redundant internal SD cards. This can be resolved by moving the system logs to either a local datastore or VSAN datastore. VSAN datastore configuration is covered in section 9 or see your ESXi documentation to create a local datastore. Resolution is documented in VMware Knowledge Base article [2032823](#).

7.5 Ensure hosts are configured for NTP

It is a best practice to use NTP on the management network to keep time synchronized in an vSphere environment. Ensure NTP is configured on ESXi hosts as follows:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, select a host.
3. In the center pane, go to **Configure > System > Time Configuration**.
4. If NTP has not been configured properly, click **Edit**.
5. In the Edit Time Configuration dialog box:
 - a. Select **Use Network Time Protocol** radio button.
 - b. Next to **NTP Service Startup Policy**, select **Start and stop with host**.
 - c. Next to **NTP servers**, enter the IP address or FQDN of the NTP server.
 - d. Click **Start** to start the NTP client followed by **OK** to close the dialog box.
6. The **Time Configuration** page for the host should appear similar to Figure 25.

esx01.vsphere.local: Edit Time Configuration

Specify how the date and time on this host should be set.

☐ Manually configure the date and time on this host

02/21/2017 5:07 PM

☒ Use Network Time Protocol (Enable NTP client)

NTP Service Status:	Running
	<input type="button" value="Start"/> <input type="button" value="Stop"/> <input type="button" value="Restart"/>
	The NTP Service settings are updated when you click Start, Restart, or Stop.
NTP Service Startup Policy:	Start and stop with host
	Start and stop with the host system
NTP Servers:	10.0.0.101
	Separate servers with commas, e.g. 10.31.21.2, fe00::2800

OK Cancel

Figure 25 Proper NTP configuration on ESXi host

7. Repeat for remaining ESXi hosts as needed.

7.6 Create clusters and add hosts

When a host is added to a cluster, the host's resources become part of the cluster's resources. The cluster manages the resources of all hosts within it. Clusters enable features such as High Availability (HA), Distributed Resource Scheduler (DRS), and Virtual SAN (VSAN). For this guide, single cluster is created:

- vSAN-AFA

All ESXi hosts are added to the cluster.

To add clusters to the datacenter:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, right click the datacenter object and select **New Cluster**.
3. Name the cluster. Leave **DRS**, **vSphere HA**, **EVC** and **Virtual SAN** at their default settings (**Off/Disabled**). Click **OK**.

In the Navigator pane, drag and drop ESXi hosts into the cluster. The four ESXi hosts on FC430 servers are placed in the **vSAN-AFA** cluster.

When complete, each cluster (🏠) should contain its assigned hosts (🏠) as shown in Figure 26:

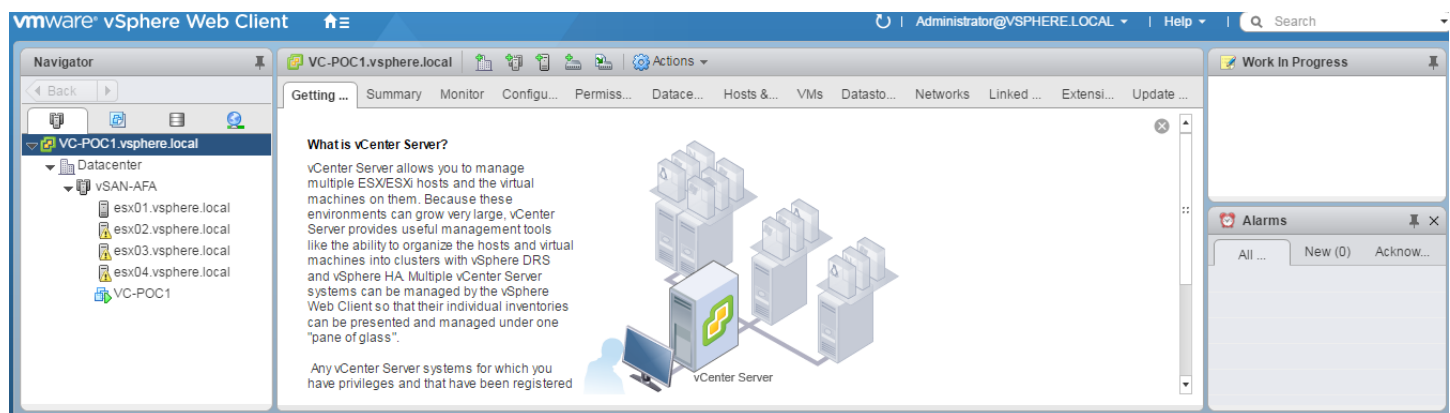


Figure 26 Clusters and hosts after initial configuration

7.7 Information on vSphere standard switches

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

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

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

1. Go to the web client **Home** page, select **Hosts and Clusters**, and select a host in the **Navigator** pane.
2. In the center pane, select **Configure > Networking > Virtual switches**.
3. Standard switch **vSwitch0** appears in the list. Click on it to view details as shown in Figure 27.

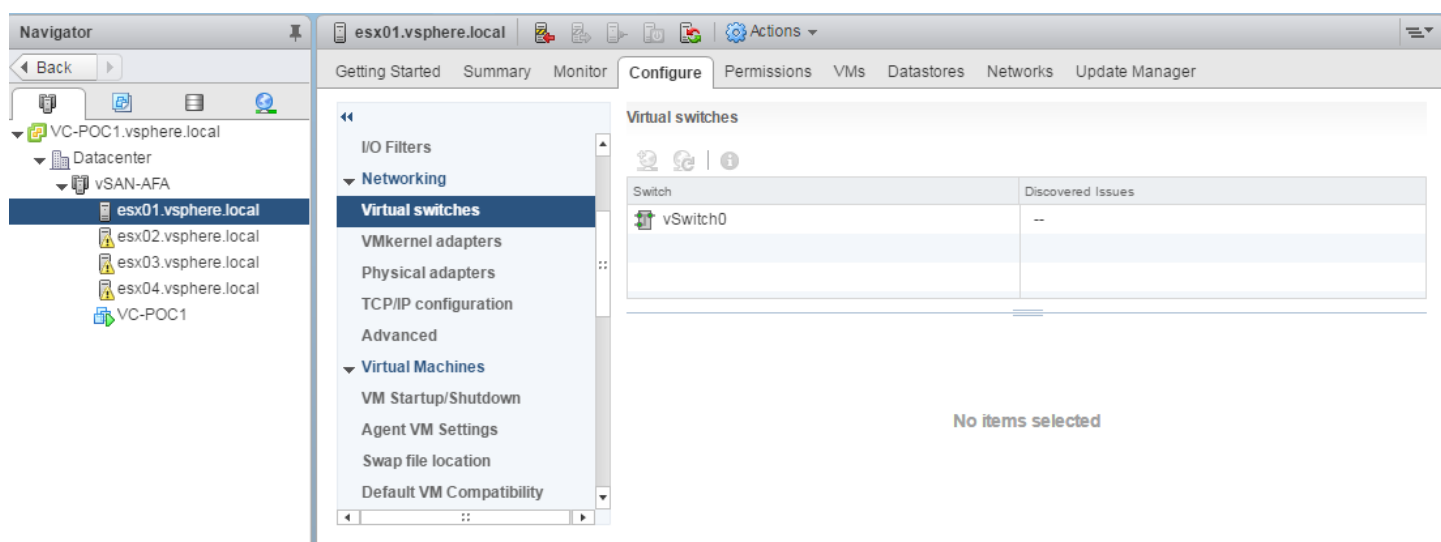


Figure 27 vSphere standard switch

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

8 Deploy 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. This lets virtual machines maintain consistent network configurations as they migrate across multiple hosts.

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

Distributed Switches contain two different port groups:

- **Uplink port group** – an uplink port group maps physical NICs on the hosts (vmnics) to uplinks on the VDS. Uplink port groups act as trunks and carry all VLANs by default.

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

- **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 each VLAN.

For this guide, one VDS is created for all cluster, and shared by all hosts in the cluster. The distributed switch used in this deployment is named:

- vSAN VDS

8.1 Create a VDS for the vSAN cluster.

Create the first VDS named **vSAN VDS**:

1. On the web client **Home** screen, select **Networking**.
2. Right click on Datacenter. Select Distributed switch > New Distributed Switch.
3. Provide a name for the first VDS, **vSAN VDS**. Click **Next**.
4. On the Select version page, select **Distributed switch: 6.5.0 > Next**.
5. On the **Edit settings** page:
 - a. Leave the **Number of uplinks** set to **4** (this field to be replaced by LAGs later).
 - b. Leave **Network I/O Control** set to **Enabled**.
 - c. **Uncheck** the **Create a default port group** box.
6. Click **Next** followed by **Finish**.
7. The VDS is created with the uplink port group shown beneath it.

When complete, the Navigator pane should look similar to Figure 28.

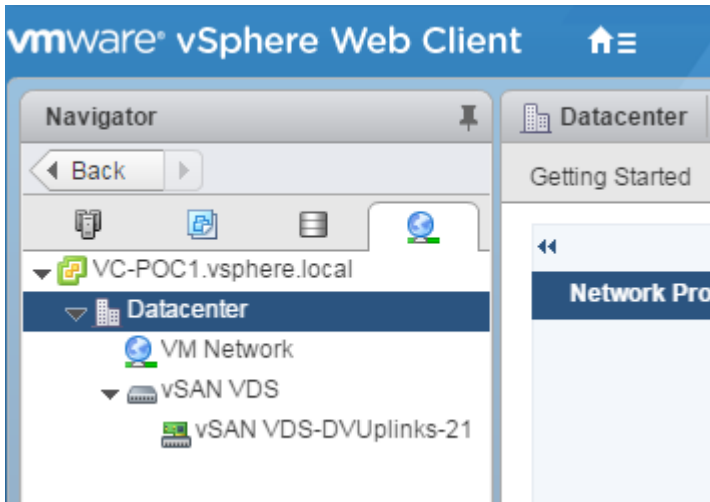


Figure 28 VDS created for each cluster

8.2 Add distributed port groups

In this section, separate distributed port groups for vMotion and VSAN traffic are added to VDS.

To create the port group for vMotion traffic on the **vSAN VDS**:

1. On the web client **Home** screen, select **Networking**.
2. Right click on vSAN VDS. Select Distributed Port Group > New Distributed Port Group.
3. On the **Select name and location** page, provide a name for the distributed port group, for example, **vMotion**. Click **Next**.
4. On the **Configure settings** page, next to **VLAN type**, select **VLAN**. Set the **VLAN ID** to **22** for the vMotion port group. Leave other values at their defaults as shown in Figure 29.
5. Click Next > Finish.




Figure 29 Distributed port group settings page – vMotion port group

Repeat steps 1-5 above to create the distributed port group for VSAN traffic, except replace "vMotion" with "VSAN" in the **port group name** and set the **VLAN ID** to **44** for the VSAN port group, and **VLAN ID** to **10** for the VMs port group.

When complete, the Navigator pane will appear similar to Figure 30.

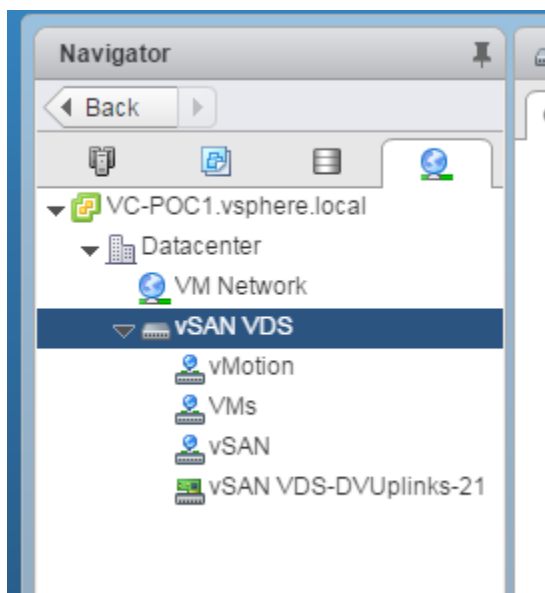


Figure 30 Distributed switches with vMotion, VMs and VSAN port groups created

8.3 Create LACP LAGs

Since Link Aggregation Control Protocol (LACP) LAGs are used in the physical network between ESXi hosts and physical switches, LACP LAGs are also configured on each VDS.

To enable LACP on **vSAN VDS**:

1. On the web client **Home** screen, select **Networking**.
2. In the **Navigator** pane, select vSAN VDS.
3. In the center pane, select **Configure > Settings > LACP**.
4. Click the **+** icon. The **New Link Aggregation Group** dialog box opens.
5. Set the number of ports equal to the number of physical uplinks on each ESXi host. In this deployment, FC430 hosts use two ports in a LAG to connect to the upstream switches so this number is set to **2**.
6. Set the **Mode** to **Active**. The remaining fields can be set to their default values as shown in Figure 31.

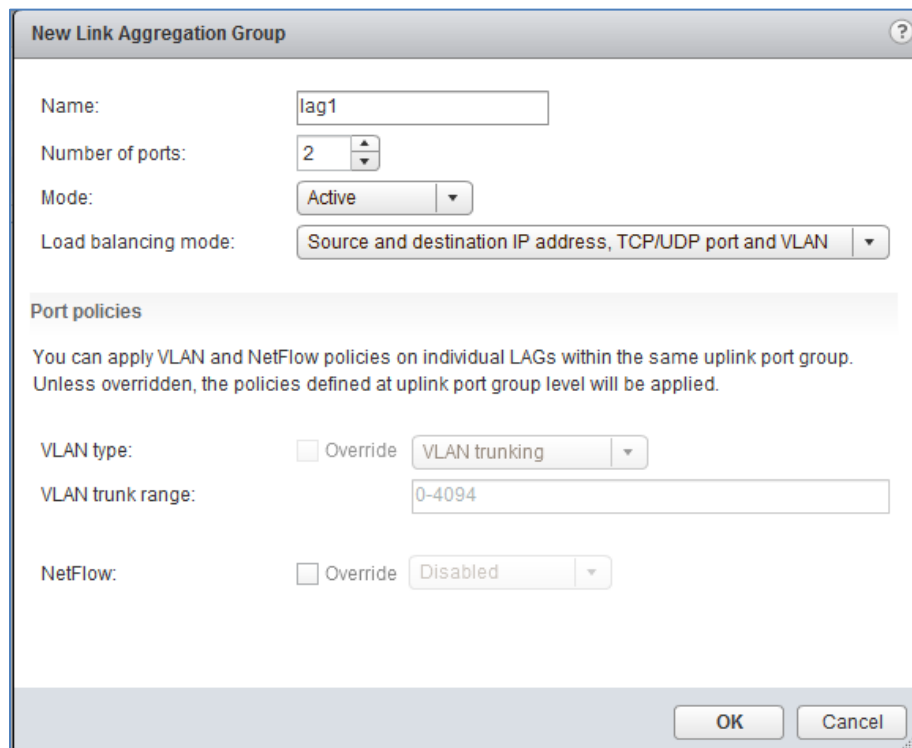
The image shows a 'New Link Aggregation Group' dialog box. It has a title bar with a question mark icon. The main area contains several fields: 'Name' with a text box containing 'lag1'; 'Number of ports' with a spinner box set to '2'; 'Mode' with a dropdown menu set to 'Active'; and 'Load balancing mode' with a dropdown menu set to 'Source and destination IP address, TCP/UDP port and VLAN'. Below these is a section titled 'Port policies' with explanatory text: 'You can apply VLAN and NetFlow policies on individual LAGs within the same uplink port group. Unless overridden, the policies defined at uplink port group level will be applied.' This section contains three rows: 'VLAN type' with an 'Override' checkbox and a dropdown set to 'VLAN trunking'; 'VLAN trunk range' with a text box containing '0-4094'; and 'NetFlow' with an 'Override' checkbox and a dropdown set to 'Disabled'. At the bottom right are 'OK' and 'Cancel' buttons.

Figure 31 LAG configuration

7. Click **OK** to close the dialog box.

This creates **lag1** on the VDS. The refresh icon (🔄) at the top of the screen may need to be clicked for the lag to appear in the table as shown in Figure 32.

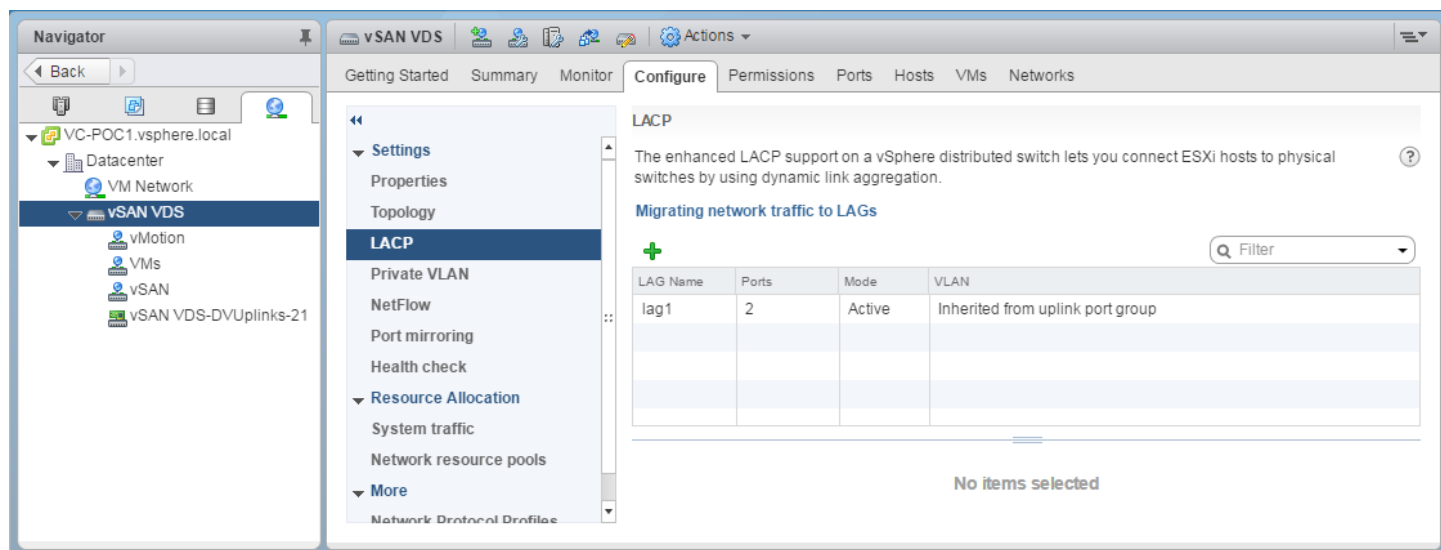


Figure 32 Lag1 created on Rack 1 Management VDS



8.4 Associate hosts and assign uplinks to LAGs

Hosts and their vmnics must be associated with each vSphere distributed switch.

Note: Before starting this section, be sure you know the vmnic-to-physical adapter mapping for each host. This can be determined by going to **Home > Hosts and Clusters** and selecting the host in the **Navigator** pane. In the center pane select **Configure > Networking > Physical adapters**. In this example, vmnics used are numbered vmnic0 and vmnic1. Vmnic numbering will vary depending on adapters installed in the host.

To add hosts to vSAN VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on vSAN 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 check box next to each host in the **vSAN-AFA** cluster. Click **OK > Next**.
 - c. On the **Select network adapters tasks** page, be sure the **Manage physical adapters** box is checked. Be sure all other boxes are unchecked. Click **Next**.
 - d. On the **Manage physical network adapters** page, each host is listed with its vmnics beneath it.
 - i. Select the first vmnic (vmnic1 in this example) on the first host and click **Assign uplink**.
 - ii. Select **lag1-0 > OK**.

- iii. Select the second vmnic (vmnic3 in this example) on the first host and click  **Assign uplink**.
- iv. Select **lag1-1** > **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 vSAN VDS should look similar to Figure 33.

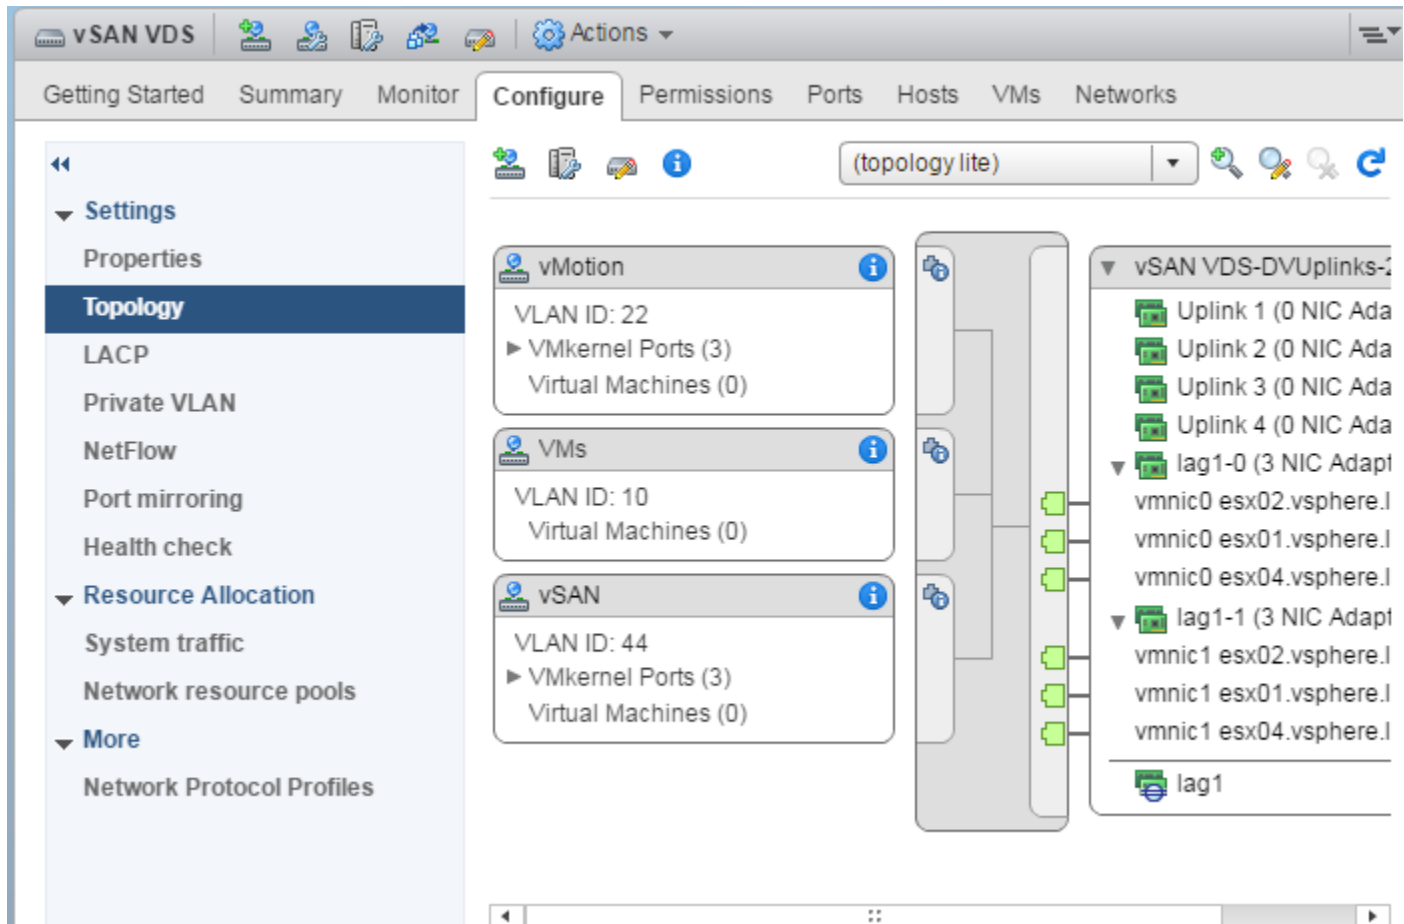


Figure 33 LAGs configured on vSAN VDS

This configuration brings up the LAGs on the upstream switches. This can be confirmed by running the `show vlt detail` command on the upstream switches as shown in the examples from Leaf-1 (Management Cluster) and FN410S-A1 (Compute Cluster) below. The Local and Peer Status columns now indicate all LAGs are UP.

Spine-1#**show vlt detail**

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
-----	-----	-----	-----	-----
2	2	UP	UP	1, 22, 44
4	4	UP	UP	1, 22, 44
6	6	UP	UP	1, 22, 44

FN410S-A1#**show vlt detail**

Local LAG Id	Peer LAG Id	Local Status	Peer Status	Active VLANs
-----	-----	-----	-----	-----
1	1	UP	UP	1, 22, 44, 55
2	2	UP	UP	1, 22, 44, 55
3	3	UP	UP	1, 22, 44, 55
4	4	UP	UP	1, 22, 44, 55
33	33	UP	UP	1, 22, 44, 55

8.5 Configure teaming and failover on LAGs

1. On the web client **Home** screen, select **Networking**.
2. Right click on vSAN VDS. Select Distributed Port Group > Manage Distributed Port Groups.
3. Select only the **Teaming and failover** checkbox. Click **Next**.
4. Click **Select distributed port groups**. Check the top box to select all port groups (vMotion and VSAN). Click **OK > Next**.
5. On the **Teaming and failover** page, click **lag1** and move it up to the **Active uplinks** section by clicking the up arrow. Move **Uplinks 1-4** down to the **Unused uplinks** section. Leave other settings at their defaults. The **Teaming and failover** page should look similar to Figure 34 when complete.

Rack 1 Management VDS - Manage Distributed Port Groups

Teaming and failover
Controls load balancing, network failure detection, switches notification, failback, and uplink failover order.

Load balancing: ⓘ

Network failure detection:

Notify switches:

Failback:

Failover order

↑ ↓

Active uplinks

lag1

Standby uplinks

Unused uplinks

Uplink 1

Uplink 2

Uplink 3

Uplink 4

Select active and standby uplinks. During a failover, standby uplinks activate in the order specified above.

Back Next Finish Cancel

Figure 34 Teaming and failover settings

6. Click **Next** followed by **Finish** to apply the settings.

8.6 Add VMkernel adapters for vMotion and VSAN

In this section, vMotion and VSAN VMkernel adapters (also referred to as VMkernel ports) will be added to each ESXi host to allow for vMotion and VSAN traffic.


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

This deployment uses the following addressing scheme for the vMotion, VSAN and VMs networks:


Table 3 VLAN and network examples

VLAN ID	Network	Gateway	Used For
22	10.22.1.0/24	10.22.1.1	vMotion
44	10.44.1.0/24	10.44.1.1	VSAN
10	10.10.1.0/24	10.10.1.1	VMs


To add VMkernel adapters to all hosts connected to the vSAN VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on vSAN VDS, and select **Add and Manage Hosts**.
3. In the Add and Manage Hosts dialog box:
 - a. On the **Select task** page, make sure **Manage host networking** is selected. Click **Next**.
 - a. On the **Select hosts** page, click  **Attached hosts**. Select all hosts. Click **OK > Next**.
 - b. On the **Select network adapter tasks** page, make sure the **Manage VMkernel adapters** box is checked and all other boxes are unchecked. Click **Next**.
 - c. The **Manage VMkernel network adapters** page opens.



vMotion adapter

- i. To add the vMotion adapter, select the first host and click  **New Adapter**.
- ii. On the **Select target device** page, click the radio button next to **Select an existing network** and click **Browse**.
- iii. Select the port group created for vMotion > **OK**. Click **Next**.
- iv. On the **Port properties** page, leave **IPv4** selected and check only the **vMotion traffic** box. Click **Next**.
- v. On the **IPv4 settings** page, if DHCP is not used, select **Use static IPv4 settings**. Set the IP address, for example 10.22.1.x, and subnet mask for the host on the vMotion network. Click **Next > Finish**.

VSAN adapter

- vi. To add the VSAN adapter, select the first host and click  **New Adapter**.
- vii. On the **Select target device** page, click the radio button next to **Select an existing network** and click **Browse**.
- viii. Select the port group created for VSAN > **OK**. Click **Next**.
- ix. On the **Port properties** page, leave IPv4 selected and check only the **Virtual SAN traffic** box. Click **Next**.
- x. Change MTU size to 9000
- xi. On the **IPv4 settings page**, if DHCP is not used, select **Use static IPv4 settings**. Set the IP address, for example 10.44.1.x, and subnet mask for the host on the VSAN network. Click **Next > Finish**.

VMs adapter

- xii. To add the VMs adapter, select the first host and click  **New Adapter**.
 - xiii. On the **Select target device** page, click the radio button next to **Select an existing network** and click **Browse**.
 - xiv. Select the port group created for VMs > **OK**. Click **Next**.
 - xv. On the **Port properties** page. Click **Next**.
 - xvi. On the **IPv4 settings page**, if DHCP is not used, select **Use static IPv4 settings**. Set the IP address, for example 10.10.1.x, and subnet mask for the host on the VMs network. Click **Next > Finish**.
- d. Repeat steps i-x for the remaining hosts, then click **Next**.
 - e. On the **Analyze impact** page, **Overall impact status** should indicate  **No impact**.
 - f. Click **Next > Finish**.

When complete, the VMkernel adapters page for each ESXi host in the vSphere datacenter should look similar to Figure 35. This page is visible by going to **Hosts and Clusters**, selecting a host in the **Navigator** pane, then selecting **Configure > Networking > VMkernel adapters** in the center pane.

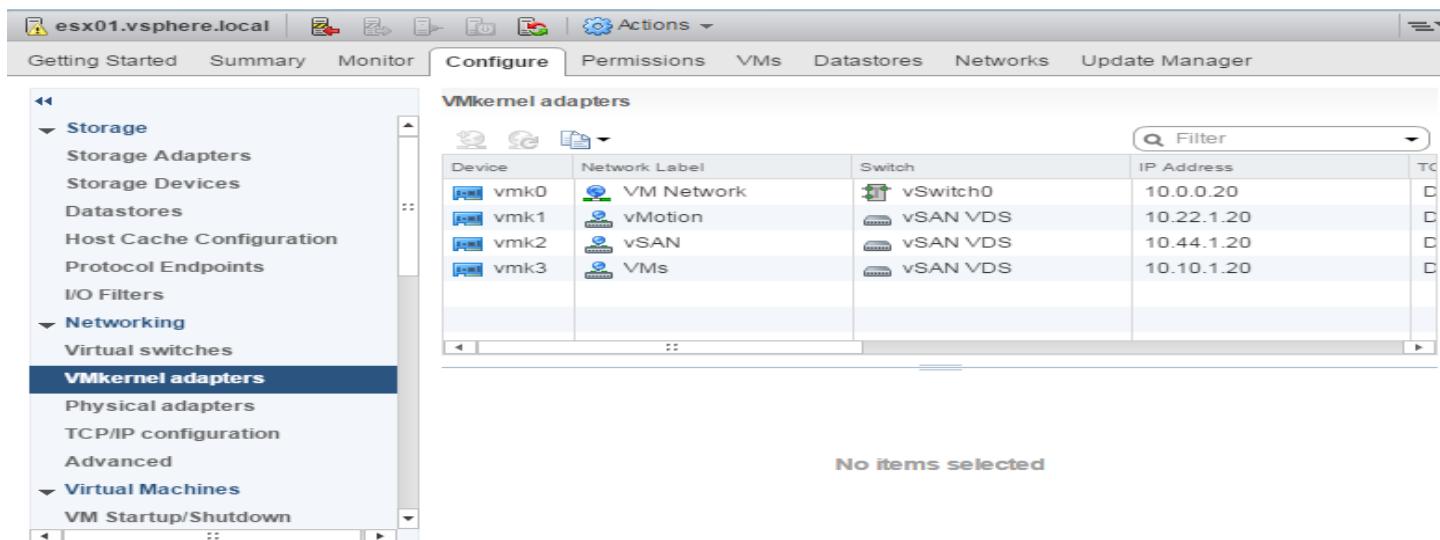


Figure 35 Host VMkernel adapters page

Adapter vmk5 was installed by default for host management. Adapters vmk1, vmk2 and vmk3 were created in this section.

To verify the configuration, ensure the vMotion adapter, **vmk1** in this example, is shown as **Enabled** in the **vMotion Traffic** column, and the VSAN adapter, **vmk2** in this example, is shown as **Enabled** in the **Virtual SAN Traffic** column. Verify the VMkernel adapter IP addresses are correct.

Verify the information is correct on other hosts as needed.

8.7 Verify VDS configuration

To verify the distributed switches have been configured correctly, the **Topology** page for each VDS provides a summary.

To view the **Topology** page for the **vSAN VDS**:

4. On the web client **Home** screen, select **Networking**.
5. In the Navigator pane, select **vSAN VDS**.
6. In the center pane, select **Configure > Settings > Topology** and click the ► icon next to **VMkernel Ports** (2 places) to expand. The screen should look similar to Figure 36.

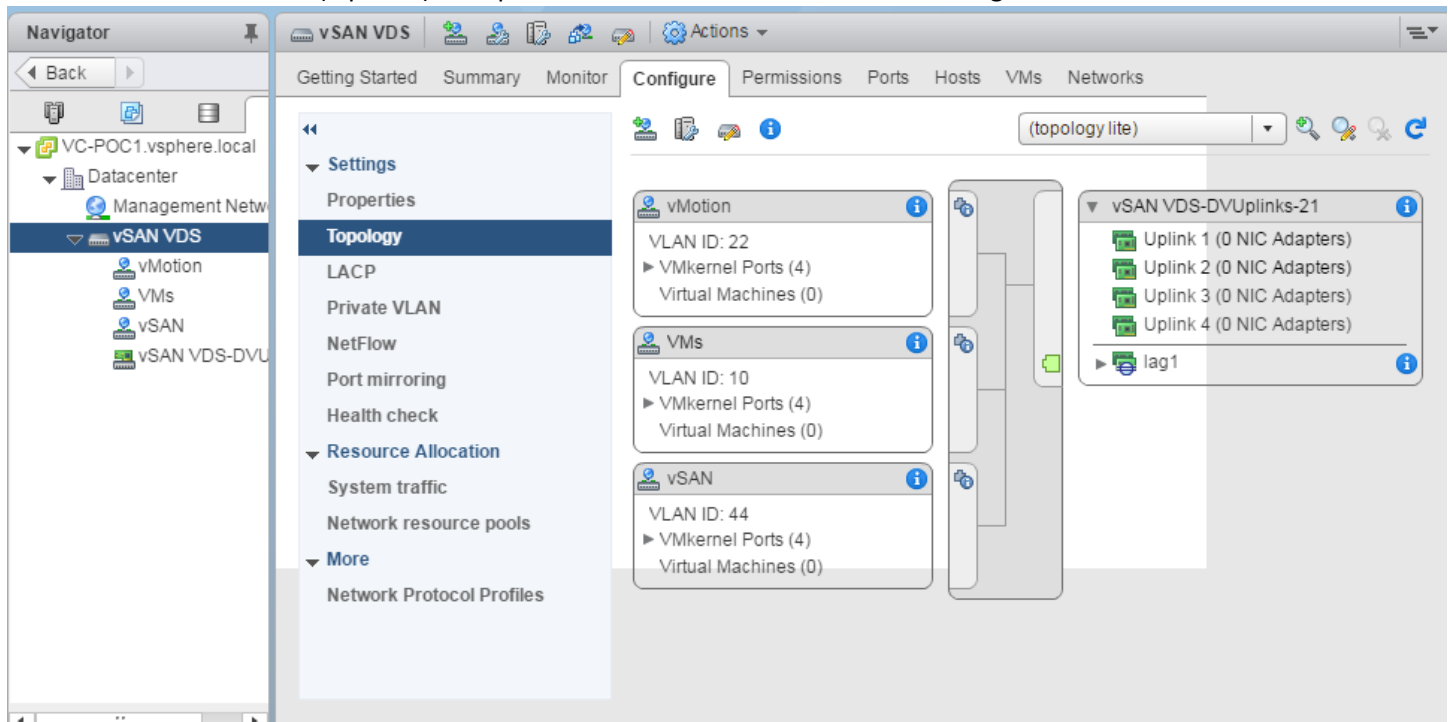


Figure 36 vSAN VDS VMkernel ports, VLANs, and IP addresses

Notice the three distributed port groups, **vMotion**, **vSAN** and **VMs** are shown in Figure 36 with their configured VLAN IDs and VMkernel ports. Since VMkernel ports were configured for all three ESXi hosts in the cluster, there are three VMkernel ports in each distributed port group.

Remember to change MTU configuration on vSAN VDS switch and on vSAN kernel adapters to 9000.

To change the **MTU** for the **vSAN VDS**:

7. On the web client **Home** screen, select **Networking**.
8. In the Navigator pane, select **vSAN VDS**.
9. In the center pane, select **Configure > Settings > Properties** and click **Edit**.
10. Click Advanced and change the MTU value to 9000

8.8 Enable LLDP

Enabling Link Layer Discovery Protocol (LLDP) on vSphere distributed switches is optional but can be helpful for link identification and troubleshooting.

Note: LLDP works as described in this section with QLogic 57810 or QLogic 57840 adapters specified in Appendix A. LLDP functionality may vary with other adapters. LLDP must also be configured on the physical switches per the switch configuration instructions provided earlier in this guide.

8.8.1 Enable LLDP on each VDS and view information sent

Enabling LLDP on vSphere distributed switches enables them to send information such as vmnic numbers and MAC addresses to the physical switch connected to the ESXi host.

To enable LLDP on each VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on a VDS, and select **Settings > Edit Settings**.
3. In the left pane of the **Edit Settings** page, click **Advanced**.
4. Under Discovery protocol, set Type to Link Layer Discovery Protocol and Operation to Both.
5. Click **OK**.

Repeat for remaining distributed switches.

To view LLDP information sent from the ESXi host adapters, run the following command from the CLI of a directly connected switch:

```
Leaf-1#show lldp neighbors
```

Loc PortID	Rem Host Name	Rem Port Id	Rem Chassis Id
Te 1/2	-	00:0a:f7:38:88:12	00:0a:f7:38:88:12
Te 1/2	localhost	00:50:56:18:88:12	vmnic1
Te 1/4	-	00:0a:f7:38:96:62	00:0a:f7:38:96:62
Te 1/4	localhost	00:50:56:18:96:62	vmnic1
Te 1/6	-	00:0a:f7:38:94:32	00:0a:f7:38:94:32
Te 1/6	localhost	00:50:56:18:94:32	vmnic1
Fo 1/49	Spine-1	fortyGigE 1/1/1	4c:76:25:e7:41:40
Fo 1/50	Spine-2	fortyGigE 1/1/1	4c:76:25:e7:3b:40

8.8.2 View LLDP information received from physical switch

LLDP configuration is part of the physical switch configurations covered in Section 5. The switches are configured to send information (host name, port number, etc.) via LLDP to the ESXi host network adapters.

To view LLDP information sent from the physical switch:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, select a host.
3. In the center pane, select **Configure > Networking > Physical adapters**.
4. Select a connected physical adapter, **vmnic1** for example.
5. Below the adapter list, select the **LLDP** tab. Information similar to that shown in Figure 37 is provided by the switch.

The screenshot shows the VMware vSphere Web Client interface. The left sidebar (Navigator) is expanded to 'Physical adapters' under the 'Networking' section. The main pane displays the 'Physical adapters' configuration for a host named 'esx01.vsphere.local'. A table lists two adapters: 'vmnic0' and 'vmnic1'. Below the table, the 'Physical network adapter: vmnic0' is selected, and the 'LLDP' tab is active. The LLDP section shows the 'Link Layer Discovery Protocol' details for the selected adapter.

Device	Actual Speed	Configured Speed	Switch	MAC Address
Broadcom Corporation QLogic 57810 10 Gigabit Ethernet Adapter				
vmnic0	10000 Mb	Auto negotiate	vSAN VDS	48:4d:7e:93:27:91
vmnic1	10000 Mb	Auto negotiate	vSAN VDS	48:4d:7e:93:27:94

Link Layer Discovery Protocol	
Chassis ID	f4:8e:38:44:9e:ea
Port ID	TenGigabitEthernet 0/1
Time to live	94
TimeOut	30

Figure 37 Information sent from physical switch to vmnic via LLDP

9 Configure a VSAN datastore in each cluster

9.1 VSAN Overview

VMware Virtual SAN virtualizes the local physical storage resources of ESXi hosts in a single cluster and turns them into pools of storage that can be divided and assigned to virtual machines and applications. VSAN is implemented directly in the ESXi hypervisor.

VSAN eliminates the need for external shared storage and simplifies storage configuration and virtual machine provisioning activities. VMware features such as HA, vMotion and DRS require shared storage.

Hosts must meet the following criteria to participate in a VSAN:

- A Virtual SAN cluster must contain a minimum of 3 and a maximum of 64 hosts that contribute capacity to the cluster.
- A host that resides in a Virtual SAN cluster must not participate in other clusters.
- If a host contributes its local capacity devices to the Virtual SAN datastore, it must provide at least one device for flash cache and at least one device for capacity, also called a data disk.
- All storage devices, drivers, and firmware versions in the Virtual SAN configuration must be certified and listed in the Virtual SAN section of the [VMware Compatibility Guide](#)

9.2 Configure VSAN

Before proceeding, ensure each host in the cluster has a properly configured VMkernel adapter enabled for VSAN traffic (covered in Section 8.6).

To configure VSAN on a cluster:

1. Go to Home > Hosts and Clusters.
2. In the **Navigator** pane, select a cluster.
3. In the center pane, select **Configure > Settings**. Under **Virtual SAN**, select **General**.
4. Click the **Configure** button to launch the **Configure Virtual VSAN** wizard.
 - a. Set **Add disks to storage** to **Manual**. Leave the remaining options at their defaults as shown in Figure 38 and click **Next**.

vSAN-AFA - Configure Virtual SAN

1 Virtual SAN capabilities

2 Network validation

3 Claim disks

4 Ready to complete

Virtual SAN capabilities
Select how you want your Virtual SAN cluster to behave.

Disk Claiming

Add disks to storage: Manual

Requires manual claiming of any new disks on the included hosts to the shared storage.

Deduplication and Compression

☒ Enable ⓘ

☐ Allow Reduced Redundancy ⓘ

Fault Domains and Stretched Cluster

☒ Do not configure ⓘ

☐ Configure two host Virtual SAN cluster ⓘ

☐ Configure stretched cluster ⓘ

☐ Configure fault domains ⓘ

Back Next Finish Cancel

Figure 38 Configure Virtual VSAN – Select VSAN capabilities page.

Note: Check Deduplication only in all flash configuration

- b. On the **Network validation** page, the VMkernel ports configured for VSAN traffic are shown with their IP addresses and a green check in the VSAN enabled column. Click **Next**.

vSAN-AFA - Configure Virtual SAN

1 Virtual SAN capabilities
2 Network validation
 3 Claim disks
 4 Ready to complete

Network validation
 Check the Virtual SAN network settings on all hosts in the cluster.

View: Virtual SAN VMkernel adapters Filter

Name	Network	IP Address	VSAN Enabled
esx02.vsphere.local			✓ Yes
vmk2	vSAN	10.44.1.21	Yes
esx01.vsphere.local			✓ Yes
vmk2	vSAN	10.44.1.20	Yes
esx04.vsphere.local			✓ Yes
vmk2	vSAN	10.44.1.23	Yes
esx03.vsphere.local			✓ Yes
vmk2	vSAN	10.44.1.22	Yes

8 items Export Copy

✓ All the hosts in this cluster have a VMkernel adapter with VSAN traffic enabled. Review the list below for more details.

Back Next Finish Cancel

Figure 39 Configure Virtual SAN - VMkernel adapter confirmation

- c. On the **Claim disks** page, set **Group by** to **Host** and expand the hosts to view available disks. One disk group will be configured for each host. A disk group should have 1 disk claimed for the **Cache Tier** and the remaining disks claimed for the **Capacity Tier**.

Note: A maximum of eight disks per disk group are allowed. Up to five disk groups can be configured per host.

- d. In Figure 40, four disk groups are created, one for each host. For the first disk on the first host, Cache Tier is selected. Capacity Tier is selected for the remaining seven disks. When all groups have been configured, make sure there is a green checkmark in the Configuration validation box as shown in Figure 40.
- e. Click Next > Finish to apply the configuration.




Figure 40 Configure Virtual SAN disk group management

Once this has been completed, a VSAN datastore is automatically created and attached to all participating hosts in the cluster.

Note: For more information see the [Designing and Sizing a Virtual SAN Cluster](#) section of the vSphere 6.5 online documentation.

9.3 Verify VSAN configuration

VSANs are viewed by going to the web client **Home** page and selecting **Storage**. In the **Navigator** pane, in addition to local storage on each host, a vsanDatastore will be listed for each VSAN created.

The default names are vsanDatastore, vsanDatastore (1), etc. The hosts associated with each VSAN can be viewed by selecting a **vsanDatastore** in the Navigator pane. In the center pane, select **Related Objects > Hosts**.

It is a good idea to give datastores more user-friendly names to make them easier to work with. This is done by right clicking on the datastore name in the **Navigator** pane and selecting **Rename**.

In Figure 41, the **vsanDatastore** has been selected in the left pane. In the right pane, the **Related Objects > Hosts** tab shows the hosts and cluster associated with this VSAN.

Name	State	Status	Cluster	Consumed
esx01.vsphere.local	Connected	Normal	vSAN-AFA	2
esx02.vsphere.local	Connected	Warning	vSAN-AFA	0
esx03.vsphere.local	Connected	Warning	vSAN-AFA	0
esx04.vsphere.local	Connected	Warning	vSAN-AFA	0

Figure 41 VSAN Hosts page

Additional VSAN monitoring and management is done on the **Summary**, **Monitor**, and **Configure** tabs.

9.4 Check VSAN health and resolve issues

A VSAN health check is run as follows:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, select a cluster such as **vSAN-AFA**
3. In the center pane, select Monitor > Virtual SAN > Health> Retest
4. Verify all health tests pass as shown in Figure 42.

Test Result	Test Name
✓ Passed	Cluster
✓ Passed	Hardware compatibility
✓ Passed	Network
✓ Passed	Data
✓ Passed	Limits
✓ Passed	Physical disk
✓ Passed	Performance service

Figure 42 Virtual SAN health monitoring test results

If all tests pass, repeat for remaining clusters that have a VSAN configured. If there are warnings or failures, see the following sections to resolve common issues.

After all tests have passed on all VSANs, proceed to Section 9.5.

9.4.1 Failure: Virtual SAN HCL DB up-to-date

If this error is seen, select the failed test, and do one of the following to update the VMware VSAN Hardware Compatibility List (HCL):

- Online option: Click the **Get latest version online** button that appears when the failed HCL test is selected. When the file has been installed, click **Retest**. The test should pass.
- Local File option: If unable to connect online, you can upload from a local file. The HCL DB is a .json file available at <http://partnerweb.vmware.com/service/vsan/all.json>. Download the file to a

workstation. In the web client, click the **Upload from file** button and follow the prompts. Click **Retest**. This test should pass.

9.4.2 Warning: Controller Driver / Controller Release Support

This error may be seen with the PERC H730 (in R630 servers) or PERC FD33xD (in FC430 servers). If so, the drivers need to be updated per the following VMware Knowledge Base article: [Best practices for VSAN implementations using Dell PERC H730 or FD332-PERC storage controllers \(2109665\)](#)

Note: A step-by-step driver update video is available here: [Updating FD332 PERC driver in VMware ESXi on PowerEdge FX2 chassis](#). The same procedure applies to the PERC H730 in R630 servers.

Install the updated driver on all hosts in the VSAN cluster and reboot the hosts. When the hosts have come back online, click **Retest**. This test should pass.

9.4.3 Warning: Performance Service / Stats DB object

The warning should disappear after the VSAN performance service is enabled.

To enable the VSAN performance service:

1. In the web client, go to **Hosts and Clusters** and select a cluster containing a VSAN.
2. In the center pane, go to **Configure > Settings > Virtual SAN > Health and Performance**.
3. Next to Performance Service is Turned Off, click Edit. Check the Turn On box and click OK.

Repeat as needed for other VSAN-enabled clusters, then click **Retest**. This test should pass.

9.5 Verify IGMP snooping functionality

On directly connected physical switches, the `show ip igmp snooping groups vlan 44` command can be issued to verify IGMP snooping is functioning properly.

There should be three groups at any given time. Multicast address group 224.1.2.3 is the default VSAN group for master nodes. Multicast group 224.2.3.4 is the member group and should contain all ESXi host-connected interfaces (port channels 1-4 in this example).

```
FN410S-A1#show ip igmp snooping groups vlan 44
```

```
IGMP Connected Group Membership
```

Group Address	Interface	Mode	Uptime	Expires	Last Reporter
224.1.2.3	Vlan 44	IGMPv2	1w0d	00:01:36	10.44.2.54
Member Ports: Po 2, Po 3					
224.2.3.4	Vlan 44	IGMPv2	1w0d	00:01:41	10.44.2.53
Member Ports: Po 1, Po 2, Po 3, Po 4					
239.255.255.253	Vlan 44	IGMPv2	1w0d	00 :01:39	10.44.2.53
Member Ports: Po 1, Po 2, Po 3, Po 4					

10 Scaling guidance

10.1 VSAN sizing

The largest single VSAN cluster size recommended for this deployment is 32 Dell FC430 servers installed in 8 PowerEdge FX2s enclosures per rack.

While the current maximum size for a VSAN cluster is 64 nodes, this deployment uses a 32-node cluster for the following reasons:

- Limiting a VSAN cluster to a single rack avoids the increased complexity of a layer 3 multicast deployment.
- Eight FX2s chassis per rack allows for moderate power consumption.

Note: See the [VMware Virtual SAN Design and Sizing Guide](#) and the [VMware VSAN Compatibility Guide](#) for more information.

10.2 Port count and oversubscription

The following table outlines the connections for single rack with two spine switches with 10Gb interconnect speeds.

Table 4 Oversubscription Information

	PowerEdge FC430	FN410S IOM (2 per FX2s)	IOM links to leaf switches (8 FX2s per rack)
Connections	2 NIC ports	4 uplink interfaces	8 chassis * 4 = 32 uplinks
Port bandwidth	10Gb	10Gb	10Gb
Total theoretical bandwidth	2 * 10 = 20Gb	40Gb	32 * 10Gb = 320Gb per rack (160Gb per Spine switch)

This example provides for an oversubscription rate of 2:1 for connectivity.

A Dell EMC validated hardware and components

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

A.1 Switches

Qty	Item	Firmware Version
2	S4048-ON Leaf switch	DNOS 9.11(0.0)
3	N1548 Management switch	DNOS 6.3.1.8a17

A.2 PowerEdge FX2s chassis and components

This guide uses one FX2s chassis with four FC430 servers and two FD332 storage sleds in the Compute cluster.

Qty per chassis	Item	Firmware Version
1	FX2s Chassis Management Controller	1.41
4	FC430 servers. Each server contains: <ul style="list-style-type: none">• 2 - Intel Xeon E5-2695 v3 2.3GHz CPU, 14 cores• 8 - 32GB DIMMS (256 GB total)• 8 - 800 GB SAS SSD (provided by FD332 storage sled)• 2 - 16 GB Internal SD Cards• 1 - QLogic 57810 10GbE DP LOM• 1 - PERC FD33xD Dual Storage Controller• FC430 BIOS• FC430 iDRAC with Lifecycle Controller	<ul style="list-style-type: none">• -• -• -• -• 08.07.26• 25.5.0.0018• 2.3.5• 2.41.40.40 (07)
2	FD332 storage sled with 16 x 800 GB SAS SSD	-
2	FN410S IOM	DNOS 9.11(0.0)
4	Intel I350-T 1GbE DP LP PCIe adapter	17.5.10

B Dell EMC validated software and required licenses

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

B.1 Software

Item	Version
VMware ESXi	6.5.0, 4564106 - Dell EMC customized image version A00
VMware vCenter Server Appliance	6.5.0.5100
vSphere Web Client	Version 6.5.0 Build 4602587

B.2 Licenses

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

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

- VMware vSphere 6.5 Standard – 16 CPU sockets
- vCenter 6.5 Server Standard – 1 instance
- VSAN Advanced – 16 CPU sockets

VMware product licenses can be centrally managed by going to the vSphere web client **Home** page and clicking **Licensing** in the center pane.

C Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

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

C.1 Dell EMC product manuals and technical guides

[Manuals and documentation for Dell Networking N1548](#)

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

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

[Dell TechCenter Networking Guides](#)

[PowerEdge FX2 – FN I/O Module – VLT Deployment Guide](#)

[Dell EMC NSX Reference Architecture - FC430 Compute Nodes with VSAN Storage](#)

C.2 VMware product manuals and technical guides

[VMware vSphere 6.5 Documentation Center](#)

[VMware vCenter Server 6.5 Deployment Guide](#)

[vSphere 6.5 Installation and Setup](#)

[VMware Virtual SAN Design and Sizing Guide](#)

[VMware Compatibility Guide](#)

[VMware VSAN Compatibility Guide](#)

D Support and Feedback

Contacting Technical Support

Support Contact Information

Web: <http://Support.Dell.com/>

Telephone: USA: 1-800-945-3355