PowerEdge FX2s VSAN Guide

An FN I/O Module stack configuration example

Dell Networking Solutions Engineering
December 2016

Revisions

Date	Description	Authors
December 2016	Initial release	Ed Blazek, Curtis Bunch, Dennis Dadey

Table of contents

1 Introduction

This guide presents configuration steps to get a VMware Virtual SAN (VSAN) running for evaluating VMware VSAN in a PowerEdge FX2s with stacked FN410T switches. Dell EMC designed the configuration in this guide for VSAN redundancy and fault tolerance. The configuration allows a single component to fail without disabling storage service to connected VMs. This guide focuses on the VSAN network configuration and assumes that the reader knows how to install and configure VMware in a traditional VMware environment. Chapter 2 of this document summarizes the components of a VSAN proof-of-concept example using the FX2s chassis and associated server, storage and networking modules. It also provides a brief overview of VMware VSAN. Chapters 3 and 4 contain the topology and wiring connections for GbE PCIe adapters and management networks. Chapter 5 describes stack wiring and the details of the FN IOM switch. Chapters 6 through 10 cover server configuration and the ESXi VSAN environment.

1.1 VMware VSAN 6.2

VSAN aggregates the local or direct-attached storage devices of ESXi hosts presenting them as a single pool of shared storage across all hosts. VSAN introduces a converged storage and compute platform. Virtual machines run on ESXi hosts as usual. A small percentage of CPU and memory resources on each host serve the VSAN storage needs of the virtual machines.

VMware VSAN is software-defined storage (SDS) that provides a shared datastore for VM consumption. It is Enterprise-class, native storage for VMware hyper-converged software solutions. Embedded in the hypervisor, VSAN delivers high-performance, flash-optimized, hyper-converged storage for any virtualized application. It places server-attached flash devices and/or hard disks in clusters that aggregate the shared datastore across multiple hosts. The clusters provide a flash-optimized, highly resilient, shared datastore suitable for a variety of workloads, including the following:

- Business-critical applications
- Virtual desktops
- Remote IT
- Disaster recovery (DR)
- DevOps infrastructure

The example installation in this document consists of three FX2s chassis. Each chassis has two FC630 servers and two FD332 storage sleds. The chassis has two FN410T switches supporting VSAN and fault tolerant (FT) networks for network connectivity. In addition, four Intel Server Adapter i350 network interface cards support connections to the management and vMotion networks. For this application, the FN410T switches are configured in stack mode.

1.2 Typographical conventions

This document uses the following typographical conventions:

`Monospace text`	CLI examples
`Bold monospace text`	Commands entered at the CLI prompt
`Underlined monospace text`	CLI examples that word wrap (text is entered as a single command)
`Italic monospace text`	Variables in CLI examples

2 Hardware overview

This chapter briefly describes the primary hardware used to validate this deployment. Appendix B provides a complete listing of hardware validated for this guide.

2.1 PowerEdge FX2s enclosure

The PowerEdge FX2s enclosure is a 2-rack unit (RU) computing platform chassis that contains cooling fans, power supplies and a Chassis Management Controller (CMC). See Figure 1. There are FX2s chassis options for two FC830 full-width servers, four FC630 half-width servers or eight FC430 quarter-width servers. An enclosure provides slots for a combination of servers and storage sleds. The FX2s enclosure in this guide contains two FC630 servers (Section 2.1.1) and two FD332 storage sleds (Section 2.1.2).

Figure 1 Dell PowerEdge FX2s (front) with one PowerEdge FC630 server installed

The back of the FX2s enclosure includes two I/O networking modules (IOMs) slots and eight PCIe expansion slots. See Figure 2:

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

2.1.1 PowerEdge FC630 server

The PowerEdge FC630 server is a half-width, two-socket server. See Figure 3: The deployment in this guide uses two FC630 servers installed in the top half of the FX2s enclosure. In combination with the FD332 storage sleds, they form the compute cluster.

Figure 3 PowerEdge FC630

2.1.2 PowerEdge FD332 Storage Sled

The PowerEdge FD332 is a half-width, direct-attached storage sled with up to 16 drives. See Figure 4. It combines with FC-series servers to build flexible storage solutions. This deployment uses two FD332 storage sleds installed in the bottom half of the FX2s enclosure.

Figure 4 PowerEdge FD332

2.1.3 Dell PowerEdge FN410 I/O Module

The FN410 I/O Module (IOM) is available in two models: the 410S, see Figure 5 and 410T, see Figure 6. Dell EMC designed the IOM specifically for the PowerEdge FX2s converged-infrastructure chassis, part of the PowerEdge FX architecture. The FX2s converged infrastructure supports up to two FN IOMs per chassis, including eight 10GbE internal ports and four external ports per IOM.

Figure 5 PowerEdge FN410S

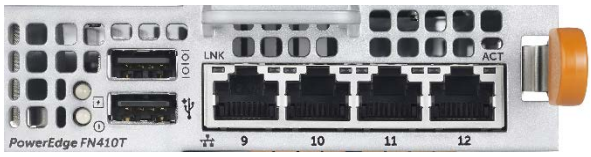

Figure 6 PowerEdge FN410T

2.2 Dell Networking S3048-ON

The S3048-ON is a 1RU switch with forty-eight 1GbE base-T ports and four 10GbE SFP+ ports. This guide uses redundant S3048-ON switches for management and VMware vSphere vMotion traffic. See Figure 7:

Figure 7 Dell Networking S3048-ON

3 Topology

This chapter provides an overview of the physical and virtual topology this deployment uses.

3.1 Networks

There are four distinct networks in this example.

- Storage and FT network – Up to six FX2s chassis with stacked FN410 switches form the network.
- Management network – Each FC630 has dual 1Gbps links to the Out Of Band (OOB) management switch.
- VMware vMotion network – Each FC630 has two 1Gbps links to the OOB management switch.
- Production network – Each FC630 has two dual-port 1GbE NICs connecting to a production network.

Note: The production and management network designs are beyond the scope of this document.

3.2 Storage network

The storage network uses six IOM switches to provide a closed, layer 2 switching environment. These six switches are split equally into two switch-stacks. Each stack member is installed into an FX2s chassis (all three FX2s chassis). The result is two stacks, Stack A1 and Stack A2. Inter-stack link (ISL) configurations provide connectivity between the stacks, which are not shown. See Figure 8:

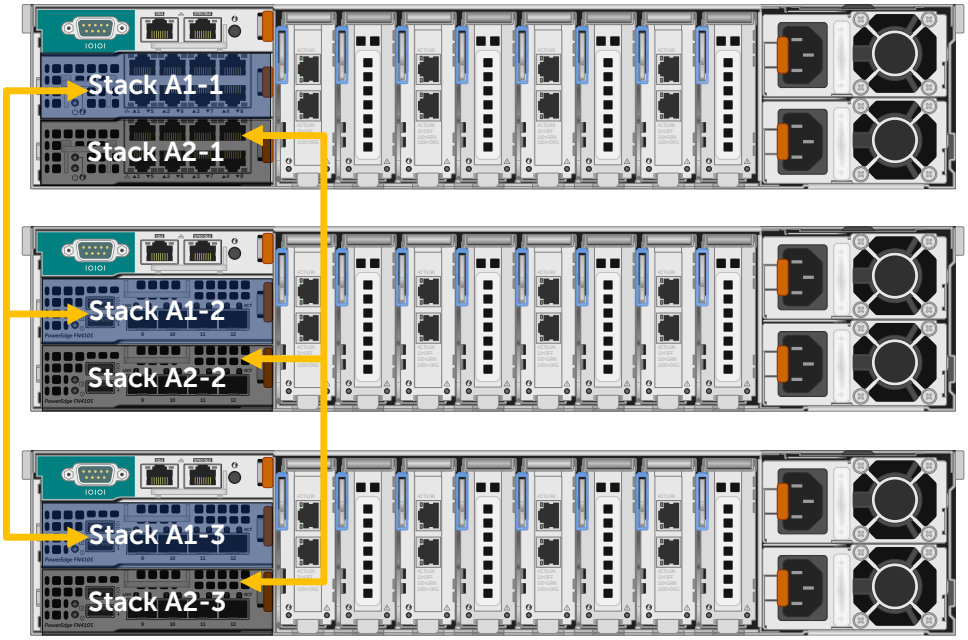

Figure 8 Physical storage network

3.3 Management network

This guide uses a redundant network for management traffic, which is isolated from the storage network. An S3048-ON switch installed at the top of rack (ToR) provides connectivity to the management network.

Each FX2s chassis has four add-in, dual-port Gb PCIe network adapters for ESXi host management and a built-in CMC for OOB management. Each PCIe network adapter connects internally to an FC630 server. See Figure 9:

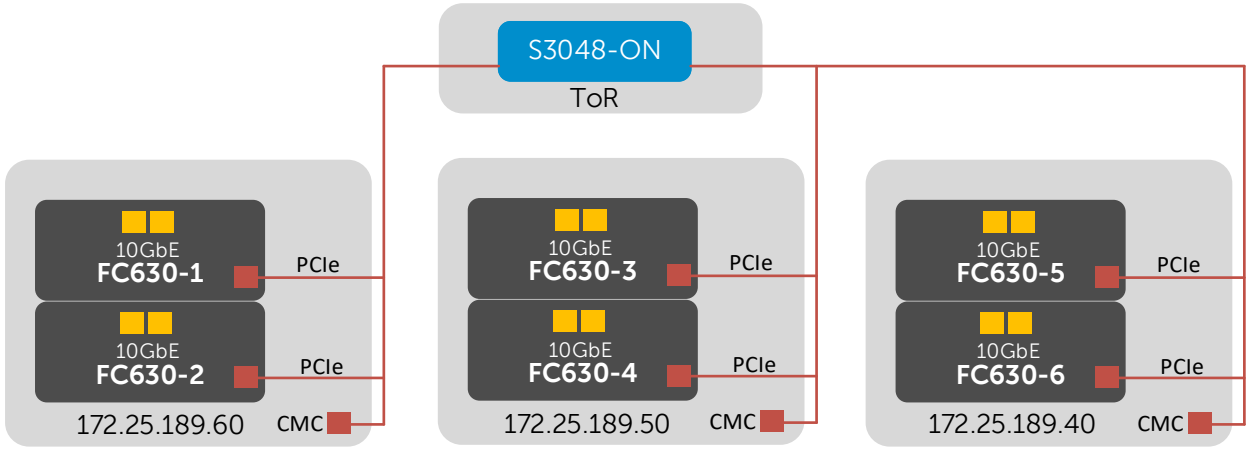

Figure 9 Physical layout of CMC and ESXi management interfaces

4 Network connections

This chapter details the physical network connections made for the vMotion and management networks. Administrators physically connect the FN410 switches after completing initial switch configuration in chapter 5.

4.1 Management connections

For management traffic, each FC630 server connects to a pair of S3048-ON switches via Intel i350-t add-in adapters in the FX2s chassis. The FX2s CMC management port also connects to the management TOR. Below, Figure 10 shows four Intel i350 cards per FX2s chassis. The left two cards support slot 2 while the right two cards support slot 1. This provides redundant connections for each server.

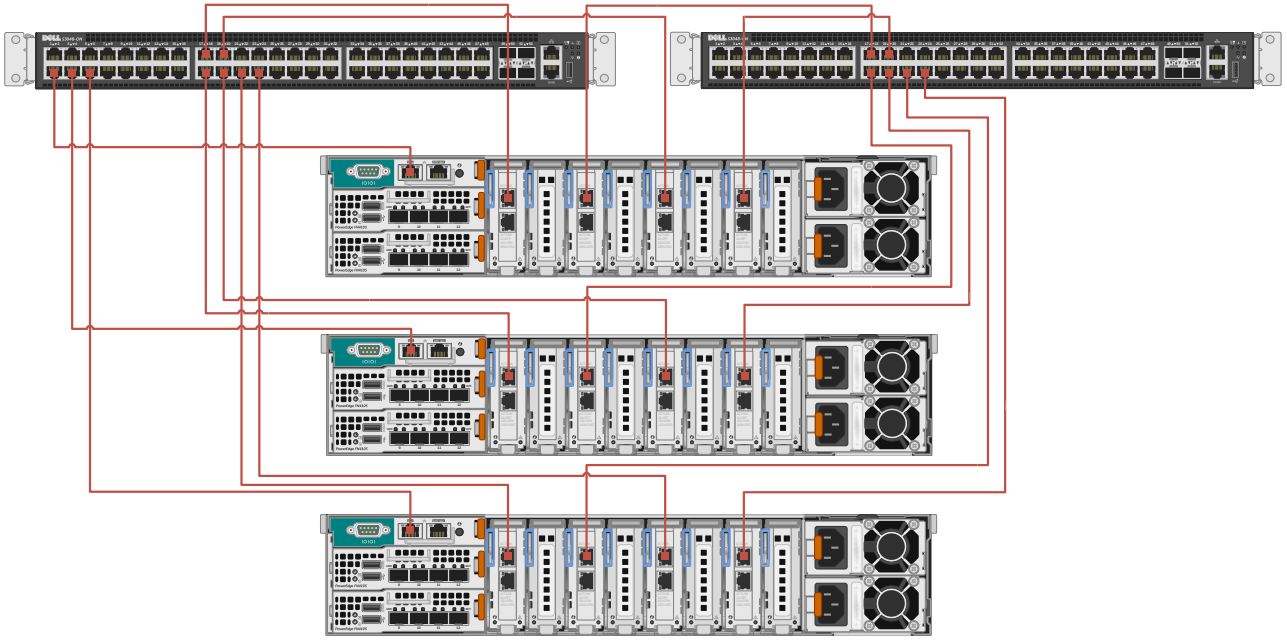

Figure 10 Physical cabling management connections

4.2 vMotion connections

The two S3048-ON switches are used to handle management connections described in section 4.1 are also utilized for infrequent vMotion traffic. Figure 11 shows these connections in blue.

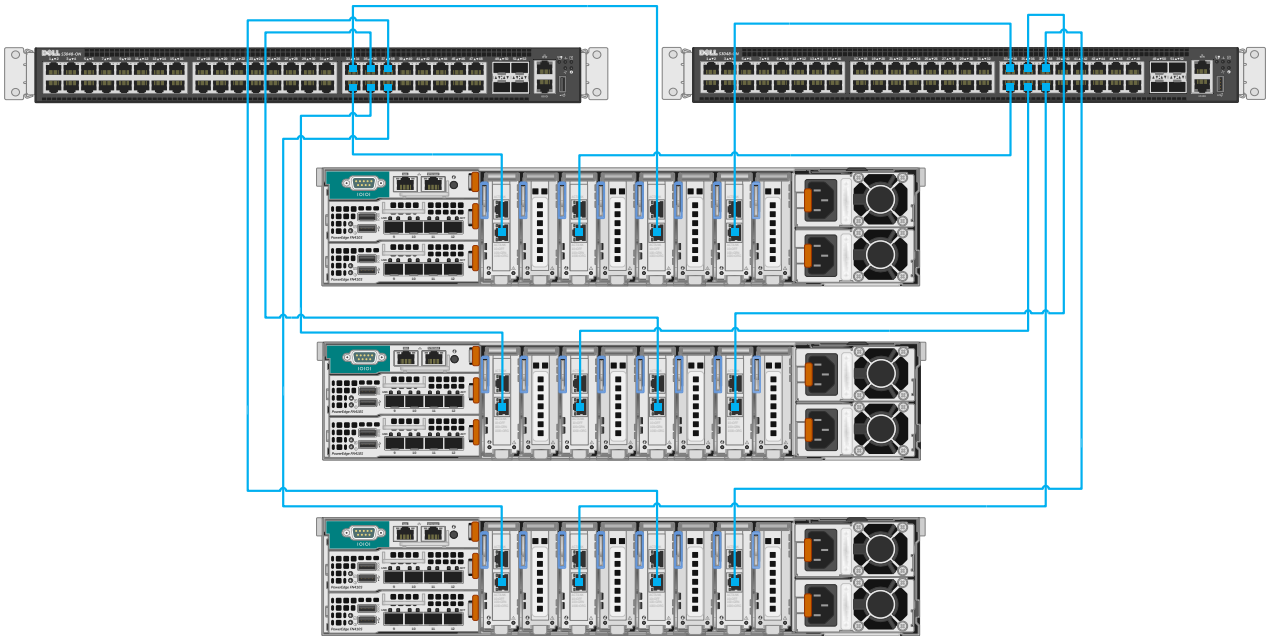

Figure 11 Physical cabling for vMotion connections

5 Physical switch configuration

This chapter contains switch configuration details with explanations for the steps involved.

5.1 Factory default settings

The configuration commands in the sections below assume that the switches start at their factory default settings. Reset any switch in this guide to factory defaults using the following commands:

```
Dell#delete startup-config.bak
Dell#restore factory-defaults stack-unit unit# clear-all
Proceed with factory settings? Confirm [yes/no]:yes

********************************************************************
*  Warning - Restoring factory defaults will delete the existing   *
*  startup-config and resets all persistent settings (stacking,    *
*  fanout, etc.) and boot environment variables (boot config, console *
*  baud rate, management interface settings, etc.)                 *
*  After restoration the unit(s) will be powercycled immediately.  *
*  Proceed with caution !                                          *
********************************************************************
```

The preceding procedure restores the switch to its factory default settings and leaves the switch ready for configuration.

5.2 FN IOM switch configuration

Table 1 below contains configuration details related to each of the FN IOM switches discussed in the next section. For example, the configuration for switch A1-1 uses a stack-unit value of 0 with a priority of 10. Priority is set to determine which switch in the stack will become the management unit of the stack. The configuration for switch A2-3 uses a stack-unit value of 2 with a priority of 6. In normal stack operation conditions stack unit 0 should always be the management switch, stack unit 1 the standby management switch and all other switches assigned to be member units.

Table 1 Stack unit and priority table

Source Switch	Chassis Number	Stack Unit Number	Priority	Unit Type
A1-1	1	0	10	Management
A1-2	2	1	8	Standby
A1-3	3	2	6	Member
A2-1	1	0	10	Management
A2-2	2	1	8	Standby
A2-3	3	2	6	Member

5.2.1 Stack Mode Configuration

This configuration requires the FN IOM switches to be in stack mode. The following section outlines the configuration commands for the FN IOM switches. The switches start at their factory default settings per section 5.1.

After the FN IOM switches boot to their default settings, use the following commands to place them in stack mode:

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

% You are about to stack your IOA module in 10G uplink Mode. Please reload
the IOA and then plug in the stacking cable to effect the changes.

Dell(conf)#stack-unit 2 priority 6
Dell(conf)#uplink-state-group 1
Dell(conf-uplink-state-group-1)#no enable
Dell(conf-uplink-state-group-1)#end
```

Note: The preceding configuration steps are specific to switch FN IOM-A1-3. The configuration for the remaining five switches is similar and a forced power-cycle of the switch takes place for renumbered switches. Unit number 0 is default skip renumber step for Unit 0 switches. See attached switch configurations for further information.

```
Dell#write
Dell#reload
```

Repeat the steps in this section for the remaining five switches. Use Table 1 above and the attached switch configuration files to complete the configuration. After all switches have the appropriate, assigned stack unit numbers and priorities, complete physical cabling.

5.3 Storage network connections

Do not physically cable the FN IOM switches before completing the previous section. If completed out of order, the switch stack may not initiate properly, which could lead to the wrong switch being considered the management unit.

5.3.1 Stack A1

The top FN IOM switches from each FX2s chassis comprise Stack A1. Their wiring uses a stack ring topology. See Figure 12. Table 2 contains the information required to wire the stack in this configuration with labels A through C to identify the cable in Figure 12.

Note: Figure 12 omits Stack A2 for clarity. Section 5.3.2 describes the configuration of Stack A2.

Table 2 Physical cabling for Stack A1

Source switch	Source interface	Destination switch	Destination interface	Label
A1-1	Te1/9	A1-2	Te1/10	A
A1-2	Te1/9	A1-3	Te1/10	B
A1-3	Te1/9	A1-1	Te1/10	C

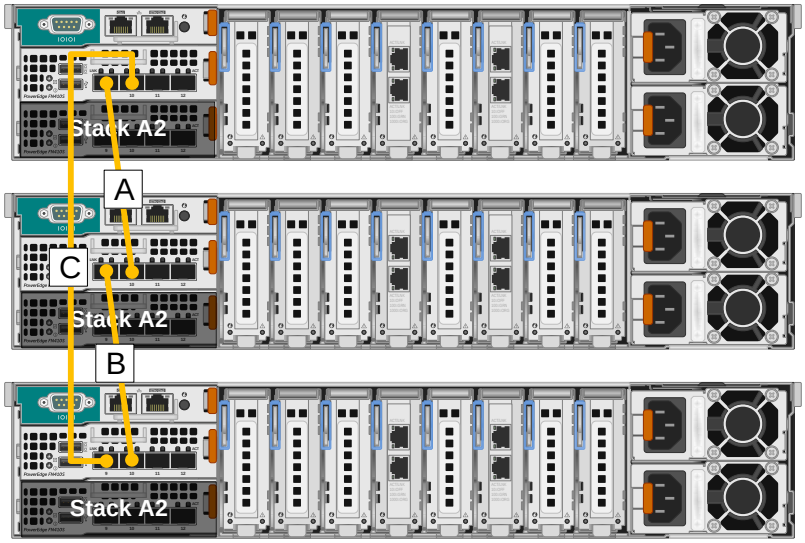

Figure 12 Physical cabling for Stack A1

5.3.2 Stack A2

The bottom FN IOM switches from each FX2s chassis comprise Stack A2. Their wiring uses a stack ring topology. See Figure 13. Table 3 contains the information required to wire the stack in this configuration with labels A through C to identify the cable in Figure 11.

Note: Figure 13 omits Stack A1 for clarity. Section 5.3.1 describes the configuration of Stack A1.

Table 3 Physical cabling for Stack A2

Source Switch	Source Interface	Destination Switch	Destination Interface	Label
A2-1	Te1/9	A2-2	Te1/10	D
A2-2	Te1/9	A2-3	Te1/10	E
A2-3	Te1/9	A2-1	Te1/10	F

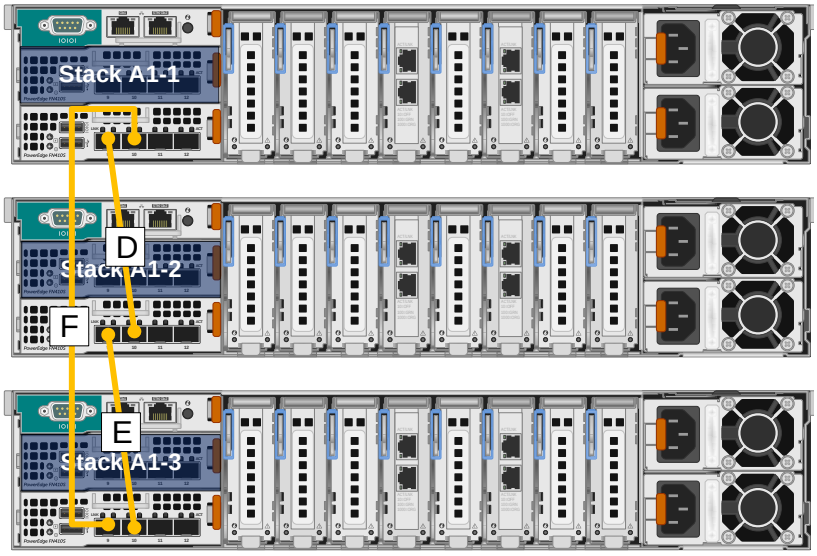

Figure 13 Physical cabling for Stack A2

5.3.3 Interswitch Links

To provide communication between Stack A1 and Stack A2, an interswitch link (ISL) connects each corresponding pair of FN IOM switches (e.g. A1-1 and A2-1). All three of these connections, by default, become part of a port channel bond. Table 4 contains the information required to connect the ISLs in this configuration, with labels G through I identifying the physical cables. See also Figure 14.

Table 4 Stack A1 and Stack A2 ISL connections

Source Switch	Souce Interface	Destination Switch	Destination Interface	Label
A1-1	Te1/12	A2-1	Te1/12	G
A1-2	Te1/12	A2-2	Te1/12	H
A1-3	Te1/12	A1-3	Te1/12	I

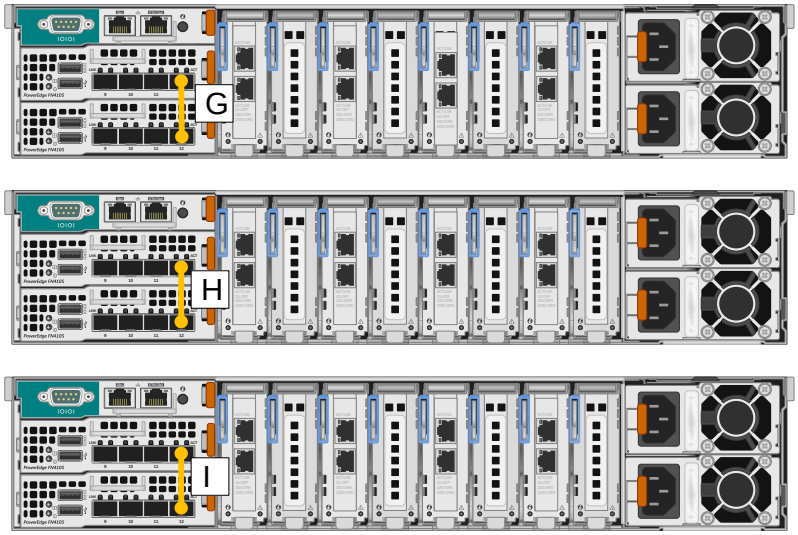

Figure 14 Physical cabling ISL connections

5.4 Verify stack connectivity

After completing the cable connections, reload all of the switches to synchronize the stack. Reloading the current master also reloads the secondary and member stack units.

Use the following commands from the master switch to reload the entire stack:

```
Dell> enable
Dell# reload
```

Use the following commands to verify stack mode connections:

```
Dell# show system stack-ports

Topology: Ring
Interface  Connection    Link Speed      Admin     Link
                           (Gb/s)        Status    Status
-----------------------------------------------------------
  0/9         1/10          10             up        up
  0/10        2/9           10             up        up
  1/9         2/10          10             up        up
  1/10        0/9           10             up        up
  2/9         0/10          10             up        up
  2/10        1/9           10             up        up

Dell# show system stack-unit all iom mode

Unit    Boot-Mode               Next-Boot
-----------------------------------------------
```

```
0        stack                 stack
1        stack                 stack
2        stack                 stack
3           Not Present
4           Not Present
5           Not Present
```

Use the following commands to display stacking details:

```
Dell#show system

Stack MAC                  : f8:b1:56:6a:ed:29

--  Unit 0 --
Unit Type                  : Management Unit
Status                     : online
Next Boot                  : online
Required Type              : PE-FN-410T-IOM - 12-port GE/TE (FN)
Current Type               : PE-FN-410T-IOM - 12-port GE/TE (FN)
Master priority            : 10
Hardware Rev               : A00
Num Ports                  : 12
Up Time                    : 24 min, 30 sec
Dell Networking OS Version : 1-0(0-5199)
Jumbo Capable              : yes
POE Capable                : no
FIPS Mode                  : disabled
Burned In MAC              : f8:b1:56:6a:ed:29
No Of MACs                 : 3

--  Unit 1 --
Unit Type                  : Standby Unit
Status                     : online
Next Boot                  : online
Required Type              : PE-FN-410T-IOM - 12-port GE/TE (FN)
Current Type               : PE-FN-410T-IOM - 12-port GE/TE (FN)
Master priority            : 8
Hardware Rev               : A00
Num Ports                  : 12
Up Time                    : 24 min, 14 sec
Dell Networking OS Version : 1-0(0-5199)
Jumbo Capable              : yes
POE Capable                : no
FIPS Mode                  : disabled
Burned In MAC              : f8:b1:56:7d:42:21
No Of MACs                 : 3
```

```
--  Unit 2 --
Unit Type                  : Member Unit
Status                     : online
Next Boot                  : online
Required Type              : PE-FN-410T-IOM - 12-port GE/TE (FN)
Current Type               : PE-FN-410T-IOM - 12-port GE/TE (FN)
Master priority            : 6
Hardware Rev               : A00
Num Ports                  : 12
Up Time                    : 24 min, 12 sec
Dell Networking OS Version : 1-0(0-5199)
Jumbo Capable              : yes
POE Capable                : no
FIPS Mode                  : disabled
Burned In MAC              : f8:b1:56:7d:42:99
No Of MACs                 : 3


Dell# show redundancy

--  Stack-unit Status  --
-----------------------------------------------
 Mgmt ID:                           0
 Stack-unit ID:                     0
 Stack-unit Redundancy Role:        Primary
 Stack-unit State:                  Active
 Stack-unit SW Version:             1-0(0-5199)
 Link to Peer:                      Up

--  PEER Stack-unit Status  --
-----------------------------------------------
 Stack-unit State:                  Standby
 Peer Stack-unit ID:                1
 Stack-unit SW Version:             1-0(0-5199)

--  Stack-unit Redundancy Configuration  --
-----------------------------------------------
 Primary Stack-unit:                mgmt-id    0
 Auto Data Sync:                    Full
 Failover Type:                     Hot Failover
 Auto reboot Stack-unit:            Enabled
 Auto failover limit:               3 times in 60 minutes

--  Stack-unit Failover Record  --
-----------------------------------------------
 Failover Count:                    0
```

```
 Last failover timestamp:          None
 Last failover Reason:             None
 Last failover type:               None

--  Last Data Block Sync Record:  --
------------------------------------------------
 stack-unit Config:         succeeded  Oct 25 2016 17:27:16
             SSMGR:         succeeded  Oct 25 2016 17:27:16
   Start-up Config:         succeeded  Oct 25 2016 17:27:16
 Runtime Event Log:         succeeded  Oct 25 2016 17:27:16
    Running Config:         succeeded  Oct 25 2016 17:27:16
           ACL Mgr:         succeeded  Oct 25 2016 17:27:16
              LACP:         no block sync done
               STP:         no block sync done
              SPAN:         no block sync done
         CRYPTOMGR:         succeeded  Oct 25 2016 17:27:16
```

5.5 Configure server ports

This section sets up VLANs on the server facing ports.

5.5.1 Configure Stack A1 ports

Configure the downstream interfaces with the next set of commands.

```
Dell#configure
Dell(conf)#interface range te 0/1,te 0/3, te 1/1,te 1/3, te 2/1,te 2/3
Dell(conf-if-range-te-0/1,te-0/3…)#description VSAN active links
Dell(conf-if-range-te-0/1,te-0/3…)#vlan tagged 20
Dell(conf-if-range-te-0/1,te-0/3…)#exit
Dell(conf)#interface range te 0/2,te 0/4, te 1/2,te 1/4, te 2/2,te 2/4
Dell(conf-if-range-te-0/2,te-0/4…)#description FT standby links
Dell(conf-if-range-te-0/2,te-0/4…)#vlan tagged 40
Dell(conf-if-range-te-0/2,te-0/4…)#exit
Dell(conf)#exit
Dell#write
```

5.5.2 Configure Stack A2 ports

```
Dell#configure
Dell(conf)#interface range te 0/1,te 0/3, te 1/1,te 1/3, te 2/1,te 2/3
Dell(conf-if-range-te-0/1,te-0/3…)#description VSAN standby links
Dell(conf-if-range-te-0/1,te-0/3…)#vlan tagged 20
Dell(conf-if-range-te-0/1,te-0/3…)#exit
Dell(conf)#interface range te 0/2,te 0/4, te 1/2,te 1/4, te 2/2,te 2/4
Dell(conf-if-range-te-0/2,te-0/4…)#description FT active links
Dell(conf-if-range-te-0/2,te-0/4…)#vlan tagged 40
```

```
Dell(conf-if-range-te-0/2,te-0/4…)#end
Dell#write
```

6 Server preparation

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

Note: Exact iDRAC console steps in this chapter 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 Ensure CPU Virtualization is enabled in BIOS

Note: CPU Virtualization is typically enabled by default in PowerEdge server BIOS. This document provides these steps for reference in case this required feature becomes 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 the Virtualization Technology setting is **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 Ensure 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 **Default** and **Yes** to load the default settings. Click **OK**.
7. To save the settings, click **Finish** and **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 Install ESXi

Dell EMC recommends using the latest Dell EMC customized ESXi .iso image available on support.dell.com. This image includes the correct drivers for your PowerEdge hardware.

Install ESXi on all servers scheduled to be part of your deployment. The example in this guide includes six FC630 servers.

A simple way to install ESXi on a PowerEdge server remotely uses the iDRAC to boot the server directly to the ESXi .iso image. Complete the following steps:

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 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 begins installation.
7. Follow the prompts to install the OS.
8. After installation completes, click Virtual Media > Disconnect Virtual Media > Yes.
9. Reboot the system when prompted.

6.4 Configure the ESXi management network connection

Be sure the host is physically connected to the management network. For this deployment, the Intel i350-t GbE add-in PCIe adapter provides this connection for all six FC630 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 not using DHCP, 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 your network uses DNS.
5. Press Esc to exit and Y to apply the changes.
6. From the ESXi main menu, select **Test Management Network**. Verify successful pings. 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 vCenter deployment and addition of hosts

7.1 Deploy vCenter Server

Cluster management and many other advanced vSphere features require vCenter Server. You can install vCenter Server as a Windows-based application or as a prepackaged SUSE Linux-based VM.

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

This guide uses VCSA installed on a PowerEdge FC630 server running ESXi. The FX2s considered to be chassis 1 includes that server.

Note: This section provides simplified VCSA installation instructions. The *VMware vCenter Server 6.0 Deployment Guide* provides detailed instructions and information. Find the guide at the following location: https://www.vmware.com/files/pdf/techpaper/vmware-vcenter-server6-deployment-guide.pdf

1. Mount the VCSA image on a Windows workstation connected to the management network.
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** (i.e. vctr01) 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** (i.e. pct.lab) and **SSO Site name** (i.e. site).
   g. Select an **Appliance size** depending on your requirements. The configuration in this guide selected **Medium (up to 400 hosts, 4000 VMs)**.
   h. Select a datastore. Optionally, if space is limited, select the **Enable Thin Disk Mode** check 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**). The configuration in this guide used a static address.
      iii. If selecting **static**, 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 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 Network Time Protocol (NTP) server. Since the ESXi hosts are not yet configured for NTP, select Synchronize appliance time with ESXi host. Section 9.5 covers ESXi host configuration for NTP.

v. **Enable SSH** is optional. Click **Next.** Click **OK** if a fully qualified domain name (FQDN) recommendation message displays.

k. Dell EMC recommends joining the **VMWare Customer Experience Improvement Program**, but this is optional. Select an option and click **Next.**

l. Review the summary page and click **Finish** if all settings are correct.

vCenter Server installs as a virtual machine on the ESXi host. When complete, the message shown in Figure 15 displays.

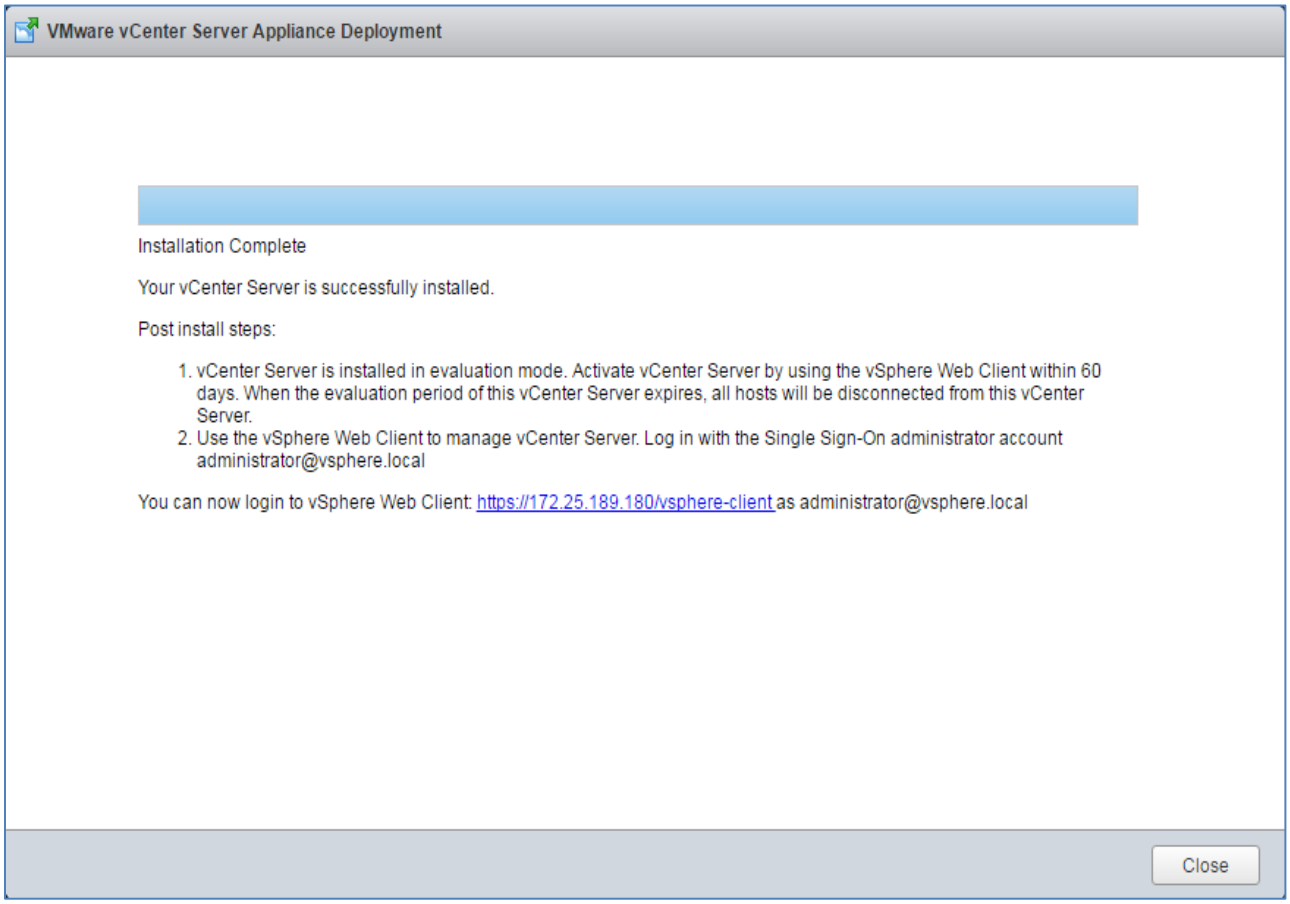

Figure 15 vCenter Server installation complete.

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 to display the web client home page. See Figure 16.

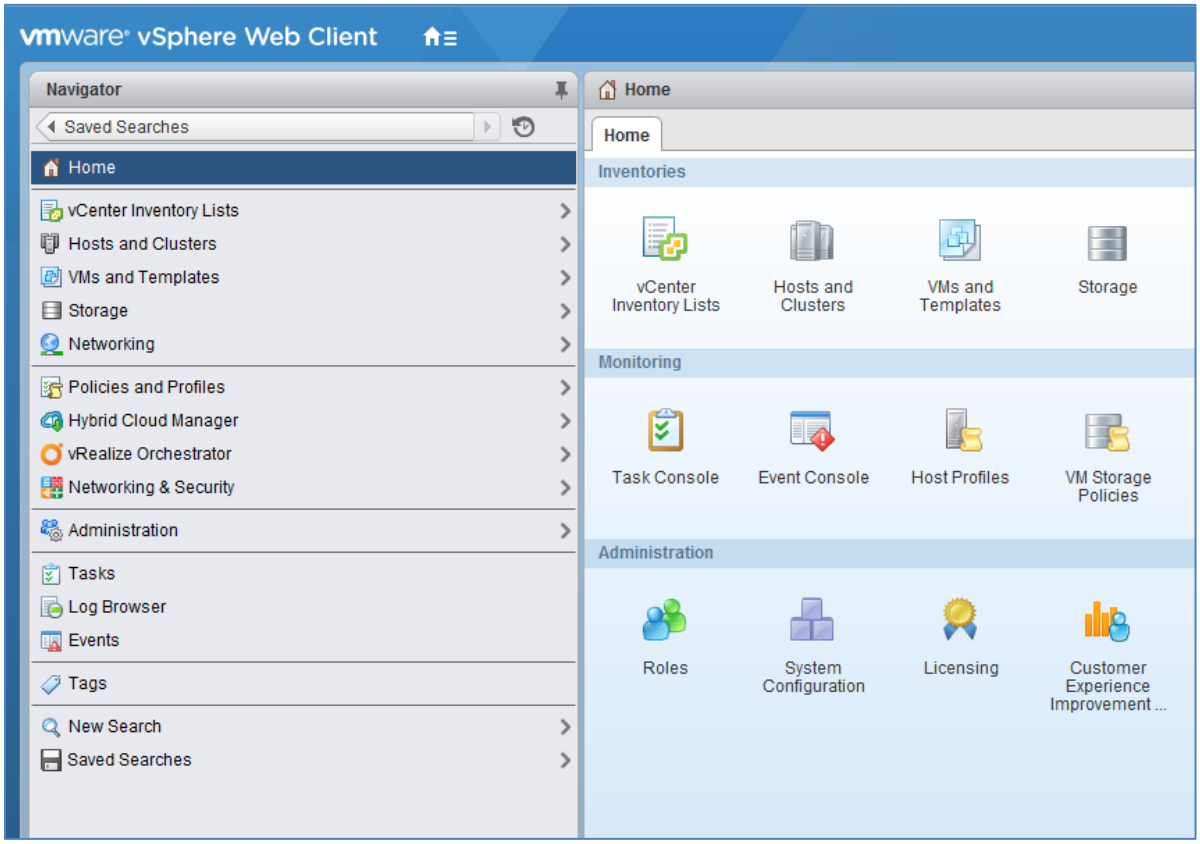

Figure 16 vSphere Web Client home page

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

7.3 Install VMware licenses

Appendix C.3 lists the VMware licenses required for this deployment. 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, complete the following steps:

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

7.4 Create a datacenter object and add hosts

You must create a datacenter object before adding hosts. This guide uses a single datacenter object named Datacenter.

To create a datacenter object, complete the following steps:

1. Go to the Web Client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, right click the vCenter Server object and select **New Datacenter**.
3. Provide a datacenter name and click **OK.**

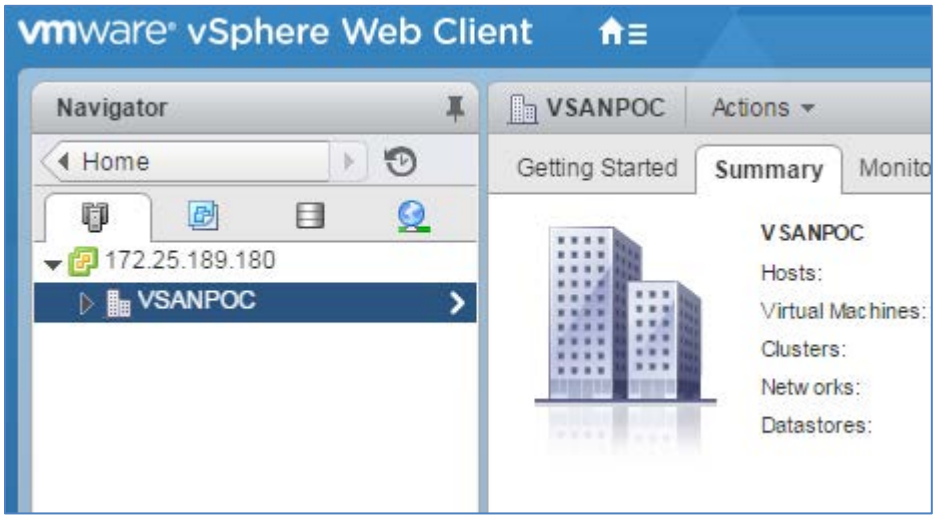

Figure 17 vSphere Web Client screen: Datacenter created

To add ESXi hosts to the datacenter, complete to the following steps:

1. Go to the Web Client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, right click the datacenter object and select **Add Host**.
3. Specify the **IP address** of an ESXi host (or the **host name** if your network is configured for DNS). Click **Next**.
4. Enter the credentials for the ESXi host and click **Next**. If a security certificate warning displays, 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 Standard 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 the datacenter and click **Next**.
9. On the **Ready to complete** screen, click **Finish**.

Repeat for all future VSAN environment servers running ESXi. This deployment example uses six FC630 servers running ESXi. When complete, the datacenter includes all ESXi hosts. See Figure 18:

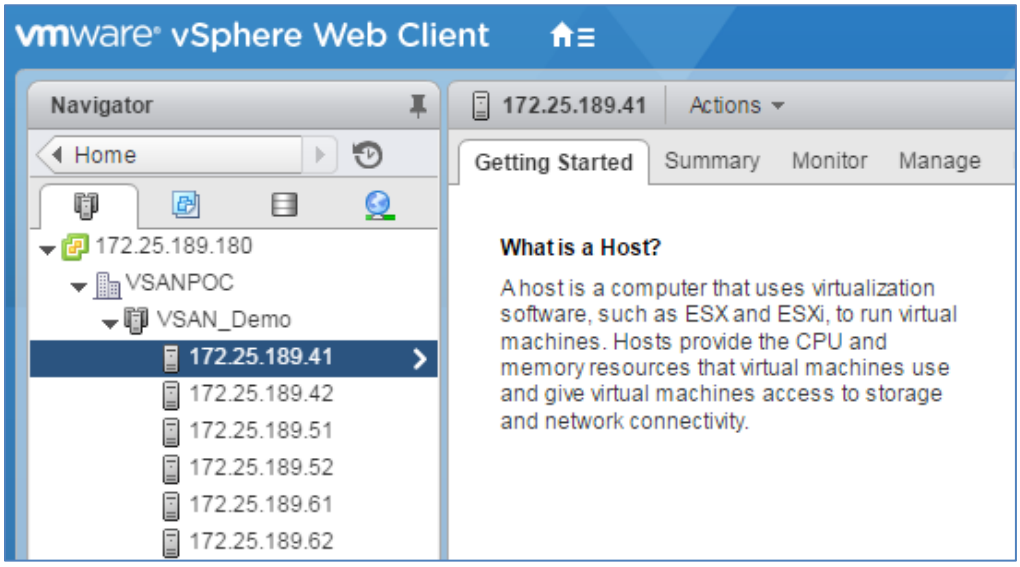

Figure 18 vSphere Web Client screen: ESXi hosts added

Some hosts may have a warning icon ( ). View the warning message by selecting the host and going to the **Summary** tab. In this case the warning indicates that the ESXi Shell and SSH are enabled. If the behavior is desired, click **Suppress Warning**. The host icon returns to normal.

7.5 Ensure hosts are configured for NTP

It is a best practice to use Network Time Protocol (NTP) on the management network to keep time synchronized in an NSX environment. Ensure NTP is configured on ESXi hosts as follows:

1. Go to the Web Client **Home** screen and select **Hosts and Clusters**.
2. In the **Navigator** pane, select a host.
3. In the center pane, go to **Manage > Settings > Time Configuration.** If the information shown is correct (see Figure 19), skip to step 7. Otherwise, continue to step 4.
4. If the information shown is incorrect, NTP has not been configured properly. Click **Edit.**
5. In the Edit Time Configuration dialog box, complete the following steps:
   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 fully qualified domain name (FQDN) of the NTP server.
   d. Click **Start** to start the NTP client. Click **OK** to close the dialog box.
6. The Time Configuration page for the host appears similar to Figure 19:

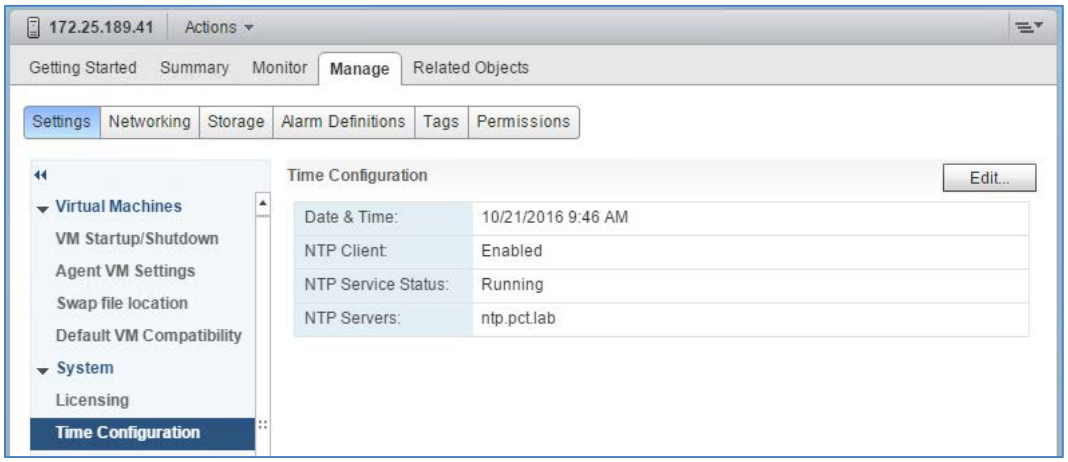

Figure 19 Proper NTP configuration on ESXi host

7. Repeat for remaining ESXi hosts as needed.

7.6 Create Clusters and add hosts

When adding a host 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). The example in this guide creates one cluster called **VSAN_Demo**.

This example adds all ESXi hosts to this cluster.

To add clusters to the datacenter, complete the following steps:

1. Go to the Web Client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, right-click the datacenter object and select **New Cluster**.
3. Name the cluster. This example names the cluster **VSAN_Demo**. Leave **DRS, vSphere HA, EVC** and **Virtual SAN** at their default settings (**Off/Disabled**). Click **OK.**

Note: In-tdepth vSphere HA and EVC cluster configuration are outside the scope of this guide. For detailed information on these features, see the VMware vSphere 6.0 Documentation. Chapter 9 of this guide covers Virtual SAN configuration.

Go to the Navigator pane. Drag and drop ESXi hosts into the appropriate clusters. Place the six ESXi hosts on FC630 servers in the VSAN_Demo cluster.

When complete, the cluster ( ) should contain its assigned hosts ( ) as shown in Figure 20:

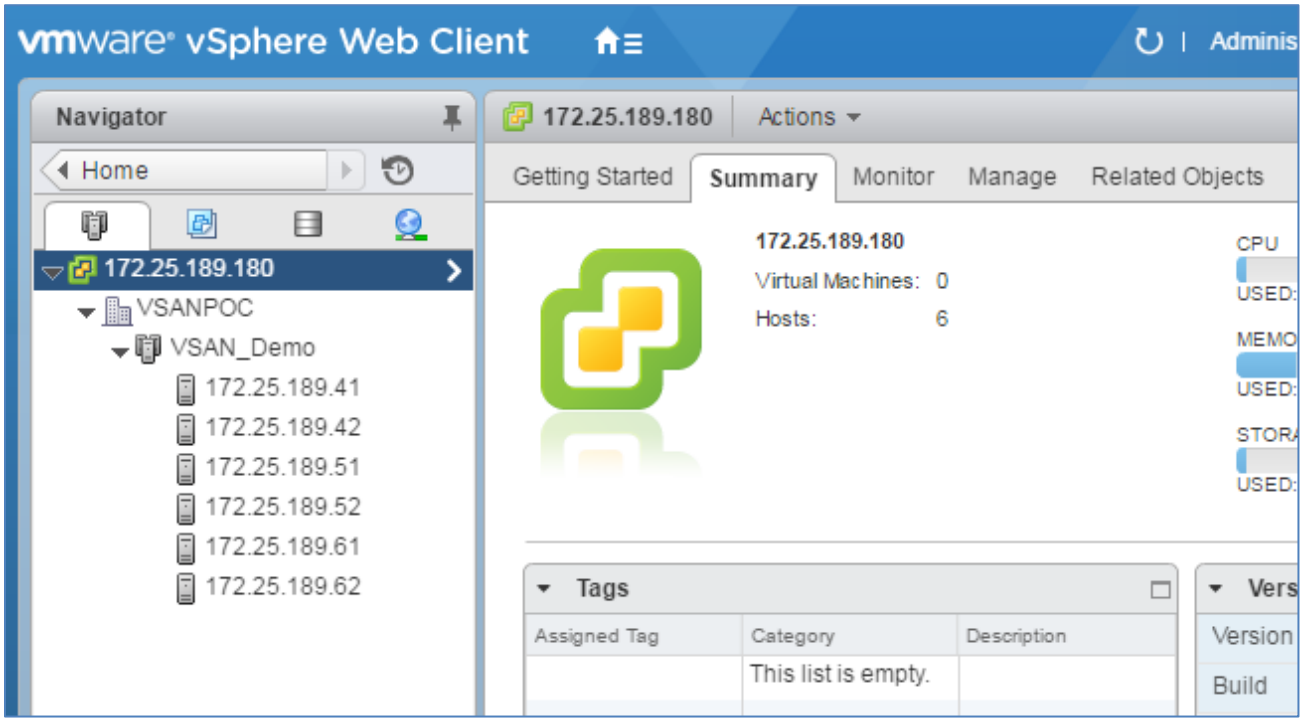

Figure 20 Clusters and hosts after initial configuration

8 Configuring vSphere standard switches

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

ESXi host installation automatically creates a standard switch named vSwitch0 on the host to provide connectivity to the management network. Add three additional standard switches to all participating hosts to provide network services required across the FN IOM switches.

To view and configure standard switches, complete the following steps:

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 Manage > Networking > Virtual switches.
3. Standard switch vSwitch0 appears in the list. Click on it to view details. See Figure 21:

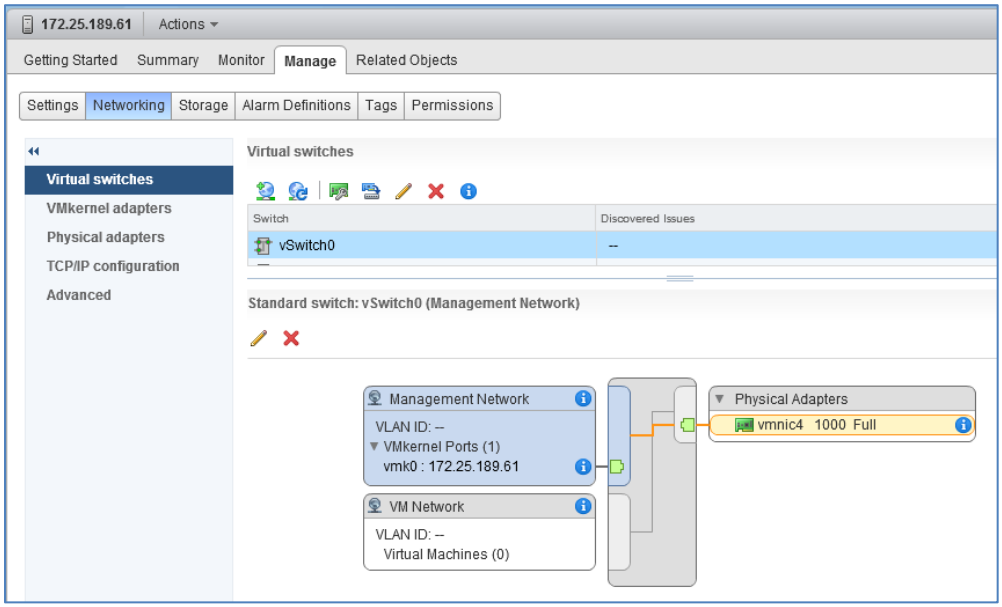

Figure 21 vSphere standard switch

8.1 Create VSAN, FT and vMotion VMKernel interfaces

In addition to the default management VSS, the example in this guide deploys three additional VMkernel services: VSAN, Fault Tolerance (FT) and vMotion. These three VMkernel interfaces use three separate standard switches. Each host has the exact same configuration. The four standard switches (including the Management switch) and corresponding VMkernels that this deployment uses are as follows:

- vSwitch0 – Management
- vSwitch1 - VSAN
- vSwitch2 – Fault Tolerance
- vSwitch3 - vMotion

To create new standard switches and associated VMkernel interfaces, complete the following steps:

1. Go to the web client Home page, select **Hosts and Clusters** and select host 172.25.189.61 in the Navigator pane.
2. In the center pane, select Manage > Networking > Virtual switches.
3. Click the  icon. The **Add Networking** dialog box opens.
4. Select VMkernel Network Adapter and click Next.
5. On the Select target device New standard switch and click **Next**
6. In the **Create a Standard Switch** dialog box, complete the following steps:
    a. Click the  icon under **assigned adapters**.
    b. In the **Add Physical Adapters to the Switch** screen:
        i. Leave the Failover order group set to **Active adapters.**
        ii. Select **vmnic0** from the **Network Adapters** list and click **OK**.
        iii. Repeat step 6 to add vmnic1 to the **Active adapters list,** then click **Next**.
7. On the **Port properties** screen, complete the following steps:
    a. Type **VSAN** in the **Network label** field.
    b. Type **20** for the **VLAN ID**.
    c. Select **Virtual SAN traffic** for the **Available service**.
8. Click **Next**.
9. In the **IPv4 settings** dialog box, complete the following steps:
    a. Select the **Use static IPv4 settings** radio button.
    b. Type **192.168.20.61** for the **IPv4 address**.
    c. Type **255.255.255.0** for the **subnet mask**.
10. Click **Next.**
11. Review the information in the **Ready to complete** dialog box and click **Finish**.

Repeat the previous steps for the remaining 11 interfaces using Table 5 below to substitute the values for steps 6, 7 and 9.

Table 5 VSAN, vMotion and FT VMkernel Information

Host	Vmnics used	Network	Available Service	IPv4 address	VLAN
172.25.189.61	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.61	20
172.25.189.61	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.61	40
172.25.189.61	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.61	30
172.25.189.62	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.62	20
172.25.189.62	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.62	40
172.25.189.62	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.62	30
172.25.189.51	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.51	20
172.25.189.51	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.51	40
172.25.189.51	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.51	30
172.25.189.52	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.52	20
172.25.189.52	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.52	40
172.25.189.52	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.52	30
172.25.189.41	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.41	20
172.25.189.41	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.41	40
172.25.189.41	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.41	30
172.25.189.42	vmnic0 / vmnic1	VSAN	Virtual SAN traffic	192.168.20.42	20
172.25.189.42	vmnic2 / vmnic3	FT	Fault Tolerance logging	192.168.40.42	40

Host	Vmnics used	Network	Available Service	IPv4 address	VLAN
172.25.189.42	vmnic4 / vmnic5	vMotion	vMotion traffic	192.168.30.42	30

8.2 Add additional standby vNIC to vSwitch0 (Management)

Installation of ESXi typically includes selection of a single management interface to allow further configuration. This environment provides two interfaces for redundant management. To add a second interface to the default VSS, vSwitch0, complete the following steps:

1. Go to the web client **Home** page, select **Hosts and Clusters** and select host **172.25.189.61** in the **Navigator** pane.
2. In the center pane, select Manage > Networking > Virtual switches.
3. In the middle pane, select **vSwitch0** and click the icon.
4. In the **Manage Physical Network Adapters for vSwitch0** dialog box, complete the following steps:
   Click the icon under **assigned adapters**.
   In the **Add Physical Adapters to the Switch** dialog box, complete the following steps:
   i. Change the **Failover order group** to **Standby adapters**.
   ii. Select **vmnic7** from the **Network Adapters** list and click **OK**.
   Click **OK**.

8.3 Configure teaming and failover on a standard switch

In this deployment, each VMkernel interface previously created uses an active/standby interface configuration. This ensures that, under normal conditions, different FN IOM switches handle VSAN and FT traffic. However, in the event of maintenance or hardware failure, the standby interfaces initialize to allow use of the secondary FN IOM switch. To configure teaming and failover, complete the following steps:

1. Go to the web client **Home** page, select **Hosts and Clusters** and select host **172.25.189.61** in the Navigator pane.
2. In the center pane, select Manage > Networking > Virtual switches.
3. In the middle pane select **vSwitch1**.
4. Under Standard switch: vSwitch1 (no item selected) select the **VSAN** network label.
5. Click the icon under Standard switch: vSwitch1 (VSAN).
6. In the VSAN – Edit Settings dialog box click **Teaming and failover**.
7. Under the **Teaming and failover** dialog box, complete the following steps:
   Click **Override** under **Failover order**.
   Select **vmnic0** and click the icon to move the vmnic under **Active adapters**.
   c. Select **vmnic1** and click the icon to move the vmnic under **Standby adapters**.
8. Leave all other settings as defaults and click **OK**.

Repeat the previous steps for the remaining 11 interfaces using Table 6 below to substitute the values for steps 3, 4 and 7.

Note: Teaming and failover configuration applies to the VMkernel interface, not the vSwitch itself.

Table 6 Teaming and failover settings for each VSS

Host	vSwitch	Interface	Active	Standby
172.25.189.61	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.61	vSwitch2	FT	vmnic3	vmnic2
172.25.189.61	vSwitch3	vMotion	vmnic5	vmnic4
172.25.189.62	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.62	vSwitch2	FT	vmnic3	vmnic2
172.25.189.62	vSwitch3	vMotion	vmnic5	vmnic4
172.25.189.51	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.51	vSwitch2	FT	vmnic3	vmnic2
172.25.189.51	vSwitch3	vMotion	vmnic5	vmnic4
172.25.189.52	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.52	vSwitch2	FT	vmnic3	vmnic2
172.25.189.52	vSwitch3	vMotion	vmnic5	vmnic4
172.25.189.41	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.41	vSwitch2	FT	vmnic3	vmnic2
172.25.189.41	vSwitch3	vMotion	vmnic5	vmnic4
172.25.189.42	vSwitch1	VSAN	vmnic0	vmnic1
172.25.189.42	vSwitch2	FT	vmnic3	vmnic2
172.25.189.42	vSwitch3	vMotion	vmnic5	vmnic4

When complete, the VMkernel adapters page for each ESXi host in the vSphere datacenter should look similar to Figure 22. View this page by going to **Hosts and Clusters**, selecting a host in the **Navigator** pane and selecting **Manage > Networking > VMkernel adapters** in the center pane.

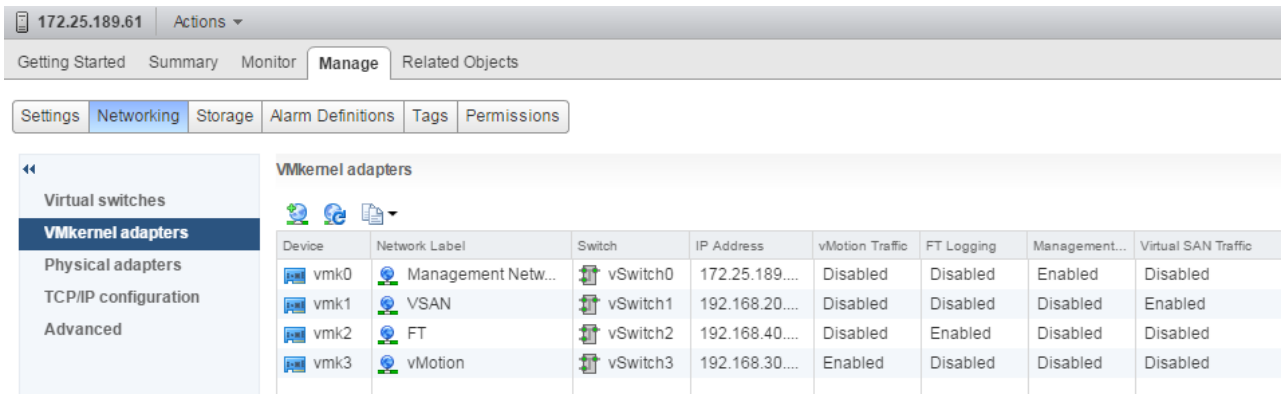

Figure 22 Host VMkernel adapters page

To verify teaming and failover, select a VMkernel adapter from above and click the **Policies** tab. See Figure 23:

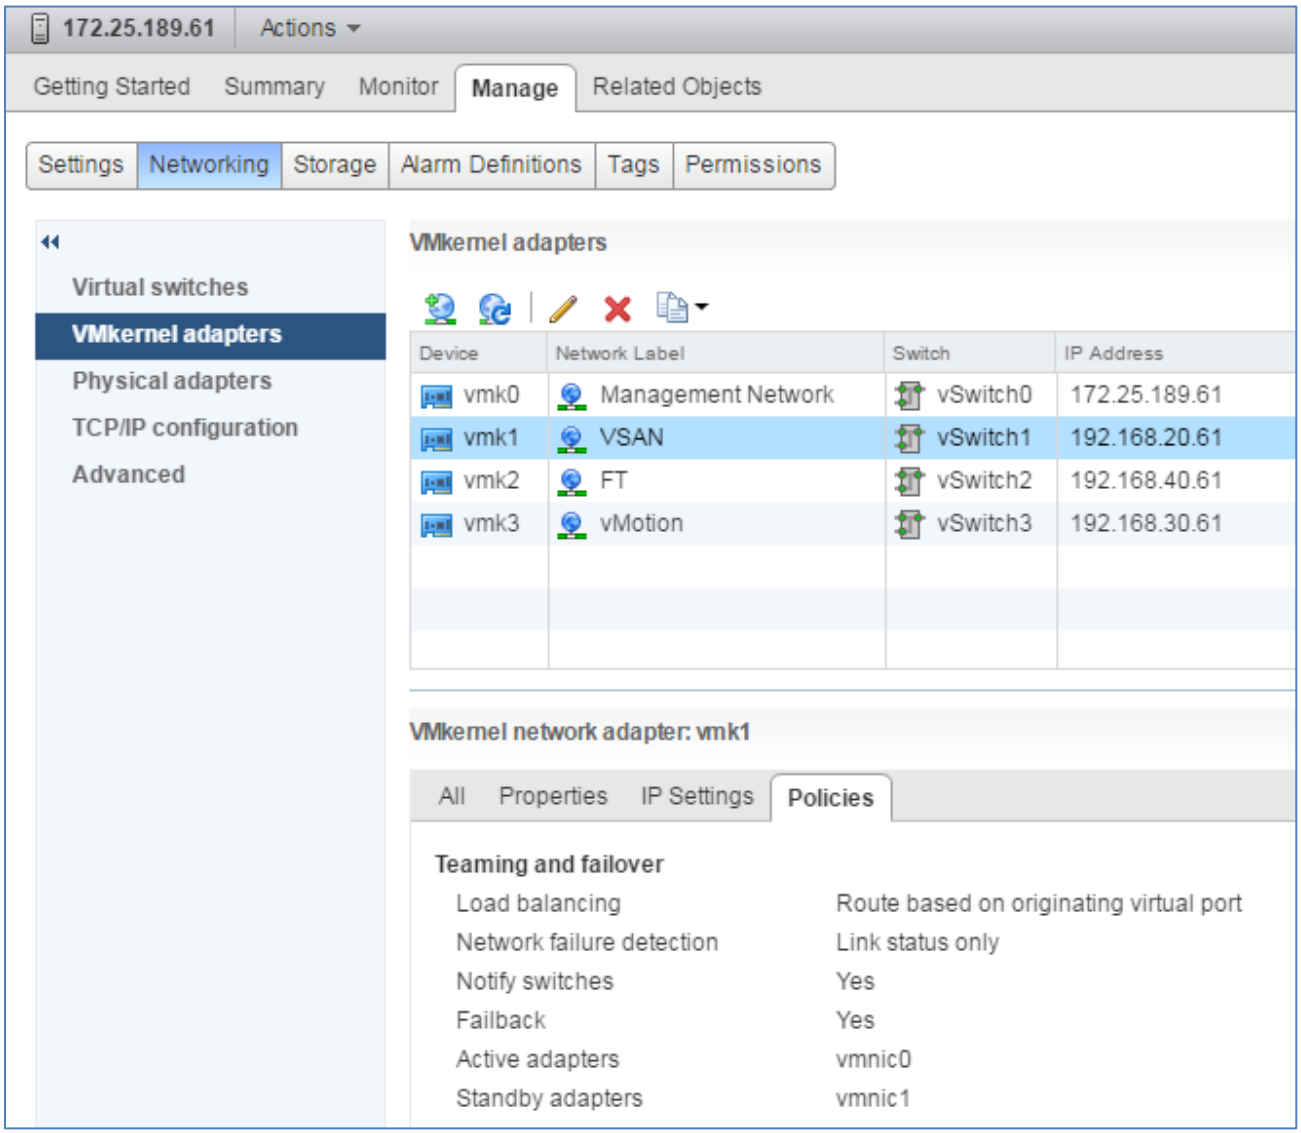

Figure 23 VSAN virtual switch policies.

9 Configure VSAN

9.1 VSAN overview

VMware Virtual SAN virtualizes local, physical storage resources of ESXi hosts, turning them into pools of storage for distribution to specific virtual machines and applications. The ESXi hypervisor includes direct implementation of VSAN.

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.

A host must meet the following criteria to participate in a VSAN:

- A VSAN 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 VSAN cluster must not participate in other clusters.
- If a host contributes its local capacity devices to the VSAN 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 VSAN section of the VMware Compatibility Guide

To configure VMware VSAN on a PowerEdge FX2s using PowerEdge FC630s, complete the following steps:

1. Confirm the physical switch configuration is in place.
2. Place all hosts in a single cluster using the web client. See Figure 20.
3. Create a VSAN-enabled VMkernel on each participating host.
4. Enable VSAN on the cluster.
5. Add disk groups to each host in the cluster.

9.2 VSAN configuration

Before proceeding, ensure that each host in the cluster has a properly configured VMkernel adapter enabled for VSAN traffic. See Chapter 8 for details.

To configure VSAN on a cluster:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, select the **VSAN_Demo**.
3. In the center pane, Select Manage > Settings > Virtual SAN > General.
4. Click **Configure** to launch the Configure Virtual VSAN wizard.
   a. Set **Add disks to storage** to **Manual**. Leave the remaining options at their defaults and click **Next**. See Figure 24,

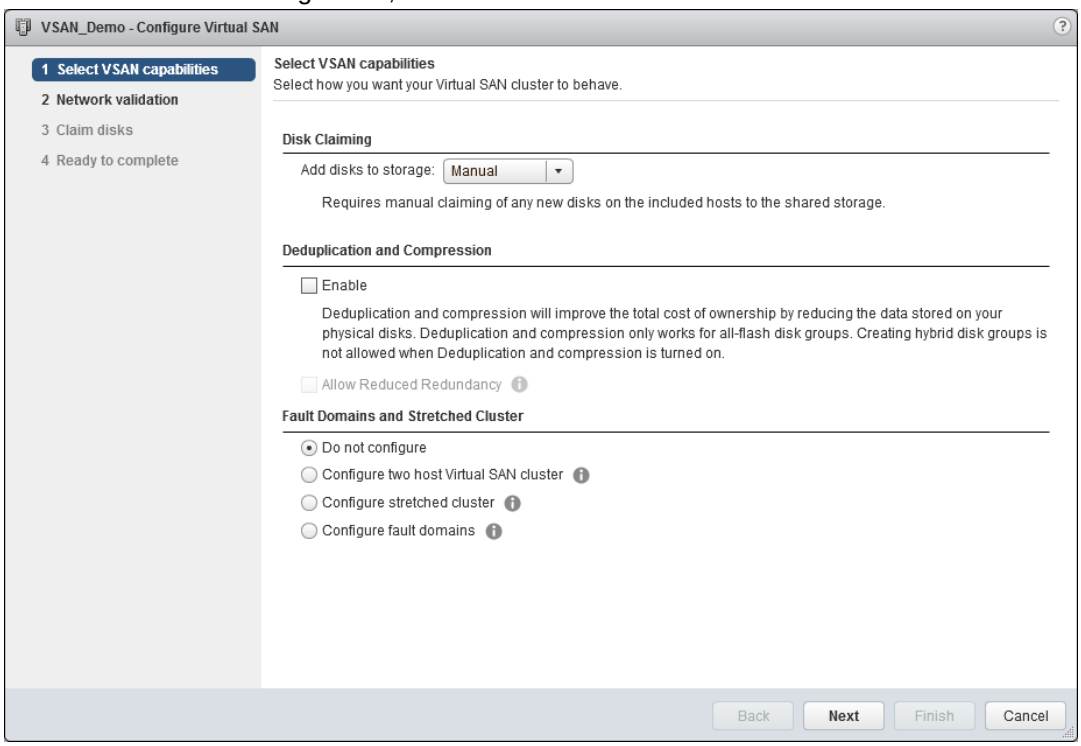

Figure 24 Configure Virtual VSAN – Select VSAN capabilities page

b. The **Network validation** page shows the VMkernel ports configured for VSAN traffic with their IP addresses and a green check in the VSAN enabled column. Click **Next.**

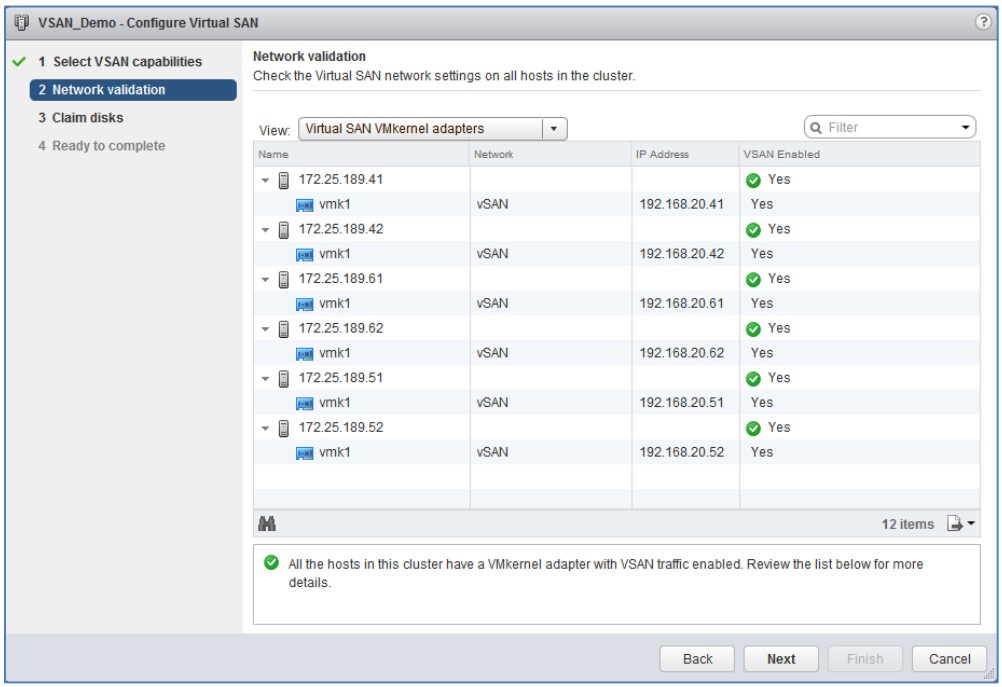

Figure 25 Configure Virtual SAN VMkernel confirmation

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

Note: VSAN configuration allows a maximum of eight disks per disk group. Each host can include up to five configured disk groups.

Figure 26 shows disk group created for host 41. Create a disk group for each host.

Select **Cache Tier** for the first disk. Select **Capacity Tier** for the remaining seven disks. After group configuration, ensure that there is a green checkmark in the **Configuration validation** field. See Figure 26. Click **Next > Finish** to apply the configuration.

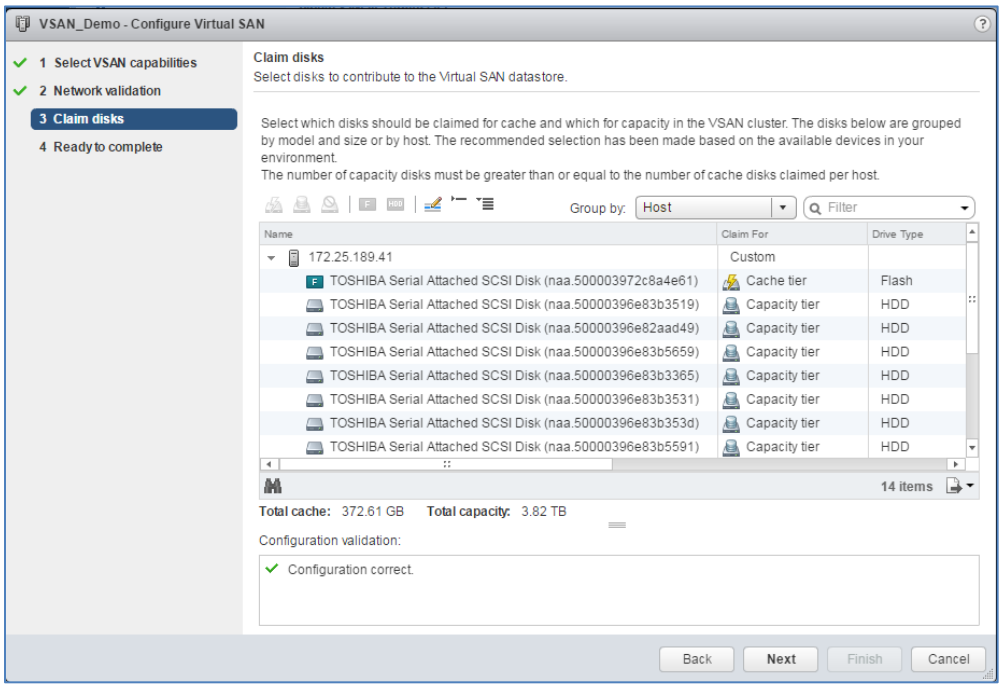

Figure 26 Configure Virtual SAN disk group management

Completion of the preceding steps automatically creates a VSAN datastore and attaches it to all hosts participating in the cluster.

Note: For more information see the Designing and Sizing a Virtual SAN Cluster section of the vSphere 6.0 online documentation.

9.3 Verify VSAN configuration

To view VSANs, go to the web client **Home** page and select **Storage**. The **Navigator** pane, in addition to local storage on each host, lists a vsanDatastore for each VSAN created.

The default datastore name is vsanDatastore. View the hosts associated with the VSAN by completing the following steps:

1. Select a vsanDatastore in the Navigator pane.
2. Select **Related Objects > Hosts** in the center pane.

Note: Rename the local datastores for usability. This example renames local datastores LDS (local datastore) plus the last octet # of the host address.

Figure 27 shows the vsanDatastore selection in the left pane. In the right pane, the Related Objects > Hosts tab shows the hosts and cluster associated with this VSAN.

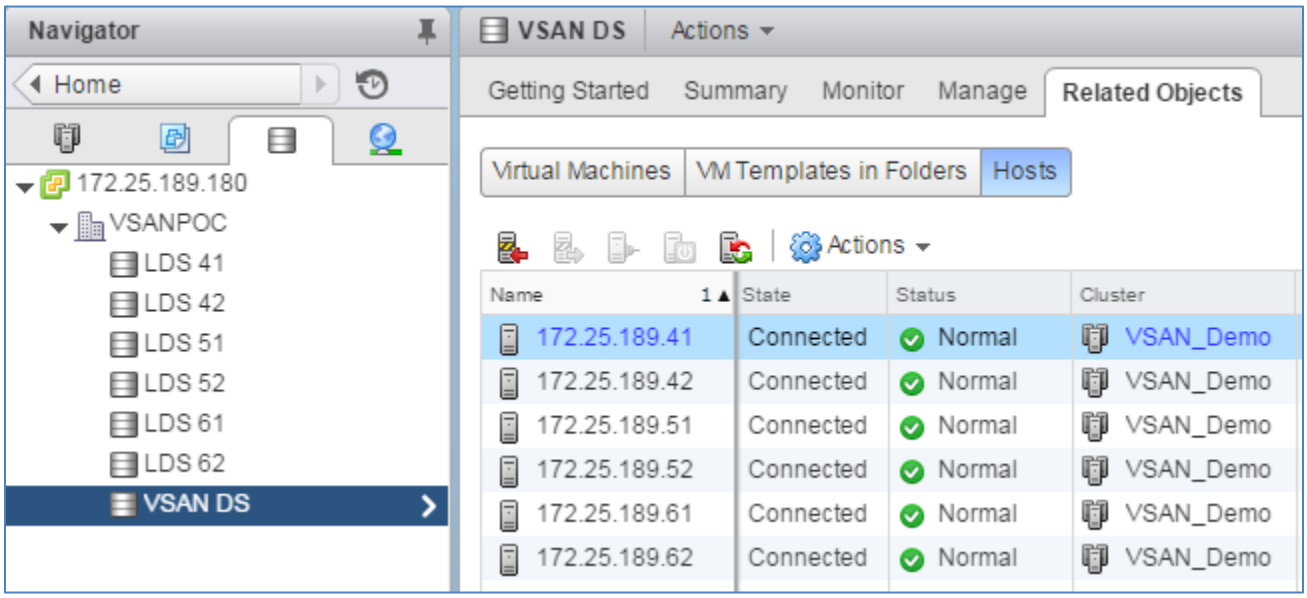

Figure 27 VSAN hosts page

Additional VSAN monitoring and management is available on the Summary, Monitor and Manage tabs.

9.4 Check VSAN health and resolve issues

Run a VSAN health check by completing the following steps:

1. Go to the web client Home screen and select **Hosts and Clusters**.
2. In the Navigator pane, select a cluster such as **VSAN_Demo**.
3. In the center pane, select Monitor > Virtual SAN > Health.

4. Verify all health tests pass. See Figure 28.

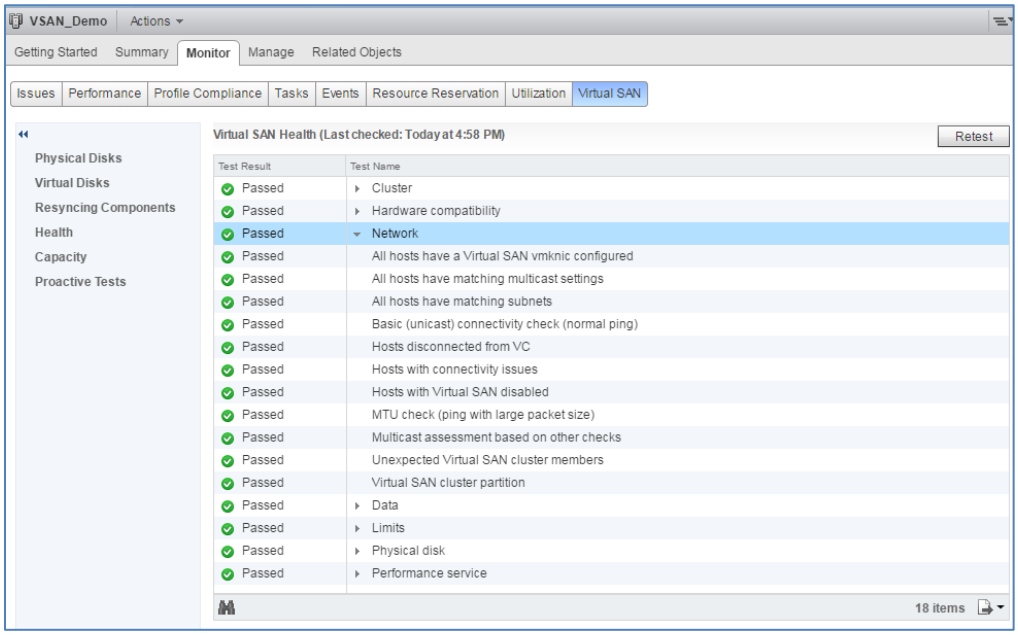

Figure 28 Virtual SAN health monitoring test results

If all tests pass, repeat for remaining clusters that have one or more VSANs configured. See the following sections to resolve common issues if there are warnings or failures.

After all tests pass on all VSANs, proceed to section 9.5.

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

Respond to this Failure error by selecting the failed test and completing the applicable option to update the VSAN HCL:

1. Online option: Select the failed HCL test and click the **Get latest version online** button. After completion of file installation, 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 **Upload from file** and follow the prompts. Click **Retest**. The test should pass.

9.4.2 Warning: Controller Driver / Controller Release Support

This error may occur with the PERC FD33xD (in FC630 servers). If so, update the drivers per the following VMWare Knowledge Base article: Best practices for VSAN implementations using Dell PERC H730 or FD332-PERC storage controllers (2109665)

Note: For video instructions on correcting this warning, please see the following YouTube video: https://www.youtube.com/watch?v=cLj2UVlFBFo

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

9.4.3 Warning: Stats DB object

The warning should disappear after enabling the VSAN performance service.

To enable the VSAN performance service complete the following steps:

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

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

9.5 Verify IGMP snooping functionality

On directly connected, physical switches, enter the **`show ip igmp snooping groups vlan 20`** command to verify proper functioning of igmp snooping.

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.

Note: The following output includes formatting modifications.

```
Dell#show ip igmp snooping groups vlan 20
IGMP Connected Group Membership
Group Address    Interface   Mode     Uptime    Expires   Last Reporter
224.1.2.3        Vlan 20     EXCLUDE  01:22:32  00:01:44  192.168.20.42
224.2.3.4        Vlan 20     EXCLUDE  01:22:44  00:02:01  192.168.20.51

Dell#show ip multicast-cam stack-unit 0 port-set 0
Vrf Id Group Address   Source Address  Vlan Id IPMC index   L2 Ports   L3 Ports
------ --------------- --------------- ------- ------------ ---------- --------
0      224.1.2.3       0.0.0.0         20      1135         1,3,11,12          -
0      224.2.3.4       0.0.0.0         20      1135         1,3,11,12          -
```

10 Configure Fault Tolerance

vSphere FT provides a higher level of business continuity than vSphere High Availability (HA). When a secondary VM replaces its Primary VM counterpart, the Secondary VM immediately takes over the Primary VM's role. This preserves the entire state of the virtual machine including running applications and data stored in memory with no need for re-entry or reload. This differs from a failover provided by HA, which restarts the virtual machines affected by a failure.

FT requires the following cluster configurations prior to its enablement:

- FT and VMotion VMkernels (See section 8.1.)
- High Availability (See section 10.1.)

10.1 Configure HA

Before enabling FT, you must enable vSphere HA at the cluster level. To enable vSphere HA at the cluster level, complete the following steps:

1. Go to the web client **Home** screen.
2. In the **Navigator** pane, select the cluster **VSAN_Demo**.
3. In the center pane select **Manage > Settings > vSphere HA** and click **Edit**.
4. In the **Edit Cluster Settings** dialog box, complete the following steps:
   a. Select **Turn on vSphere HA**.
   b. Leave all other settings as default and click **OK**.

10.2 Deploy virtual machines

This example deploys two VMs in the VSAN_Demo cluster. These two VMs represent application servers named App-VM1 and App-VM2. Figure 29 shows the added VMs:

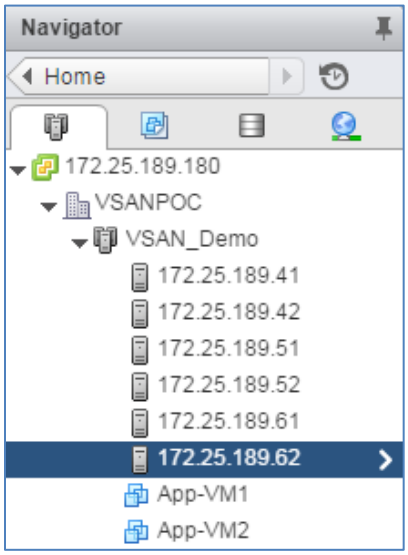

Figure 29 Hosts and Clusters view - virtual machines deployed

10.3 Configure Distributed Resource Scheduling (DRS)

10.3.1 Enable Enhanced vMotion Compatibility for Intel Hosts (optional)

Enhanced vMotion Compatibility (EVC) is optional for using Fault Tolerance (FT) with DRS. This process allows FT-enabled VMs to benefit from better initial placement.

To enable EVC, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the **Navigator** pane, select the cluster **VSAN_Demo.**
3. In the center pane, select Manage > Settings > Configuration > VMware EVC and click Edit.
4. In the **Change EVC Mode** dialog box, complete the following steps:
   a. Under **Select EVC Mode** choose **Enable EVC for Intel Hosts**.
   b. For the **VMware EVC Mode** select **Intel "Sandy Bridge" Generation** and click **OK**

Note: The EVC Mode is dependent on the Intel CPU. Please see VMware KB 1003212 for further information.

10.3.2 Enable DRS

To enable DRS on the existing VSAN_Demo cluster, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, select the cluster **VSAN_Demo.**
3. In the center pane, select **Manage > Settings > Services > vSphere DRS** and click the **Edit** button.
4. In the Edit Cluster Settings dialog box:
5. Select the DRS **Turn on vSphere DRS** check box, complete the following steps:
   a. Set DRS automation level to **Fully Automated**
   b. Leave the rest of the options as default and click **OK**

10.3.3 Create host DRS groups

A VM-host affinity rule establishes an affinity (or anti-affinity) relationship between a virtual machine DRS group with a host DRS group(s). Due to the single failure point of the FX2s chassis, it is important not to place the Primary and Secondary VMs in the same FX2s chassis.

The following criteria define two host DRS groups :

- FT_Hosts_Odd - Any host where the last octet of the management IP is odd. (e.g. 172.25.189.61)
- FT_Hosts_Even - Any host where the last octet of the management IP is even. (e.g. 172.25.189.62)

Any VM assigned to either affinity rule will run on two FC630 hosts that are installed in separate FX2s chassis.

To create a host DRS group, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, select the cluster **VSAN_Demo.**
3. In the center pane, select Manage > Settings > Configuration > VM/Host Groups.
4. On the center screen, click the **Add** button.
5. In the **Create VM/Host Group** dialog box, complete the following steps:
   a. For the name, type **FT_Hosts_Odd**.
   b. For the type, select **Host Group**.
   c. Click **Add**.
   d. Check all ESXi hosts with an odd last octet (.41, .51, .61) and click **OK**.
   e. Click **OK**.

Repeat the preceding steps for the **FT_Hosts_Even** group. Use Table 7 below to verify that the host DRS groups include the appropriate hosts:

Table 7 Host DRS group association

Host DRS group name	Member 1	Member 2	Member 3
FT_Hosts_Odd	172.25.189.41	172.25.189.51	172.25.189.61
FT_Hosts_Even	172.25.189.42	172.25.189.52	172.25.189.62

10.3.4 Create virtual machine DRS groups

To create a virtual machine DRS group, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, select the cluster **VSAN_Demo.**
3. In the center pane, select Manage > Settings > Configuration > VM/Host Groups.
4. On the center screen click the **Add** button.
5. In the **Create VM/Host Group** dialog box, complete the following steps:
   a. For the name, type **FT_VM_Odd**.
   b. For the type select **VM Group**.
   c. Click the **Add** button.
   d. Select **App-VM1** and click **OK**.

Repeat the preceding steps for the **FT_VM_Even** group. Use Table 8 below to verify that the appropriate host DRS groups include the corresponding VMs:

Table 8 VM DRS group association

VM DRS group name	Member
FT_VM_Odd	App-VM1
FT_VM_Even	App-VM2

10.3.5 Create DRS affinity rules

After defining both required DRS groups, create the following two affinity rules:

- FT_Odd
- FT_Even

Each rule combines the corresponding DRS groups to ensure that even-numbered hosts only run VMs defined as even.

To create a DRS affinity rule, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, select the cluster **VSAN_Demo.**
3. In the center pane, select Manage > Settings > Configuration > VM/Host Rules.
4. In the center pane under **VM/Hosts Rules** click **Add**.
5. In the **Edit VM/Host Rule** dialog box, complete the following steps:
   a. For the name type **FT_Odd**.
   b. For the type select **Virtual Machines to Hosts**.
   c. For the **VM Group** select **FT_VM_Odd**.
   d. Set the conditional rule to **Must run on hosts in group**.
   e. For the **Host Group** select **FT_Host_Odd**.
   f. Click **OK**.

Repeat the previous steps for the **FT_Even** affinity rule. Use Table 9 below to verify that both affinity rules contain the two correct groups.

Table 9 DRS affinity rule associations

Rule Name	VM DRS Group	Conditional	Host DRS Group
FT_Odd	FT_VM_Odd	Must run on hosts in group	FT_Host_Odd
FT_Even	FT_VM_Even	Must run on hosts in group	FT_VM_Even

Note: As part of ongoing maintenance of this solution, any additional VMs and/or hosts must be added to the appropriate DRS group manually to ensure proper initial and ongoing placement of the VMs.

10.4 Enable Fault Tolerance for virtual machines

To enable FT for VMs, complete the following steps:

1. Go to the web client **Home** screen and select **Hosts and Clusters**.
2. In the Navigator pane, right-click the virtual machine **App-VM1**.
3. Select Fault Tolerance > Turn On Fault Tolerance.
4. In the **Turn on Fault Tolerance** dialog box, complete the following steps:
   a. For each of the Secondary VM objects, click **Browse**.
   b. In the **Select a datastore** dialog box, select **vsanDatastore** and click **OK**.
   c. Ensure that all three files have the **vsanDatastore** selected for storage and click **Next**.
5. In the **Select host** dialog box select a host to place the Secondary VM and click **Next**.
6. Review your selections and click **Finish**.

Note: In step 5 above, select any host regardless of the DRS affinity rules established in section 10.3.5. This is expected behavior and DRS migrates the Secondary VM automatically to conform to affinity rules.

At this point, fault tolerance starts to create a Secondary VM for the targeted virtual machine. After completion of Secondary VM creation, navigate to the VM and view the summary tab to monitor FT status.

11 Conclusion

This guide presents the benefits of the FX2s chassis as a VSAN solution. It described a step-by-step process for assembling and configuring the components involved to evaluate VSAN for your application. As a highly integrated, hyper-converged infrastructure appliance, the Dell-FX2s chassis solution provides cost savings and complete virtualization using reliable Dell hardware. Additionally, the Dell-FX2s chassis solution provides flexibility in terms of starting small and adding additional nodes as needed. The FX2s chassis solution can easily scale up to six FX2s chassis providing up to 12 FC630 ESXi server nodes and 28TB usable VSAN storage.

A Slot-to-switch port assignments

Half-width servers – quad-port CNAs

In half-width servers with quad-port CNAs installed, the CNA ports map to two ports on each IOM. Figure 30 and Table 10 present the port mapping for half-width servers with quad-port CNAs:

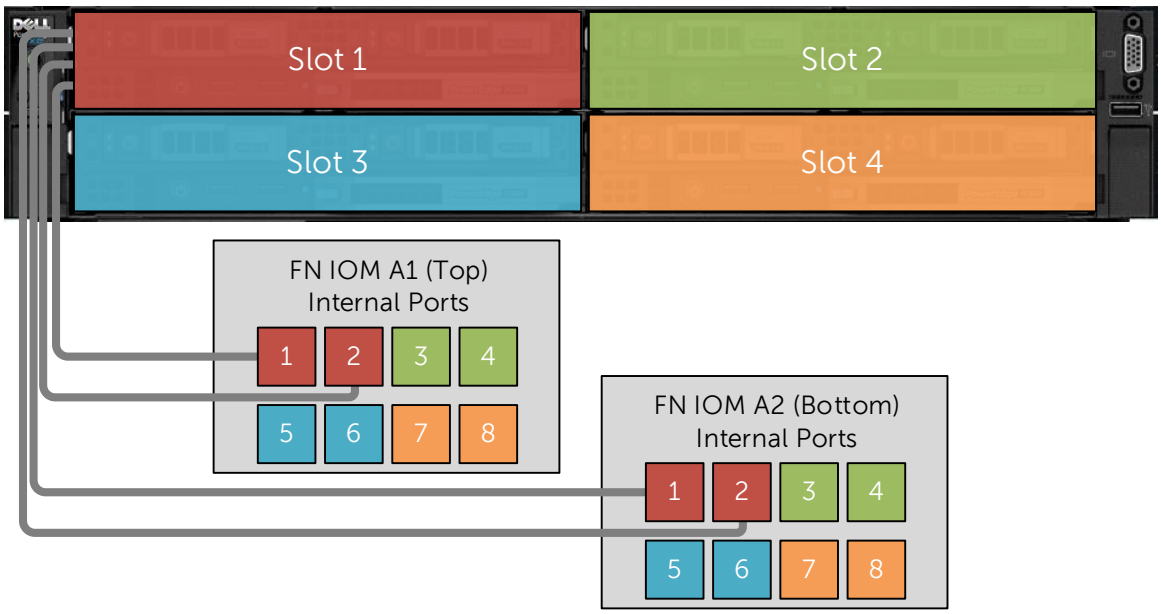

Figure 30 Half-width slots with quad-port CNAs

Table 10 Half-width slots with quad-port CNAs

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

B Dell-validated hardware and components

B.1 PowerEdge FX2s chassis

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

Qty per chassis	Item	Firmware Version
1	FX2s Chassis Management Controller	1.40.200
2	FC630 servers. Each server contains: • 2 - Intel Xeon E5-2630 v3 2.4GHz CPU, 14 cores • 8 - 32GB DIMMS (256 GB total) • 8 - 800 GB SAS SSD (provided by FD332 storage sled) • 1 - Intel 10GbE 4P x710-k bNDC LOM • 1 - PERC FD33xD Single Storage Controller FC630 BIOS FC630 iDRAC with Lifecycle Controller	- - - - • 17.5.11 • 25.4.1.0004 • 2.1.7 • 2.30.30.30
2	FD332 single controller storage sled with: • 32 - 800 SAS SSD	
2	FN410S or FN410T IOM	DNOS 9.10.0.1P13 CPLD ver 9
4	Intel I350-t 1Gb DP LP PCIe adapter	17.5.10

C Dell-validated software and required licenses

C.1 Software

Item	Version
VMware ESXi	6.0.0 Update 2 - Dell customized image version A00
VMware vSphere Desktop Client	6.0.0 build 3562874
VMWare vCenter Server Appliance	6.0.0 Update 2 - build 3634788
vSphere Web Client	6.0.0 build 3617395 (included with VCSA above)

C.2 Specific ESXi drivers

Component	Version
PERC HBA Driver	lsi-mr3-6.903.85.00-1OEM.600.0.0.2768847.x86_64.vib
Intel x710 Driver	neti40e-1.4.28-1OEM.550.0.0.1331820.x86_64.vib

C.3 Licenses

The vCenter Server uses a per-instance license. The number of CPU sockets in participating hosts determines the allocation of remaining licenses.

The topology built in this guide requires the following licenses:

- VMware vSphere 6 Enterprise – 12 CPU sockets
- vCenter 6 Server Standard – 1 instance
- VSAN Standard – 12 CPU sockets

Centrally manage VMware product licenses by clicking **Licensing** in the center pane of the vSphere web client **Home** page.

D 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.

D.1 Dell EMC product manuals and technical guides

Manuals and documentation for Dell Networking S3048-ON

Manuals and Documentation for PowerEdge FX2/FX2s and Modules

Dell TechCenter Networking Guides

D.2 VMware product manuals and technical guides

VMware vSphere 6.0 Documentation Center

VMware vCenter Server 6.0 Deployment Guide

VMware Virtual SAN Design and Sizing Guide

VMware Compatibility Guide

VMware VSAN Compatibility Guide

VMware Virtual SAN Diagnostics and Troubleshooting Manual

E PMUX Mode Configuration

This configuration requires the FN IOM switches to be in PMUX mode. The following section outlines the configuration commands given to the FN IOM switches. The switches start at their factory default settings per section 5.1.

To place FN IOM switches in PMUX mode, boot them to their default settings and enter the following commands:

```
Dell> enable
Dell# configure
Dell(conf)# stack-unit 0 iom-mode programmable-mux
Dell> exit
The stack unit number is renumbered

Dell# stack-unit 0 renumber 1
Proceed[confirm yes/no] yes
System configuration has been modified. Save? [yes/no] yes
```

Note: The following configuration details are specific to switch FN IOM-A1-2. Skip renumber step for Unit 0 switches. The configuration for the remaining five switches is similar and a forced power-cycle of the switch takes place. See attached switch configurations for further information.

Initial configuration involves setting set stack unit priority and enabling stack ports. Stack-group 0 and 1 enable te1/9 and te1/0 as respective stack interfaces.

```
Dell> enable
Dell# configure

Dell(conf)# stack-unit stack unit number priority priority
Dell(conf)# stack-unit stack unit number stack-group 0
Setting ports Te…[comfirm yes/no]: yes
Dell(conf)# stack-unit stack unit number stack-group 1
Setting ports Te…[comfirm yes/no]: yes
```

Save the configuration and reload the switch.

```
Dell(conf)#do reload

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

Repeat the steps in this section to complete the remaining five switches. Use Table 1 in Section 5.2 and the attached switch configuration files to complete the configuration.

Note: Complete physical cabling before continuing with the deployment. Reference the physical cabling diagrams from section 4 before continuing.

Initial configuration involves setting the hostname, the management interface and default gateway. Finally, disable DCB, iSCSI optimization and UFD following best practice guidelines.

```
Dell> enable
Dell# configure

Dell(conf)# hostname FN410-A1-2

Dell(conf)# interface ManagementEthernet 0/0
Dell(conf-if-ma-0/0)# ip address 172.25.189.65 255.255.255.0
Proceed with Static IP [confirm yes/no]: yes
Dell(conf-if-ma-0/0)# no shutdown

Dell(conf)# management route 0.0.0.0/0 172.25.189.254

Dell(conf)# no dcb enable
Dell(conf)# no iscsi enable
Dell(conf)# uplink-state-group 1
Dell(conf-uplink-state-group 1)# no enable

Next, configure the two required VLANs with short descriptions applied and
jumbo frames (MTU of 9216) set.

Dell(conf)# interface vlan 20
Dell(conf-if-vl-20)# description VSAN
Dell(conf-if-vl-20)# mtu 9216

Dell(conf)# interface vlan 40
Dell(conf-if-vl-40)# description FT
Dell(conf-if-vl-40)# mtu 9216
Dell(conf)# exit
```

Configure the downstream interfaces in the next set of commands. Configure the interfaces as layer 2 interfaces, enable jumbo frames and tag the appropriate VLAN interfaces.

```
Dell(conf)# interface range te0/1, te0/3, te1/1, te1/3, te2/1, te2/3
Dell(conf-if-range-te-0/1,te-0/3…)# shutdown
Dell(conf-if-range-te-0/1,te-0/3…)# description VSAN active links
Dell(conf-if-range-te-0/1,te-0/3…)# mtu 9216
Dell(conf-if-range-te-0/1,te-0/3…)# switchport
Dell(conf-if-range-te-0/1,te-0/3…)# vlan tagged 20
Dell(conf-if-range-te-0/1,te-0/3…)# protocol lldp
Dell(conf-if-range-te-0/1,te-0/3…)# no shutdown
Dell(conf-if-range-te-0/1,te-0/3…)# exit

Dell(conf)# interface range te0/2, te0/4, te1/2, te1/4, te2/2, te2/4
Dell(conf-if-range-te-0/2,te-0/4…)# shutdown
```

```
Dell(conf-if-range-te-0/2,te-0/4…)# description FT standby links
Dell(conf-if-range-te-0/2,te-0/4…)# mtu 9216
Dell(conf-if-range-te-0/2,te-0/4…)# switchport
Dell(conf-if-range-te-0/2,te-0/4…)# vlan tagged 40
Dell(conf-if-range-te-0/1,te-0/3…)# protocol lldp
Dell(conf-if-range-te-0/2,te-0/4…)# no shutdown
Dell(conf-if-range-te-0/2,te-0/4…)# exit
```

This section describes configuring the upstream interfaces to the other FN IOM stack. Use external interface tengigabitethernet 0/12 and place it in LACP-enabled port channel 128. The port channel's configuration uses jumbo frames (MTU 9216) and enables the port channel for two configured VLAN IDs (20 and 40)

```
Dell(conf)# interface range te0/11-12, te1/11-12, te2/11-12
Dell(conf-if-range-te-0/11-12…)# description backup port channel link
Dell(conf-if-range-te-0/11-12…)# port-channel-protocol LACP
Dell(conf-if-range-te-0/11-12…-lacp)# port-channel 128 mode active
Dell(conf-if-range-te-0/11-12…)# mtu 9216
Dell(conf-if-range-te-0/11-12…)# exit

Dell(conf)# interface port-channel 128
Dell(conf-if-po-128)# mtu 9216
Dell(conf-if-po-128)# switchport
Dell(conf-if-po-128)# vlan tagged 20, 40
Dell(conf-if-po-128)# exit
```

Finally, modify flow control on all upstream and downstream interfaces.

```
Dell(conf)# interface range te0/1-12, te1/1-12, te2/1-12
Dell(conf-if-range-te-0/11-12…)# flowcontrol rx off tx off
```

Save the configuration.

```
Dell(conf-rstp)# end
Dell# write
```

F Support and feedback

Contacting Technical Support

Support Contact Information	Web: http://Support.Dell.com/
	Telephone: USA: 1-800-945-3355

Feedback for this document

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