

Dell EMC Networking - Deploying VMware vSAN with OS10 Enterprise Edition

Deployment Guide

Dell EMC Networking Infrastructure Solutions
May 2018

Revisions

Date	Description	Authors
May 2018	Initial release 1.0	Ed Blazek, Colin King

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

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

Copyright © 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA May 2018.

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

Table of contents

Revisions.....	2
1 Introduction.....	4
1.1 Dell EMC Networking S4100-ON switch series.....	4
1.2 Dell EMC PowerEdge R740xd server	5
2 Objective.....	6
3 Leaf-Spine overview.....	7
3.1 Physical vSAN cluster deployment within SDDC	7
3.2 vSAN networking recommendations	8
4 Switch configuration	9
4.1 Check switch firmware version	9
4.2 Factory default configuration	9
4.3 Switch-port profiles	9
4.4 Global switch settings.....	10
4.5 VLT configuration.....	10
4.6 VLAN configuration.....	11
4.7 Node-facing configuration.....	12
4.8 Leaf switch pair uplinks	13
5 VMware vSAN deployment	14
5.1 VMware ESXi and vCenter Server prerequisites	14
5.1.1 Prepare ESXi hosts	14
5.1.2 Create vCenter Server datacenter and cluster.....	15
5.1.3 Create a vSphere Distributed Switch (VDS).....	17
5.2 vSAN deployment.....	23
A Validated components.....	26
B Accessing the RS-232 console port	27
C Product Manuals and technical guides	28
D Support and feedback	29

1 Introduction

The Dell EMC OS10 Enterprise Edition operating system is a native Linux-based network operating system that has a wholly disaggregated software architecture. OS10 decouples the base software from the layer 2 and layer 3 protocol stack and services and brings forth the ability for open programmability and portability. This document outlines the use of OS10 Enterprise Edition running on Dell EMC Networking S4100-ON series switches.

1.1 Dell EMC Networking S4100-ON switch series

The Dell EMC Networking S4100-ON switch series (S4148U, S4148F, S4148FE, S4148T, S4128F, S4128T) is a 1RU (Rack-Unit) high-density 1/10/25/40/50/100GbE, top of rack switch. S4148-ON switches are leaf switches running OS10 in this deployment guide. The S4148-ON switch operates with up to 48 dual-speed, 1/10GbE SFP+ (S4148T-ON with BaseT) ports, and four 100GbE QSFP28 ports. This switch uses non-blocking and cut-through switching architecture to provide ultra-low-latency performance.

Note: Switch-port profiles determine the interfaces that are available.

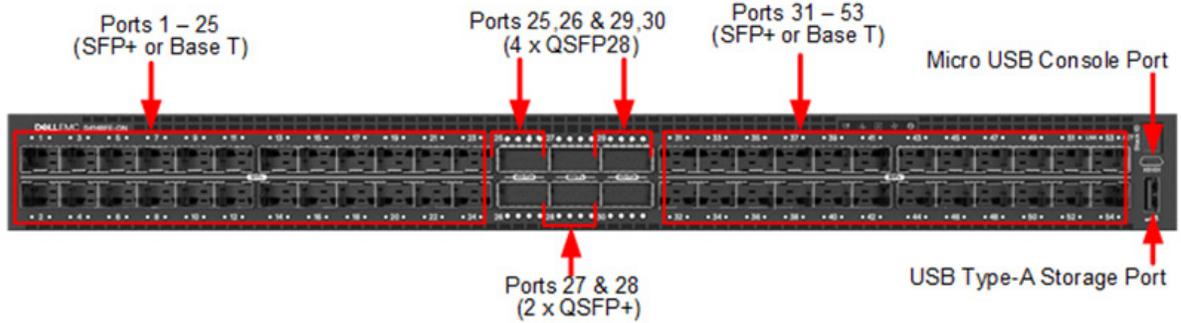


Figure 1 Dell EMC Networking S4148-ON front view

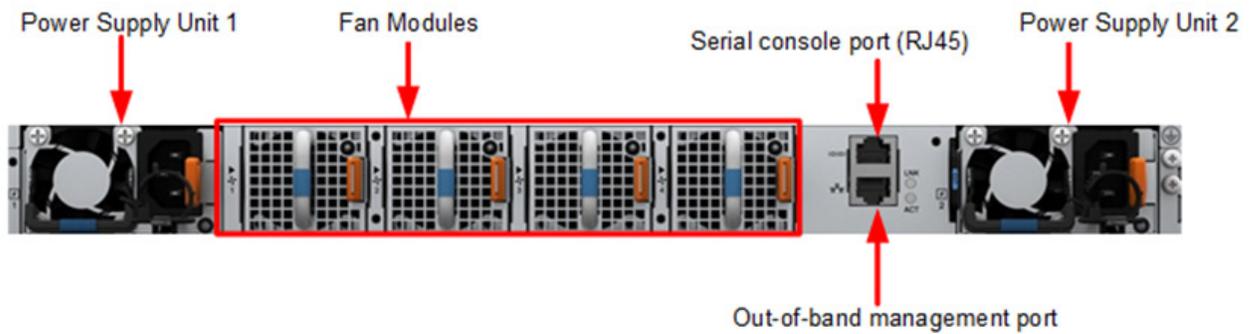


Figure 2 Dell EMC Networking S4148-ON rear view

1.2 Dell EMC PowerEdge R740xd server

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



Figure 3 Dell EMC PowerEdge R740xd server

2 Objective

This deployment guide contains information for configuring physical and virtual networking settings for the deployment of VMware vSAN. Administrators can use this guide to deploy a basic vSAN cluster on an established network.

Note: The example used in this guide is limited to vSAN subject matter and does not include a full data center networking configuration.

Table 1 Deployment and configuration guide definition

Is	Is not
<ul style="list-style-type: none">• Step by step instructions for Layer 2 switch deployment to an established network• Step by step installation for server deployment• Virtual networking configuration• vSAN configuration	<ul style="list-style-type: none">• Best practice recommendations for final configured production state• Performance recommendations for specific workloads• Production user manual• Production configuration guide

3 Leaf-Spine overview

The connections between leaf and spine switches can be layer 2 (switched) or layer 3 (routed). The terms “layer 3 topology” and “layer 2 topology” in this guide refer to these connections. In both topologies, downstream connections to servers, storage and other endpoint devices within the racks are layer 2 and connections to external networks are layer 3.

The following concepts apply to layer 2 and layer 3 leaf-spine topologies:

- Each leaf switch connects to every spine switch in the topology
- Servers, storage arrays, edge routers and similar devices always connect to leaf switches, never to spines

The layer 2 and layer 3 topologies each use two leaf switches at the top of each rack configured as a Virtual Link Trunking (VLT) pair. VLT allows all connections to be active while also providing fault tolerance. As administrators add racks to the data center, two leaf switches configured for VLT are added to each new rack.

The total number of leaf-spine connections is equal to the number of leaf switches multiplied by the number of spine switches. Bandwidth of the fabric may be increased by adding connections between the leaf and spine layer as long as the spine layer has capacity for the additional connections.

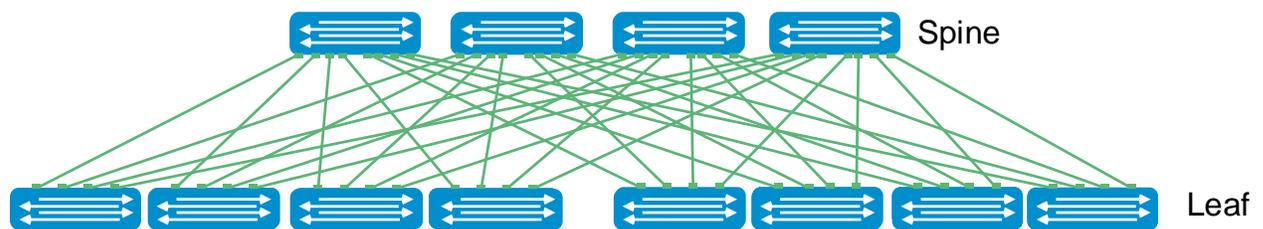


Figure 4 Leaf-Spine architecture

Note: For additional OS10 leaf-spine information, see [Dell EMC Networking L3 Design for Leaf-Spine with OS10EE](#).

3.1 Physical vSAN cluster deployment within SDDC

A vSAN cluster can be deployed in either the layer 2 or layer 3 leaf-spine topology. The deployment example presented in this document can be used in both topologies but focuses only on the relevant layer 2 configuration for vSAN operation. Administrators should reference the leaf-spine documentation linked above for the complete data center configuration.

VMware vSAN network design for consolidated software-defined data center (SDDC) recommends isolating vSAN traffic to its own VLAN. In most data center applications, a physical rack with a leaf pair will contain a single vSAN cluster. Additional vSAN clusters are typically deployed to other racks, but could also exist on the same rack using its own unique VLAN. Shown in [Figure 5](#) is a representation of the four server nodes that comprise the vSAN cluster connected to the top-of-rack (ToR) leaf pair. For simplicity, only the top server node connections are shown.

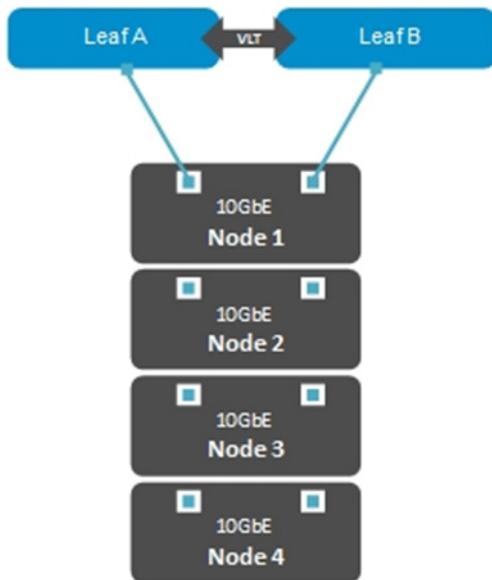


Figure 5 vSAN cluster using four server nodes connected to a leaf pair of switches

3.2 vSAN networking recommendations

This section lists the networking recommendations used in the deployment example for vSAN clusters.

These recommendations can be found in the VMware Validated Design for SDDC documentation at [VMware Validated Design Documentation](#).

The following recommendations are specifically used in the deployment example within this documentation, but do not encompass the entire set of design decisions for the entire SDDC:

- CSDDC-PHY-NET-003: Use two ToR switches for each rack to provide redundancy.
- CSDDC-PHY-NET-004: Use VLANs to segment physical network functions.
- CSDDC-PHY-NET-008: Configure the MTU size to at least 9000 bytes (jumbo frames) on the physical switch ports and distributed switch port groups that support vSAN and vMotion traffic.
- CSDDC-VI-NET-001: Use vSphere Distributed Switch (VDS).
- CSDDC-VI-NET-003: Use the route based on physical NIC load teaming algorithm for all port groups except for ones that carry VXLAN traffic. VTEP kernel ports and VXLAN traffic use route based on SRC-ID.
- CSDDC-VI-Storage-SDS-001: Use only 10 GbE for vSAN traffic.
- CSDDC-VI-Storage-SDS-003: Configure jumbo frames on the VLAN dedicated to vSAN traffic.
- CSDDC-VI-Storage-SDS-004: Use a dedicated VLAN for vSAN traffic for each vSAN enabled cluster.
- [Administering VMware vSAN](#) (VMware vSAN 6.6.1) recommends the following for networking failover and load balancing:
 - When using the teaming algorithm **Route based on physical network adapter load**, the recommended failover configuration is **Active/Active**.

4 Switch configuration

This section provides steps to configure Dell EMC Networking S4148-ON switches running Dell Networking OS 10.4.0E (R3). The process requires basic familiarity with OS10 configuration and network switches.

4.1 Check switch firmware version

Use the following command to verify that the firmware version on the switch is 10.4.0E (R3) or later. If not, visit Dell.com/support to download the latest firmware version for the switch.

```
Switch-1# show version
Dell EMC Networking OS10 Enterprise
Copyright (c) 1999-2018 by Dell Inc. All Rights Reserved.
OS Version: 10.4.0E(R3)
Build Version: 10.4.0E(R3.233)
Build Time: 2018-03-30T18:05:41-0700
System Type: S4148F-ON
Architecture: x86_64
```

4.2 Factory default configuration

Enter the following commands to set the switch to factory defaults.

```
Switch-1# delete startup-configuration
Proceed to delete startup-configuration [confirm yes/no(default)]:yes
Switch-1# reload
System configuration has been modified. Save? [yes/no]:no
Continuing without saving system configuration
Proceed to reboot the system? [confirm yes/no]: yes
```

The switch reboots with factory default settings and is ready to configure utilizing the default username and password of admin/admin.

4.3 Switch-port profiles

On the Dell EMC Networking S4148 switch series, switch-port profiles determine the enabled front-panel ports, supported breakout modes on Ethernet and unified ports. Change the port profile on a switch to customize uplink and unified port operation, and the availability of front-panel data ports.

Note: This document utilizes switch-port profile 1, that supports 100GbE uplinks.

To verify the switches switch-port profile, enter the following command:

```
Switch-1# show switch-port-profile
| Node/Unit | Current | Next-boot | Default |
|-----+-----+-----+-----|
| 1/1 | profile-1 | profile-1 | profile-1 |
```

4.4 Global switch settings

1. Power-on the switch and connect laptop's serial cable to the Dell EMC S4148-ON console port.
2. Configure the hostname and management port IP address:

Switch-1 (S4148-ON)	Switch-2 (S4148-ON)
<pre>hostname Switch-1 interface mgmt 1/1/1 no shutdown no ip address ip address 100.67.171.35/24 ipv6 address autoconfig management route 100.67.0.0/16 100.67.171.254</pre>	<pre>hostname Switch-2 interface mgmt1/1/1 no shutdown no ip address ip address 100.67.171.34/24 ipv6 address autoconfig management route 100.67.0.0/16 100.67.171.254</pre>

The following table contains example VLAN and IP address information. The addresses below will be used throughout the switch configuration steps.

Table 2 VLAN and IP addresses used in example configurations

Purpose	VLAN	Switch-1	Switch-2	VRRP (gateway)
OOB switch management	NA	100.67.171.35 /24	100.67.171.34 /24	NA
ESXi management	2030	172.20.30.251 /24	172.20.30.252 /24	172.20.30.253
vMotion	2031	172.20.31.251 /24	172.20.31.252 /24	172.20.31.253
vSAN	2032	172.20.32.251 /24	172.20.32.252 /24	172.20.32.253

Note: IP addresses are provided in the example configuration commands throughout this document. The addresses are used in a lab setting and are not intended as a recommendation for production use.

4.5 VLT configuration

In this example, configure the VLTi interconnect between the two Dell EMC S4148F-ON switches. The VLTi synchronizes layer 2 table information between the switches and enables them to appear as a single logical unit from outside the VLT domain. Compute nodes can utilize LACP LAG to the pair of VLT switches, for an active/active L2 multipathing scenario. The dedicated vSAN links do not utilize the VLT feature and make use of VMware's load balancing on a distributed port group through route based on physical NIC load.

Use VRRP for gateway redundancy with vSAN and management VLANs. VRRP is an active/standby first hop redundancy protocol (FHRP). When used among VLT peers, it becomes active/active. Both VLT peers have the VRRP virtual MAC address in their FIB table as local destination address. This allows the backup VRRP router to forward intercepted frames whose destination MAC address matches the VRRP virtual MAC address.

1. Configure the VLTi port channel for dual-switch topologies on Switch-1 using the commands in the first column of Table 2 (recommended port values shown).
2. Configure Switch-2 using the commands in column 2 provided in Table 2.

Table 3 VLTi configuration

Switch-1 (S4148-ON)	Switch-2 (S4148-ON)
<pre>interface ethernet1/1/29-1/1/30 description VLTi no shutdown no switchport vlt-domain 127 backup destination 100.67.171.34 discovery-interface ethernet1/1/29- 1/1/30</pre>	<pre>interface ethernet1/1/29-1/1/30 description VLTi no shutdown no switchport vlt-domain 127 backup destination 100.67.171.35 discovery-interface ethernet1/1/29- 1/1/30</pre>

4.6 VLAN configuration

Configure the VLAN interfaces and Virtual Router Redundancy Protocol (VRRP). VRRP will be used as a secondary form of redundancy.

Switch-1 (S4148-ON)	Switch-2 (S4148-ON)
<pre>interface vlan2030 description ESXi management no shutdown mtu 9216 ip address 172.20.30.251/24 vrrp-group 30 virtual-address 172.20.30.253 interface vlan2031 description vMotion no shutdown mtu 9216 ip address 172.20.31.251/24 vrrp-group 31 virtual-address 172.20.31.253 interface vlan2032 description vSAN no shutdown mtu 9216 ip address 172.20.32.251/24 vrrp-group 32 virtual-address 172.20.32.253</pre>	<pre>interface vlan2030 description ESXi management no shutdown mtu 9216 ip address 172.20.30.252/24 vrrp-group 30 virtual-address 172.20.30.253 interface vlan2031 description vMotion no shutdown mtu 9216 ip address 172.20.31.252/24 vrrp-group 31 virtual-address 172.20.31.253 interface vlan2032 description vSAN no shutdown mtu 9216 ip address 172.20.32.252/24 vrrp-group 32 virtual-address 172.20.32.253</pre>

4.7 Node-facing configuration

Configure the vSAN node-facing interfaces with the following steps.

Switch-1 (S4148-ON)	Switch-2 (S4148-ON)
<pre>interface ethernet1/1/43 description "vSAN node 1 Port 1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/45 description "vSAN node 2 Port 1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/47 description "vSAN node 3 Port 1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/49 description "vSAN node 4 Port 1" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge</pre>	<pre>interface ethernet1/1/43 description "vSAN node 1 Port 2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/45 description "vSAN node 2 Port 2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/47 description "vSAN node 3 Port 2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge interface ethernet1/1/49 description "vSAN node 4 Port 2" no shutdown switchport mode trunk switchport access vlan 1 switchport trunk allowed vlan 2030-2032 mtu 9216 spanning-tree port type edge</pre>

4.8 Leaf switch pair uplinks

The purpose of this document is to show the deployment of a vSAN cluster using two VLT peer switches. Administrators can use this leaf pair in a layer 3 or layer 2 leaf-spine network topology. For large deployments that require scale, the layer 3 leaf-spine architecture is commonly used. For detailed information on the layer 3 design for leaf-spine using Dell EMC Networking OS10, see [Dell EMC Networking L3 Design for Leaf-Spine with OS10EE](#).

Features configured on a layer 3 leaf-spine that are not covered in this document:

- Routing protocols – BGP and OSPF
- Equal Cost Multi-Path Routing - ECMP
- Uplink Failure Detection- UFD

Note: The attachments containing running configurations of two VLT peer switches contain a limited set of features used in a lab environment. The running configurations show a layer 3 topology with a single spine switch. See [Dell EMC Networking L3 Design for Leaf-Spine with OS10EE](#) for complete configuration examples of the layer 3 leaf-spine topology.

5 VMware vSAN deployment

This section contains information on deploying a vSAN cluster to the network. The focus is on vSAN configuration as it pertains to the leaf pair switch configuration in section 4. High level information on server preparation, ESXi, and vCenter will be provided with references for additional support.

5.1 VMware ESXi and vCenter Server prerequisites

This section contains the initial setup of the hosts and virtual networking configuration to prepare for vSAN deployment.

5.1.1 Prepare ESXi hosts

This section focuses on the network configuration to include preparing server networking interfaces for migration into a vCenter cluster. The options and choice of server hardware is out of the scope of this document, but are very important to consider when designing a vSAN cluster.

Note: For information on the server hardware requirements, see *Administering VMware vSAN*. VMware vSAN documents can be found at <https://docs.vmware.com/en/VMware-vSAN/index.html>.

This deployment uses two 10GbE ports on a single network interface card (NIC) to connect to the network. ESXi management for this deployment example is in-band, and will use the same 10GbE ports for all data. Follow the steps below to establish ESXi management connectivity:

1. Physically connect the servers to the network using one 10GbE connection to the ports configured in section 4 on each leaf switch.
2. Log into the ESXi console.

Note: The ESXi console can be reached through the iDRAC's virtual console feature.

3. Select **Configure Management Network**.
4. Select **Network Adapters**.
 - a. Choose a single 10GbE vmnic.
 - b. Press **Enter (OK)**.
5. Select **VLAN (optional)**.
 - a. Enter the ESXi management VLAN, for example *2030*.
 - b. Press **Enter (OK)**.
6. Select **IPv4 Configuration**.
 - a. Select **Set static IPv4 address and network configuration**.
 - b. Enter an **IPv4 address**, **Subnet Mask**, and **Default Gateway**. Example: *172.20.30.101*, *255.255.255.0*, *gw 172.20.30.253*.
 - c. Press **Enter (OK)**.
7. Press **Esc** to exit the network configuration menu. Answer **Y** to apply the changes.
8. From the ESXi main menu, select **Test Management Network**. Verify pings are successful.
9. Log out of the ESXi console.

The following table contains example IP address information. The addresses below will be used throughout the virtual networking configuration steps.

Table 4 VLAN and IP addresses used in example configurations

Purpose	VLAN	Host1	Host2	Host3	Host4	Gateway
ESXi management	2030	172.20.30.101 /24	172.20.30.102 /24	172.20.30.103 /24	172.20.30.104 /24	172.20.30.253
vMotion	2031	172.20.31.101 /24	172.20.31.102 /24	172.20.31.103 /24	172.20.31.104 /24	172.20.31.253
vSAN	2032	172.20.32.101 /24	172.20.32.102 /24	172.20.32.103 /24	172.20.32.104 /24	172.20.32.253

Note: IP addresses are provided in the example configuration commands throughout this document. The addresses are used in a lab setting and are not intended as a recommendation for production use.

5.1.2 Create vCenter Server datacenter and cluster

This section provides the steps to create a datacenter object and a cluster. For this example, the vCenter Server is assumed to be present within the network and can be reached by the ESXi management network on the hosts. All appropriate licenses required for host and vSAN operation are installed.

Note: vCenter Server deployment and configuration information can be found at <https://docs.vmware.com/en/VMware-vSAN/index.html>.

A datacenter object is a container for clusters, hosts, networks, datastores, and virtual machines (VMs). An existing datacenter object can be used to contain the new vSAN cluster, or a new datacenter object can be created. The following steps can be used to create a new datacenter object within vCenter:

1. Access the vCenter Server by connecting to the vSphere Web Client through a web browser.
2. Log into vCenter with your credentials.
3. On the web client **Home** screen, select **Hosts and Clusters**.
4. In the **Navigator** pane, right click the vCenter Server object and select **New Datacenter**.
5. Provide a **Name** (example: vSAN) and click **OK**.

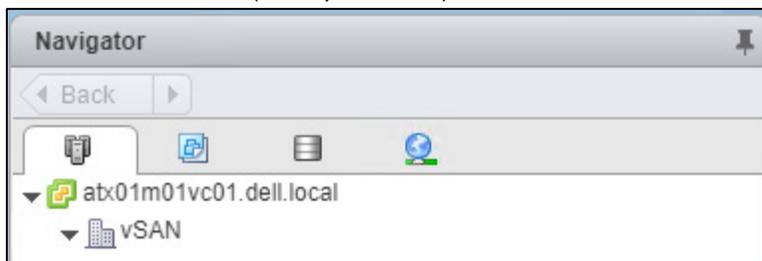


Figure 6 Create datacenter object

A cluster represents a collection of compute and memory resources for a group of physical servers. vSAN aggregates storage devices of the host cluster and creates the storage pool that is shared across all hosts in the vSAN cluster. The following steps are used to create a cluster:

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. Provide a **Name** (example: atx01-w01-vSAN) and click **OK**.

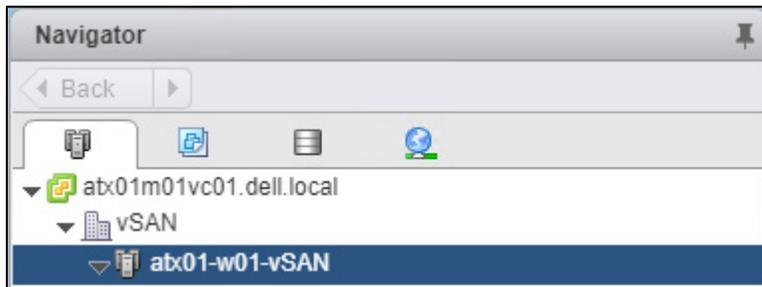


Figure 7 Create cluster

5.1.2.1 Add hosts to the cluster

This section provides steps for adding the hosts to the cluster. The example uses the four hosts that were prepared in section 5.1.1. All hosts use ESXi management VLAN 2030, and have unique IPv4 addresses for management.

Add hosts to the cluster:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. In the **Navigator** pane, right click on the cluster (example: atx01-w01-vSAN) and select **Add Host**.
3. Specify the **IP address** of an ESXi host or the host name if DNS is configured on your network, then 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 then click **Next**.
7. Select a **Lockdown mode**. This guide uses the default setting, **Disabled**. Click **Next**.
8. On the **Ready to complete** screen, select **Finish**.
9. Repeat steps 1 – 8 for each of the remaining hosts in the vSAN cluster.

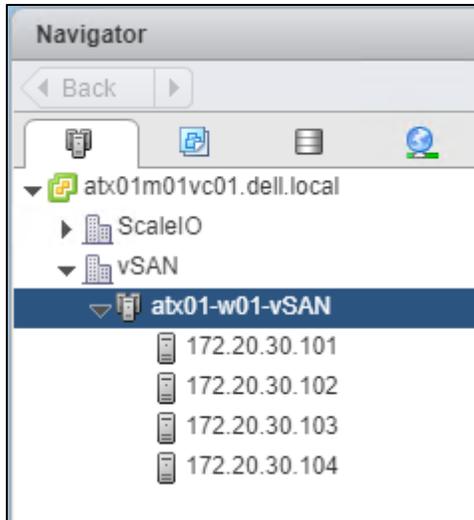


Figure 8 Add hosts to cluster

5.1.3 Create a vSphere Distributed Switch (VDS)

This section provides the steps necessary to configure a vSphere Distributed Switch (VDS). A VDS is a virtual switch that allows each host associated with it to have a consistent network configuration for both the host and virtual machines running on them. Each VDS can be assigned multiple distributed port groups, which includes the VLAN ID for the specified traffic. In this example, port groups for management, vMotion, and vSAN traffic will be defined on a single VDS.

The default configuration for an ESXi host is a single standard vSwitch containing the ESXi management VMkernel adapter. In this example the default condition of the hosts after completing section 5.1.1 is vmnic4 (10GbE NIC port) assigned to vSwitch0 with the management VMkernel adapter (vmk0). The following steps show how to create the VDS, add port groups, and migrate the management VMkernel adapter to the new VDS. The following steps can be used to create the VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on vSAN datacenter. Select **Distributed switch > New Distributed Switch**.
3. Provide a name for the distributed switch (example: atx01-vds01-vSAN). Click **Next**.
4. On the **Select version** page, select **Distributed switch: 6.5.0**. Click **Next**.
5. On the **Edit settings** page:
 - a. Leave the **Number of uplinks** set to **4**. (2 uplinks reserved for future use, can be removed at any time)
 - b. Leave **Network I/O Control** set to **Enabled**.
 - c. Uncheck the **Create a default port group** box.
6. Click **Next** followed by **Finish**.

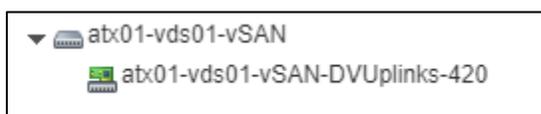


Figure 9 Create distributed switch

Note: The terms vSphere Distributed Switch, VDS, and distributed switch are used interchangeably in this document.

5.1.3.1 Add distributed port groups

This section provides steps to add the management, vMotion, and vSAN distributed port groups. Example values used are listed in Table 5.

Table 5 Example values for distributed port groups

Purpose	Distributed Port Group Name	VLAN
ESXi management	Management-vds01-vSAN	2030
vMotion	vMotion-vds-vSAN	2031
vSAN	vSan-vds-vSAN	2032

The following steps can be used to create the port groups:

1. On the web client **Home** screen, select **Networking**.
2. Right click on the distributed switch. Select **Distributed Port Group > New Distributed Port Group**.
3. On the **Select name and location** page, provide a **Name** for the distributed port group (example: Management-vds01-vSAN). Click **Next**.
4. On the **Configure settings** page, next to **VLAN type**, select **VLAN**. Set the **VLAN ID** (example:2030) for the ESXi management port group. Leave other values at their defaults.
5. Click **Next > Finish**.
6. Create two more Distributed Port Groups (vMotion and vSAN) using the values in Table 5.

After creating the distributed port groups (using example values) your configuration would look like Figure 10.

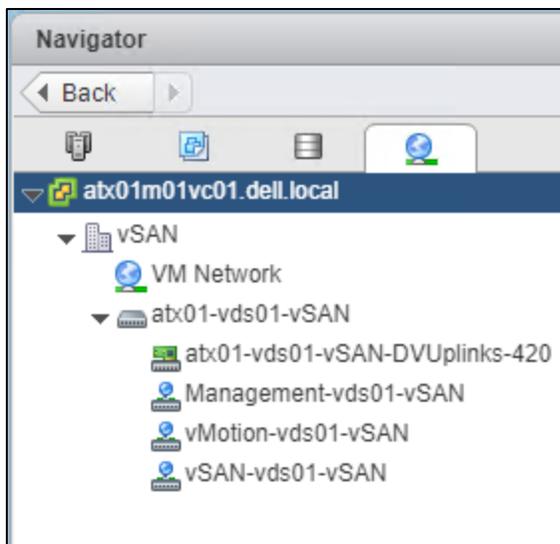


Figure 10 Completed distributed switch with port groups

5.1.3.2 Configure teaming and failover on uplinks

This section provides steps on assigning the behavior of load balancing, network failure detection, switch notification, failback, and uplink failover order. For this example, all port groups will be assigned the same behavior. The recommended setting for load balancing for vSAN, vMotion, and management traffic is route based on physical NIC load. The following steps can be used to set the teaming and failover behavior:

1. On the web client **Home** screen, select **Networking**.
2. Right click on the distributed switch, then select **Distributed Port Group > Manage Distributed Port Groups**.
3. Select only the **Teaming and failover** checkbox, then click **Next**.
4. Click **Select distributed port groups**. Check the top box to select all three port groups.
5. Click **OK**, then click **Next**.
6. On the **Teaming and failover** page:
 - a. For **Load balancing**, select **Route based on physical NIC load**.
 - b. For **Failover order**, click **Uplink 1** and move it up to the **Active uplinks** section by clicking the up arrow. Click **Uplink 2** and move it up to the **Active uplinks** section below **Uplink 1**, by clicking the up arrow. Move **Uplinks 3 & 4** down to the **Unused uplinks** section. Leave other settings at their defaults. An example is shown in Figure 11.
7. Click **Next** then **Finish** to apply the settings.

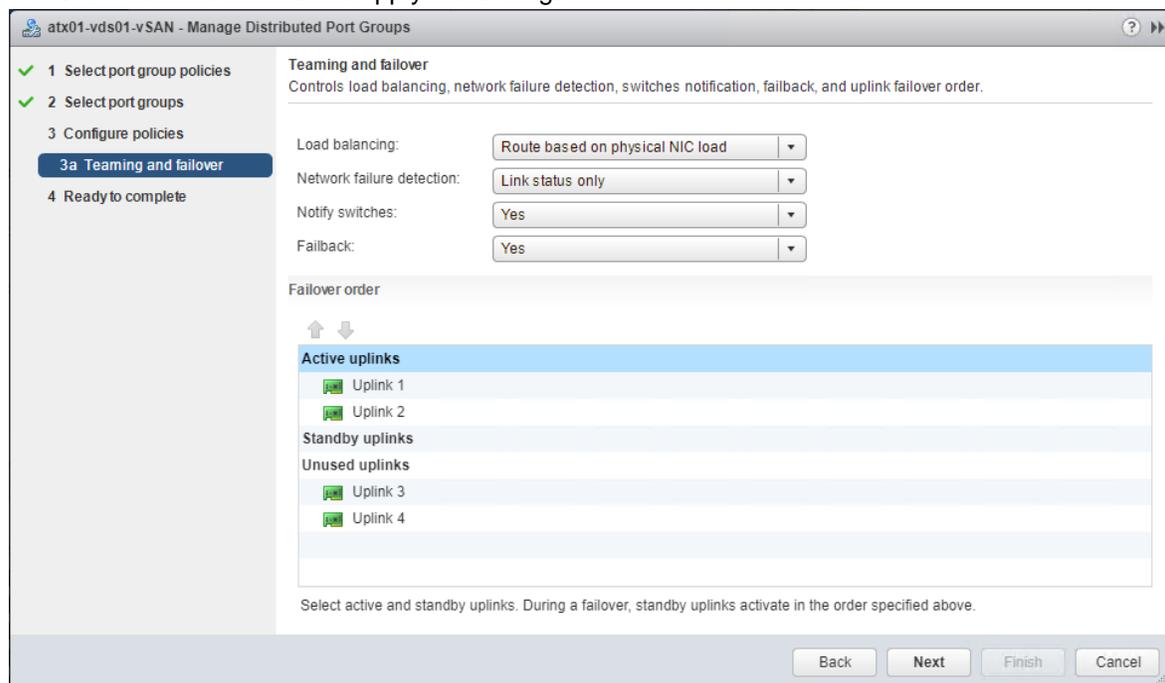


Figure 11 Teaming and failover settings for distributed port groups

5.1.3.3 Add and manage hosts

This section provides steps to add hosts to the VDS. During this process both physical adapters and VMkernel adapters will be managed. Managing physical adapters includes adding the adapters to the VDS and assigning them to uplinks. Managing VMkernel adapters includes adding the adapters to the VDS and configuring adapter settings. When the management VMkernel adapter is assigned to the VDS, it is migrated from the standard vSwitch to the VDS. The following steps can be used to add and manage hosts:

1. On the web client **Home** screen, select **Networking**.
2. Right-click on the distributed switch and select **Add and Manage Hosts**.
3. On the **Select task** page, make sure **Add hosts** is selected, then click **Next**.
4. On the **Select hosts** page, click **New hosts**, then select the check box next to each host in the vSAN cluster.
5. Click **OK**, then click **Next**.
6. On the **Select network adapters tasks** page, be sure the **Manage physical adapters** and **Manage VMkernel adapters** box is checked. Be sure the **Migrate virtual machine networking** box is unchecked.
7. Click **Next**.
8. On the **Manage physical network adapters** page, each host is listed with its vmnics beneath it.
 - a. Select an appropriate vmnic (example: vmnic4) on the first host and click  **Assign uplink**
 - b. Select **Uplink 1** then click **OK**.
 - c. Select the next appropriate vmnic (example: vmnic5) on the first host and click  **Assign uplink**
 - d. Select **Uplink 2** then click **OK**.
 - e. Repeat substeps a through d to configure the remaining hosts, then click **Next**.

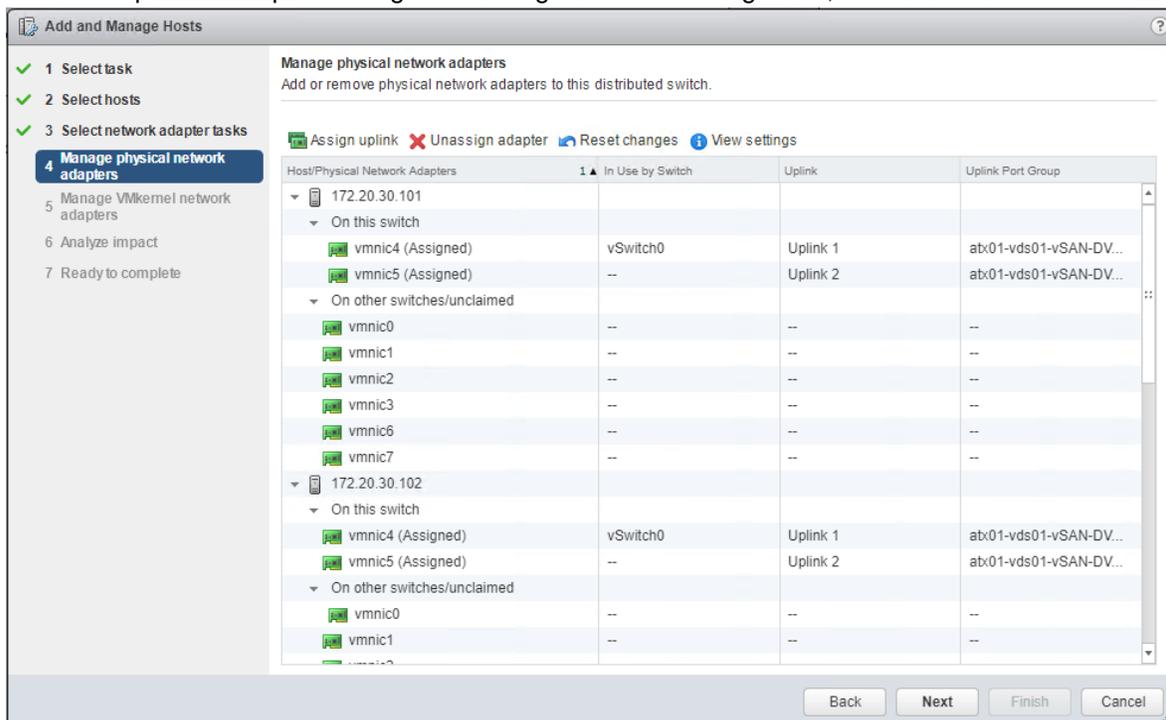


Figure 12 Manage physical network adapters

9. On the **Manage VMkernel network adapters** page, each host is listed with its VMkernel adapters beneath it. Only the default ESXi management VMkernel will be present.
 - a. Select the ESXi management VMkernel adapter, **vmk0**, on the first host and click  **Assign port group**
 - b. Choose the management port group (example: Management-vds01-vSAN). Click **OK**.
 - c. Repeat steps 9.a. - 9.b. for each of the remaining hosts in the vSAN cluster.

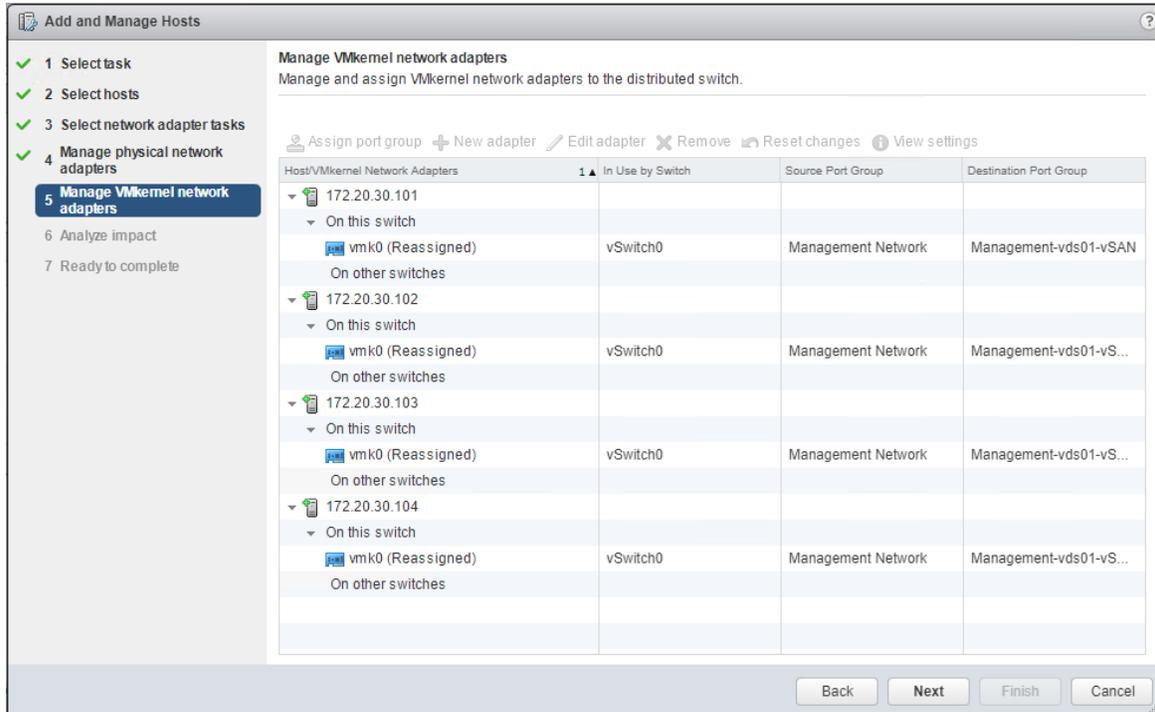


Figure 13 Manage VMkernel network adapters, step 7a-b

- d. Click on the first host, then click on  **New adapter**
 - i. On the **Select target device** page, click on the **Browse** button for **Select an existing network**.
 - ii. Choose the vMotion port group (example: vMotion-vds01-vSAN). Click **OK** > **Next**.
 - iii. On the **Port properties** page, for **Enabled services**, check the **vMotion** checkbox. Leave all other settings as default. Click **Next**.
 - iv. On the **IPv4 settings page**, if DHCP is not used, select **Use static IPv4 settings**. Set the **IP address** (example: 172.20.31.101, 255.255.255.0, gw 172.20.31.253) and **subnet mask** for the host on the vMotion network. Click **Next** > **Finish**.
 - v. Repeat step 9.d for each of the remaining hosts in the vSAN cluster.
- e. Click on the first host, then click on  **New adapter**
 - i. On the **Select target device** page, click on the **Browse** button for **Select an existing network**.
 - ii. Choose the vSAN port group (example: vSAN-vds01-vSAN). Click **OK** > **Next**.
 - iii. On the **Port properties** page, for **Enabled services**, check the **vSAN** checkbox. Leave all other settings as default. Click **Next**.
 - iv. On the **IPv4 settings page**, if DHCP is not used, select **Use static IPv4 settings**. Set the **IP address** (example: 172.20.32.101, 255.255.255.0, gw 172.20.32.253) and **subnet mask** for the host on the vSAN network. Click **Next** > **Finish**.

- v. Repeat step 9.e for each of the remaining hosts in the vSAN cluster.
- f. Click **Next**.
- g. On the **Analyze impact** page, the **Overall impact status** should indicate **No impact**.
- h. Click **Next** then **Finish**.

Note: Deployment examples in this guide only provide instructions for a single leaf switch pair. If deploying onto a routed leaf-spine, administrators may require separate gateways for vMotion or vSAN traffic. A vSAN stretched cluster deployment uses static routes to provide separate gateways. Instructions on how to add static routes can be found within the document [Administering VMware vSAN](#).

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

1. On the web client **Home** screen, select **Networking**.
2. In the **Navigator** pane, select the distributed switch.
3. In the center pane, select **Configure > Settings > Topology** and click the icon next to **VMkernel Ports** to expand. The screen should look similar to [Figure 14](#).
4. Verify each port group is properly configured.

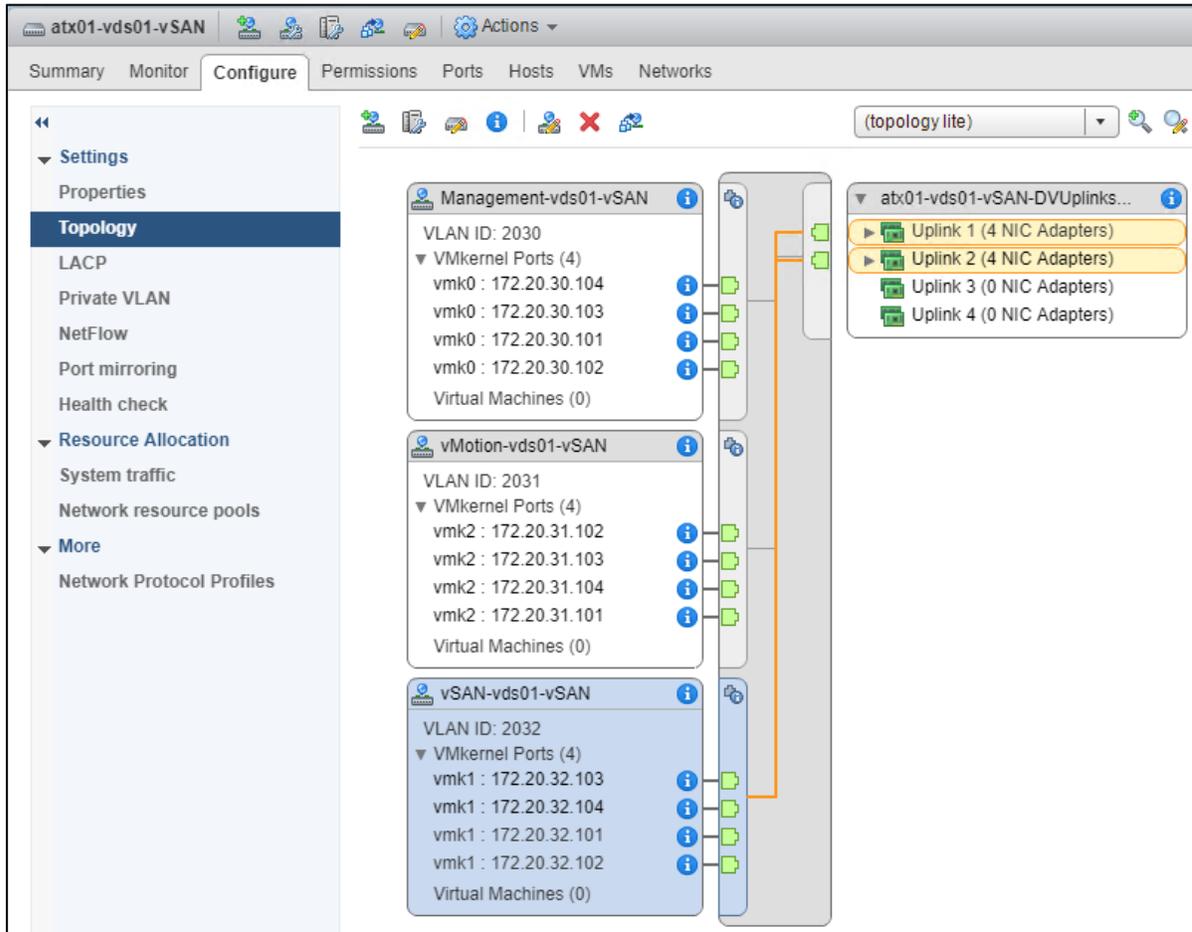


Figure 14 Final topology view for VDS with all port groups, VMkernel adapters, and uplinks configured

5.1.3.4 Enable Jumbo frames

This section provides steps to enable jumbo frames on the VDS and individual VMkernel adapters on each host. Jumbo frames are recommended for vSAN and vMotion traffic and must be configured at the physical switch, the VMkernel adapter, and the VDS. The following steps can be used to enable jumbo frames:

To enable jumbo frames at the VDS:

1. On the web client **Home** screen, select **Networking**.
2. Right click on the distributed switch and select **Settings > Edit Settings**.
3. Click on the **Advanced** settings page.
4. Change the **MTU (Bytes)** to **9000**.
5. Click **OK**.

To enable jumbo frames on the VMkernel adapters:

1. On the web client **Home** screen, select **Hosts and Clusters**.
2. Click on the first host in the vSAN cluster.
3. In the center pane, select **Configure > Networking > VMkernel adapters**.
4. Click on the vSAN VMkernel adapter. Edit the VMkernel adapter by clicking on the **Edit Settings** icon.

5. Click on the **NIC Settings** page.
6. Change the **MTU** to **9000**. Do not modify any other settings.
7. Click **OK**.
8. Repeat steps 4-7 for the vMotion VMkernel adapter on the same host.
9. Repeat steps 3-8 for all the remaining hosts in the vSAN cluster.

5.2 vSAN deployment

VMware vSAN virtualizes the local physical storage resources of ESXi hosts in a single cluster and uses them to create pools of storage that can be divided and assigned to virtual machines and applications. vSAN is implemented directly in the ESXi hypervisor. This section provides steps to enable vSAN on a cluster.

Note: The servers used in this example have 20 SSDs, all at the same capacity. The server hardware configuration depends on the requirements and workload of the cluster. Administrators should consult VMware guidelines and recommendations for vSAN clusters to assist in managing disk groups and devices. For information on designing and sizing a vSAN cluster see [VMware vSAN documentation](#). Within the document [Administering VMware vSAN](#), see the *Designing and Sizing a vSAN Cluster* section.

To configure a vSAN datastore on a cluster:

1. In the web client, go to **Home > Hosts and Clusters**.
2. In the **Navigator** pane, select the vSAN cluster.
3. In the center pane, select **Configure > Settings > vSAN > General**.
4. Click the **Configure** button to launch the **Configure vSAN** wizard.
 - a. On the **vSAN capabilities** page, leave all settings at their defaults. Click **Next**.
 - b. On the **Network validation** page, ensure all hosts show a vSAN enabled network. Click **Next**.

- c. On the **Claim disks** page, set **Group by** to **Host** and expand the hosts to view available disks.
- d. For the first disk on the first host, **Cache tier** is selected. **Capacity tier** is selected for the next seven disks as shown in [Figure 15](#). Leave the remaining disks set to **Do not claim**.

Note: The recommended flash cache to consumed capacity ratio is at least 10 percent. The VMware vSAN maximum limit is 7 capacity devices per diskgroup.

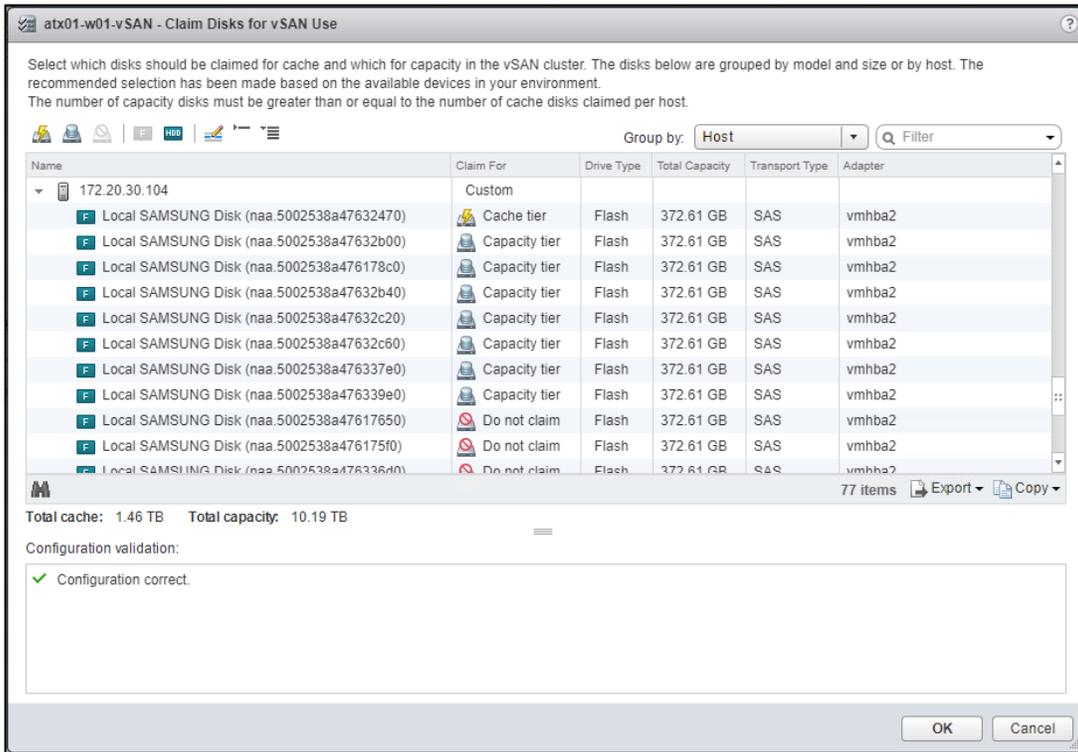


Figure 15 Claim disks - Setting Cache and Capacity tiers

- e. Scroll down to the next host and repeat the process: Set the first disk to **Cache tier**, set the next seven disks to **Capacity tier**, and leave the remaining disks on that host set to **Do Not claim**. Repeat for the remaining hosts listed.
- f. When all groups have been configured, make sure there is a green checkmark in the Configuration validation box at the bottom of the page as shown in [Figure 15](#).
- g. Click **Next**, then click **Finish** to apply the configuration.

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

To configure additional disk groups on the hosts:

1. In the web client, go to **Home > Hosts and Clusters**.
2. In the **Navigator** pane, select the vSAN cluster.
3. In the center pane, select **Configure > Settings > vSAN > Disk Management**. In this example each host has a single disk group using 8 of 20 available disks.

4. Select a host followed by the **Create a new disk group** icon 
5. In the Create Disk Group window:
 - a. Select a single disk for the **cache tier** from the upper list.
 - b. Select seven disks for the **capacity tier** from the lower list. Click **OK**.
6. Repeat steps 4 and 5 as necessary.

The vSAN datastore can be seen by navigating to **Home > Storage**. The vSAN datastore is given a default name of **vsanDatastore** and can be renamed by right clicking on it and choosing **Rename**.

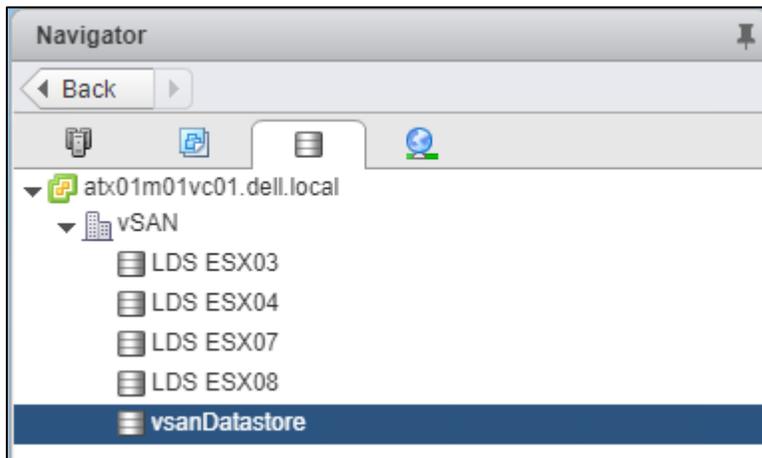


Figure 16 vSAN datastore

At this point the vSAN is deployed and ready for use.

A Validated components

Table 6 Component table example

Component	Description	Version
Server	PowerEdge R740xd	14G Monolithic 0.0.1
	OS	ESXi, 6.5.0 update 02 Build 8294253
	vCenter Server Appliance	6.5.0.1000 Build 5973321
	vSAN	6.6.1
	BIOS	1.2.11
	iDRAC	3.15.15.15
	Processor	2 x Intel® Xeon® Gold 6130 CPU @ 2.10GHz
	Memory	DDR-4 Dual Rank 2666 MHz (total 64 GB)
	Internal Disks	Solid State Disk 372.61 GB (24 drives)
	Network Devices	Intel® 2P X710/2P I350 rNDC
Switch(s)	S4148-ON	10.4.0E(R3)

B Accessing the RS-232 console port

Note: Before starting, verify that the PC has a 9-pin serial port and that a terminal emulation program is already installed and running on the PC.

If your PC does not have a DB-9 serial port connection, use a USB-to-Serial adapter.

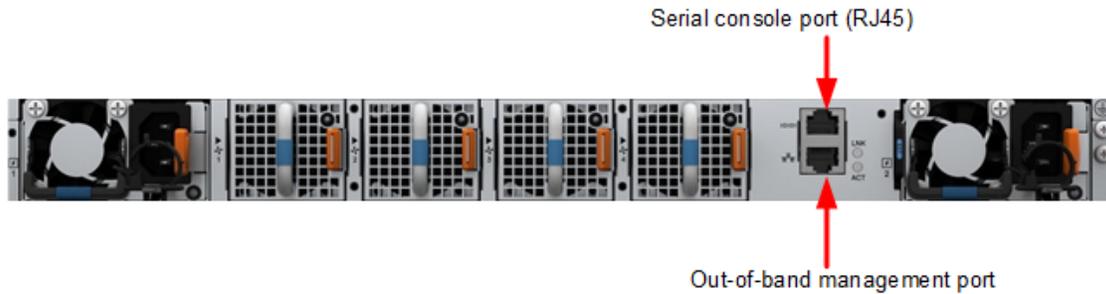


Figure 17 Dell EMC Networking S4148-ON RS-232 console ports

1. Install the RJ-45 connector side of the provided cable into the S4148-ON console port.
2. Install the DB-9 female side of the provided copper cable into your PC's serial port (or into other data terminal equipment [DTE] server hardware that you intend to use).
3. Set the following interface settings:
 - a. 115200 baud rate, where the Micro USB console port is set to 9600 baud rate
 - b. No parity
 - c. 8 data bits
 - d. 1 stop bit
 - e. No flow control

C Product Manuals and technical guides

[Dell.com/support](https://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.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.

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

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

[Manuals and documentation for the Dell Networking S4148F-ON / S4148T-ON / S4148FE-ON](#)

[Manuals and documentation for the Dell Networking S4128F-ON/S4128T-ON](#)

[VMware Validated Design Documentation](#)

[VMware vSAN documentation](#)

[Administering VMware vSAN](#)

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