VMware Virtual SAN on Cisco UCS S3260 Storage Server Deployment Guide

May 2018

Contents

Executive summary

This document describes the VMware Virtual SAN (vSAN) solution on the Cisco UCS© S3260 Storage Server. It provides step-by-step guidance for configuring this Cisco UCS S-Series Storage Server using a VMware vSAN network and storage solution.

Introduction

Virtualization is a compelling technology that helps organizations consolidate and make optimal use of IT resources. It has demonstrated the potential to link IT and business at the core for enhanced agility, quality, scalability, cost effectiveness, and speed. When storage workloads are virtualized, they function differently than storage workloads in the physical world. With virtual workloads, the relationship between applications and storage is N:1, rather than 1:1 as in a physical environment.

Virtual workloads are more mobile than physical workloads, with capabilities such as VMware Storage vMotion providing the flexibility to move workloads within the data center as required. Virtualized applications require more random I/O operations per second (IOPS) than applications hosted in a physical world, which require more sequential IOPS.

In addition, to address the emerging large scale-out designs, the underlying storage solution must be built with specific constructs such as performance or scalability in mind.

These features, plus the emergence of flash-memory storage solutions, including all-flash memory arrays, hybrid flash memory, and server-side flash memory, have created opportunities for new storage architectures, providing customers with a great deal of choice.

VMware vSphere, residing as the first piece of software code between the underlying infrastructure and the virtualized applications, has inherent knowledge of the application requirements from the underlying storage systems and a global view of the infrastructure resources available, allowing it to meet these requirements. This environment provides a unique opportunity for a hypervisor-converged storage solution such as VMware Virtual SAN (vSAN).

Using vSAN as a hypervisor-converged solution on industry-leading data center infrastructure built on the Cisco Unified Computing System™ (Cisco UCS®) and VMware vSphere platforms allows customers to continue to benefit from the extensive history of collaboration and joint innovation in the virtualized data center between Cisco and VMware. This differentiated solution is designed to deliver a scalable virtual and physical architecture, providing superior manageability, security, performance, and cost savings.

This document provides a reference architecture for a VMware Virtual SAN Release 6.6 hybrid solution on Cisco UCS S3260 M4 servers. It includes configuration details for building the joint solution, including steps for configuring manageability and availability capabilities for operating vSAN in a Cisco UCS environment.

VMware Virtual SAN

VMware vSAN is a hypervisor-converged storage solution that is fully integrated with VMware vSphere. vSAN combines storage and computing for virtual machines in a single device, with storage provided within the hypervisor, instead of using a storage virtual machine that runs alongside the other virtual machines. vSAN aggregates locally attached disks in a vSphere cluster to create a storage solution, called a shared data store, that can be rapidly provisioned from VMware vCenter Server during virtual machine provisioning operations (Figure 1).

vSAN is an object-based storage system that is designed to provide virtual machine–centric storage services and capabilities through a storage policy–based management (SPBM) platform. The SPBM platform and virtual machine storage policies are designed to simplify virtual machine storage placement decisions for vSphere administrators. vSAN is fully integrated with core vSphere enterprise features such as vSphere vMotion, vSphere High Availability (HA), and vSphere Distributed Resource Scheduler (DRS). Its goal is to provide both high availability and scale-out storage capabilities. In the context of quality of service (QoS), virtual machine storage policies can be created to define the level of performance and availability required on a per–virtual machine basis.

Many deployment options are available for vSAN. These options range from single, 2-node clusters for small implementations to multiple clusters, each with up to 64 nodes—all centrally managed by vCenter Server. Stretched clusters can easily be configured to enable cross-site protection with no downtime for disaster avoidance and rapid, automated recovery from failure of an entire site.

vSAN 6.6, the sixth generation of vSAN, is designed to help customers modernize their infrastructure by addressing three important IT needs: greater security, lower costs, and faster performance. For example, vSAN 6.6 further lowers total cost of ownership (TCO) by providing more resilient, economical stretched clusters that are easy to deploy and maintain.

Figure 1. VMware vSAN cluster: Hybrid data store

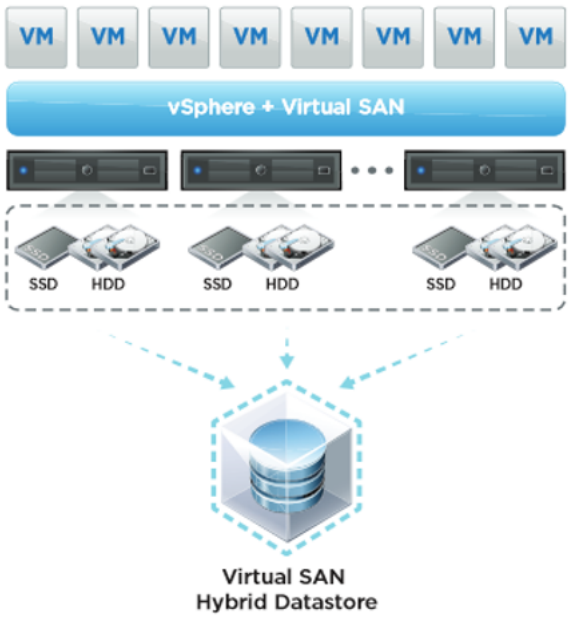

The size and capacity of the vSAN shared data store is dictated by the number of capacity disks per disk group in a vSphere host and by the number of vSphere hosts in the cluster. vSAN is a scale-out solution, in which more capacity and performance can be obtained by adding more disks to a disk group, more disk groups to a host, and more hosts to the cluster.

With vSAN, the SPBM platform plays a major role in the way that administrators can use virtual machine storage policies to specify a set of required storage capabilities for a virtual machine or, more specifically, to specify a set of requirements for the application running in the virtual machine.

The following vSAN data-store capabilities are available in vCenter Server, configurable in virtual machine storage policy:

- Number of failures to tolerate
- Number of disk stripes per object
- Flash read-cache reservation
- Object-space reservation
- Force provisioning

For more information about the capabilities and features of vSAN, see What's New in VMware Virtual SAN.

Cisco Unified Computing System

Cisco UCS is a revolutionary computing architecture designed for IT innovation and business acceleration. It enables fast IT by combining computing, networking, and storage infrastructure with management and virtualization capabilities to offer exceptional speed, simplicity, and scalability. This unique Cisco architecture provides pools of policy-based composable infrastructure that customers can optimize for traditional workloads, data analytics, and cloud-native applications, all within a common operating environment with open APIs for broad interoperability and automation. Cisco UCS has redefined computing to enhance application performance and scalability, simplify infrastructure management, reduce costs, and accelerate IT delivery to the business (Figure 2).

Figure 2. Cisco UCS platform

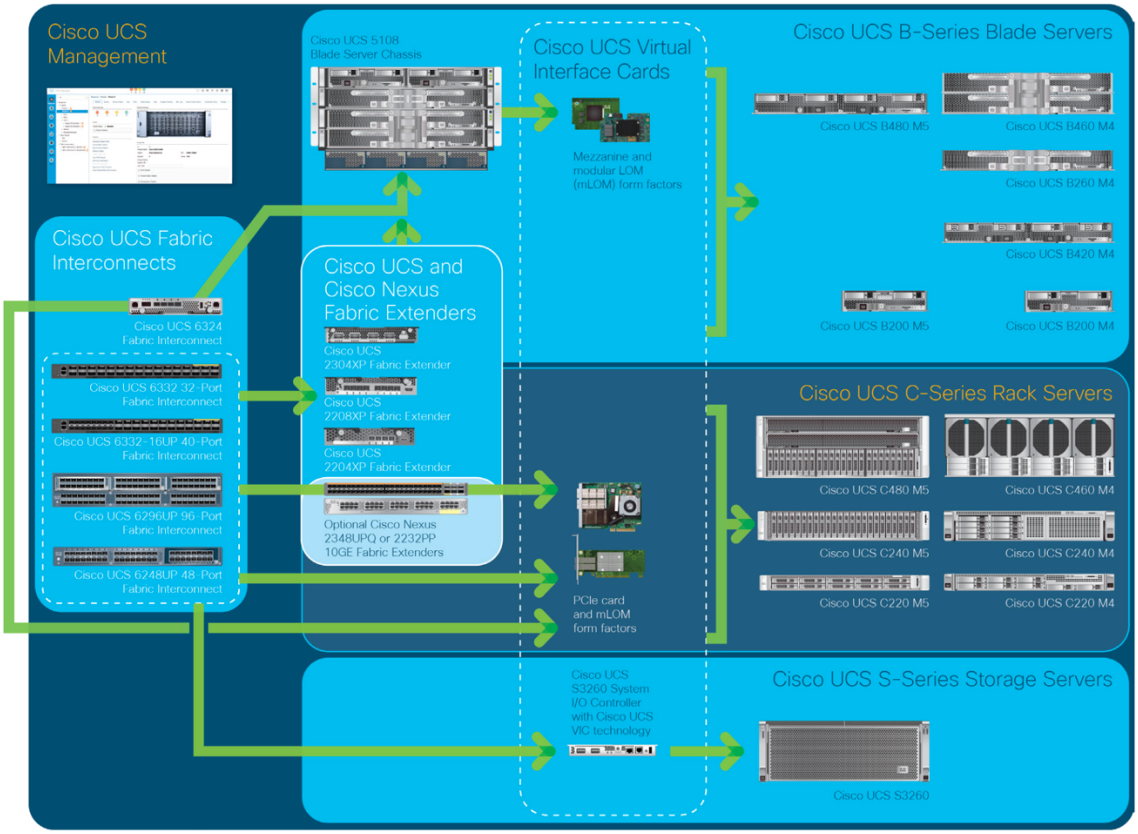

Complex data center infrastructure hampers performance, increases operating costs and risk, and uses resources otherwise available for innovation. Cisco addresses these challenges with a radically simplified, fabric-centric architecture packed with innovation designed to support a wide range of workloads and IT operating models: the Cisco UCS platform.

Cisco UCS helps organizations in many ways:

- **Less infrastructure and more intelligent servers:** This unique architecture enables end-to-end server visibility, management, and control in both bare-metal and virtual environments and facilitates the move to cloud computing and IT-as-a-service (ITaaS) with fabric-based infrastructure.
- **Consolidated resources with Cisco UCS servers:** Cisco UCS servers allow dramatic reduction in the number of devices an organization must purchase, cable, configure, power, cool, and secure. Cisco UCS servers optimize virtualized environments across the entire system. Cisco servers can support traditional operating systems and application stacks in physical environments.
- **Accelerated server deployment:** The smart, programmable infrastructure of Cisco UCS simplifies and accelerates enterprise-class application and service deployment in bare-metal, virtualized, and cloud computing environments. With Cisco UCS unified model-based management, administrators can configure hundreds of servers as quickly as they can configure a single server.
- **Simplified management:** Cisco UCS offers simplified and open management with a large partner ecosystem using Cisco UCS Manager.

The system helps reduce TCO by automating element-management tasks through the use of service profiles that enable just-in-time provisioning. Service profiles increase business agility by quickly aligning computing resources with rapidly changing business and workload requirements.

Cisco UCS S3260 Storage Server

The Cisco UCS S3260 Storage Server (Figure 3) is a modular, high-density, high-availability dual-node rack server well suited for service providers, enterprises, and industry-specific environments. It addresses the need for dense, cost-effective storage for the ever-growing amounts of data in the data center. Designed for a new class of cloud-scale applications, it is simple to deploy and excellent for big data applications, software-defined storage environments such as vSAN and other unstructured data repositories, media streaming, and content distribution. When it is used in combination with Cisco UCS Manager, organizations can easily deploy terabytes (TB) or even petabytes (PB) of storage capacity within minutes.

Figure 3. Cisco UCS S3260 Storage Server

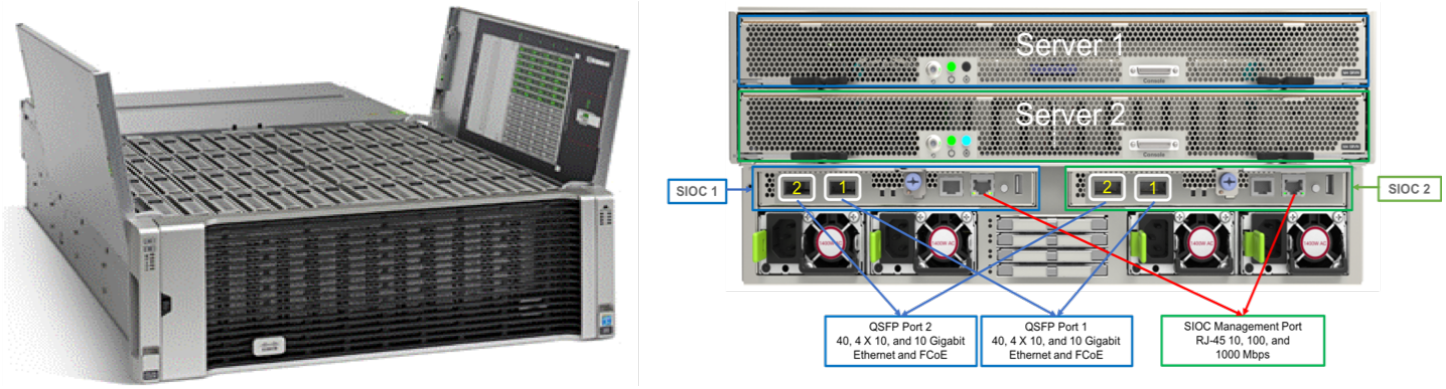

Extending the capabilities of the Cisco UCS C3000 rack server platform, the Cisco UCS S3260 Storage Server helps you achieve the highest levels of data availability. With dual-node capability that is based on the Intel® Xeon® processor E5-2600 v4 series, it offers up to 600 TB of local storage in a compact 4-rack-unit (4RU) form factor. All hard-disk drives (HDDs) can be asymmetrically

split between the dual nodes and are individually hot-swappable. The drives can be built in an enterprise-class Redundant Array of Independent Disks (RAID) redundant design or used in pass-through mode.

This high-density rack server easily fits in a standard 32-inch-depth rack, such as the Cisco® R42610 Rack.

The Cisco UCS S3260 is deployed as a standalone server in both bare-metal or virtualized environments. Its modular architecture reduces TCO by allowing you to upgrade individual components over time and as use cases evolve, without having to replace the entire system.

The Cisco UCS S3260 uses a modular server architecture that, using Cisco's blade technology expertise, allows you to upgrade the computing or network nodes in the system without the need to migrate data from one system to another. It delivers:

- Dual server nodes
- Up to 36 computing cores per server node
- Up to 60 drives, mixing a large form factor (LFF) with up to 28 solid-state disk (SSD) drives plus 2 SSD SATA boot drives per server node
- Up to 512 GB of memory per server node (1 TB total)
- Support for 12-Gbps serial-attached SCSI (SAS) drives
- A system I/O controller (SIOC) with a Cisco UCS Virtual Interface Card (VIC) 1300 platform embedded chip supporting dual-port 40-Gbps connectivity
- High reliability, availability, and serviceability (RAS) features with tool-free server nodes, SIOC, easy-to-use latching lid, and hot-swappable and hot-pluggable components

Hardware and software

This section summarizes the Cisco UCS software and hardware used in the solution described in this document.

Software distributions and versions

Table 1 lists the required software distribution versions.

Table 1. Software versions

Layer	Component	Version
Cisco UCS S3260 (chassis)	Chassis management controller	Release 3.0(3a)
	Shared adapter	Release 4.1(2d)
Server nodes: Cisco UCS C3X60 M4	BIOS	Release C3X60M4.3.0.3b.0.0325171543
	Cisco Integrated Management Controller (IMC)	Release 3.0(3a)
Network: Cisco Nexus® 9396PQ Switch	BIOS	Release 07.17
	Cisco NX-OS Software	Release 6.1(2)I3(3a)
VMware software	VMware vCenter	Release 6.5.0 Build 5973321
	VMware ESXi	Release 6.5.0, 5969303
	VMware vSAN	Release 6.6

Note: The versions listed in Table 1 reflect the vSAN ReadyNode certification requirements.

Hardware requirements

Table 2 summarizes the hardware components used in deploying vSAN 6.6 on the Cisco UCS S3260 server.

Table 2. Hardware requirements

Component	Model	Quantity	Comments

Node	Cisco UCS S3260 M4 chassis	2	• 2 x Cisco UCS S3260 M4 server nodes per Chassis (total = 4 nodes)
			• Per server node
			  • 2 x Intel Xeon processor E5-2680 v4
			  • 256 GB of RAM (8 x 32 GB at 2400 MHz)
			  • Cisco 12-Gbps SAS RAID controller
			  • 1 x 800-GB Non-Volatile Memory Express (NVMe) for VMware ESXi OS, 24 x 5.6-TB HDDs for capacity disks, and 4 x 1.6-TB SSDs for cache disks
			  • Dual-port 40-Gbps VIC

Switch	Cisco Nexus 9396PX Switch	2	–

Physical layout

Figure 4 shows the physical layout of the components used in deploying vSAN on the Cisco UCS S3260 M4 server.

Figure 4. Physical topology

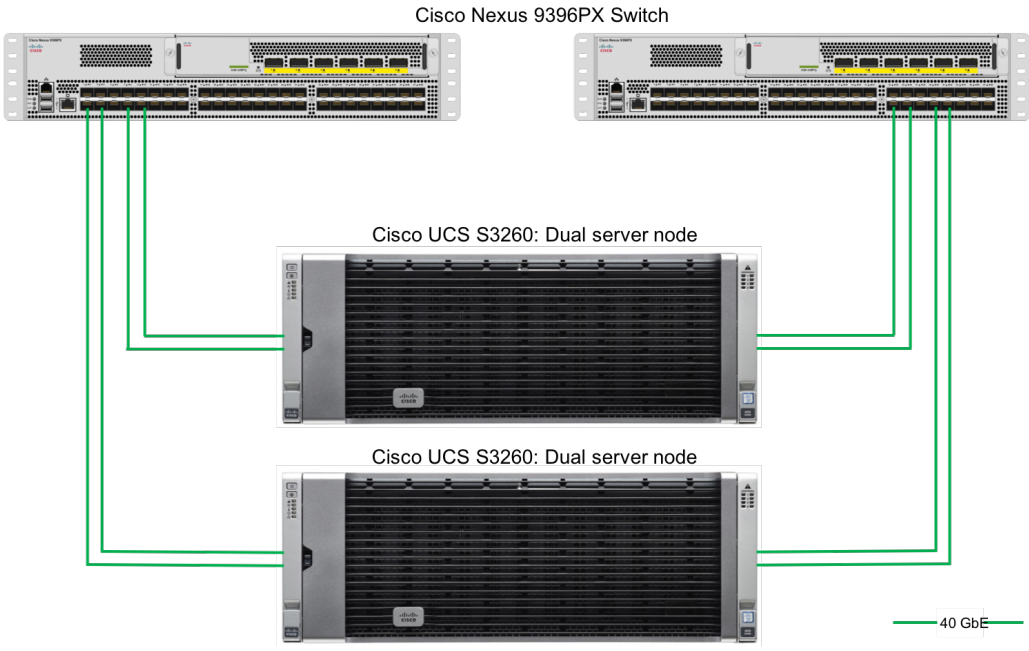

The connectivity of the solution is based on 40 Gbps. Each Cisco UCS S3260 M4 server is connected with 2 x 40-Gbps cable to each Cisco Nexus 9396PX Switch.

Table 3 shows the cabling for the Cisco UCS S3260 server and the Cisco Nexus 9396PX Switches.

Table 3. Cisco UCS S3260 Storage Server and Cisco Nexus 9396PX Switch cabling

Device	Port	Connection	Remote device	Remote port
S3260 chassis 1	SIOC 1/1	40 Gigabit Ethernet	9396PX - A	Eth 1/1
S3260 chassis 1	SIOC 1/2	40 Gigabit Ethernet	9396PX - B	Eth 1/1
S3260 chassis 1	SIOC 2/1	40 Gigabit Ethernet	9396PX - A	Eth 1/2
S3260 chassis 1	SIOC 2/2	40 Gigabit Ethernet	9396PX - B	Eth 1/2
S3260 chassis 2	SIOC 1/1	40 Gigabit Ethernet	9396PX - A	Eth 1/3
S3260 chassis 2	SIOC 1/2	40 Gigabit Ethernet	9396PX - B	Eth 1/3
S3260 chassis 2	SIOC 2/1	40 Gigabit Ethernet	9396PX - A	Eth 1/4
S3260 chassis 2	SIOC 2/2	40 Gigabit Ethernet	9396PX - B	Eth 1/4

VLANs

You need to use VLANs to separate the vSAN network traffic, management traffic, vMotion traffic, and virtual machine traffic. Table 4 lists the VLANs configured on the Cisco Nexus 9396PX Switch and the vSphere standard switch. For the Cisco Nexus 9396PX Switch configuration, visit the following link:

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/6-x/layer2/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Layer_2_Switching_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Layer_2_Switching_Configuration_Guide_chapter_011.h

Table 4. Required VLANs

VLAN name	VLAN purpose	ID used for validation in this document
Mgmt	VLAN for management interfaces	20
vSAN	VLAN for vSAN traffic	100
vMotion	VLAN for vMotion traffic	200
VM	VLAN for production virtual machine interfaces	10

Note: Set QoS policy as defined by your local network administrator. See the Cisco Nexus 9000 Series NX-OS Quality of Service Configuration Guide for more information about configuring QoS on Cisco Nexus 9000 Series Switches.

Note: In vSAN 6.6 and later releases, multicast is not required on the physical switches that support the vSAN cluster. If some hosts in your vSAN cluster are running earlier versions of the software, a multicast network still is required.

Cisco UCS S3260 configuration

This section provides a guide to help you configure the Cisco UCS S3260.

Create system IP addresses

A Cisco UCS S3260 system can have up to five IP addresses:

Note: All controllers present in the system must have IP addresses assigned so that they can communicate with each other. All IP addresses can be assigned by your Dynamic Host Configuration Protocol (DHCP) server, or you can assign static IP addresses.

- **Management IP address:** This is the overall system virtual IP address. You log in to this address when you access the system's IMC interface through your LAN connection to the active chassis management controller in SIOC 1 or SIOC 2 (see Overview of Cisco UCS S3260 Architecture).
- **SIOC 1 CMC IP address:** This is the internal address for the chassis management controller (CMC) in SIOC 1. This address can be assigned by your DHCP server, or you can set a static address by using the IMC interface.
- **SIOC 2 CMC IP address:** This is the internal address for the CMC in SIOC 2 (if installed). This address can be assigned by your DHCP server, or you can set a static address by using the IMC interface.
- **Server 1 BMC IP address:** This is the internal address for the board management controller (BMC) in server node 1. This address can be assigned by your DHCP server, or you can set a static address by using the IMC interface.
- **Server 2 BMC IP address:** This is the internal address for the BMC in server node 2 (if installed). This address can be assigned by your DHCP server, or you can set a static address by using the IMC interface.

Note: The management IP address, SIOC CMC IP addresses, and server BMC IP addresses must all be configured on the same subnet.

Create virtual network interface cards

You need to create eight virtual network interface cards (vNICs) for each server node: four on uplink port 1 and four on uplink port 2, with the properties shown in Table 5.

Table 5. vNIC properties

vNIC	Consistent device name (CDN)	Maximum transmission unit (MTU)	Uplink port	VLAN	VLAN mode
eth0	Mgmt-1	1500	0	None	Trunk
eth1	Mgmt-2	1500	1	None	Trunk
eth2	vSAN-1	9000	0	None	Trunk
eth3	vSAN-2	9000	1	None	Trunk
eth4	vMotion-1	9000	0	None	Trunk
eth5	vMotion-2	9000	1	None	Trunk
eth6	VM-1	9000	0	None	Trunk
eth7	VM-2	9000	1	None	Trunk

Follow these steps to create the vNIC on the Cisco UCS S3260 server.

1. On the S3260 home screen, click the top-right button:

2. In the menu that appears, expand the Networking section and choose Adapter Card SIOC1 and select the vNICs tab.
3. Under Host Ethernet Interfaces, click Add vNIC.
4. Enter the name, MTU value, uplink port, and VLAN mode as shown in Table 5.
5. Expand the following sections and make the following changes for vNIC eth2-eth7. The management vNIC, eth0 and eth1 will use the default values:
   - · Ethernet Interrupt
     - Interrupt Count: 32

- Ethernet Receive Queue
  - Receive Queue Count: 8
  - Receive Queue Ring Size: 4096
- Ethernet Transmit Queue
  - Transmit Queue Count: 8
  - Receive Queue Ring Size: 4096
- Completion Queue
  - Completion Queue Count: 16
- Receive Side Scaling
  - Enable TCP Receive Side Scaling: Selected

6. Verify that your settings match the settings shown here.

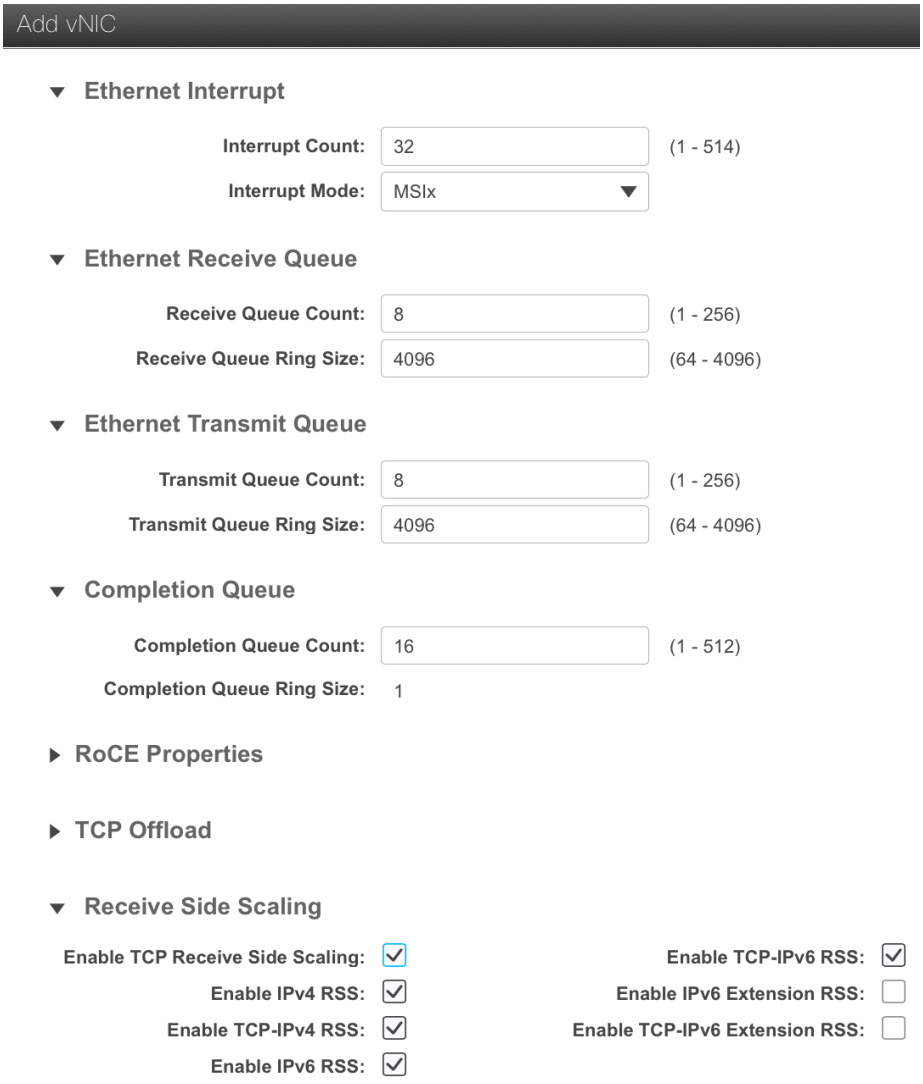

7. Click Add vNIC to create the vNIC.
8. Create all the vNICs as shown in Table 5.

Configure disk zoning

You can assign disk drives to the server nodes using disk zoning. Disk zoning can be performed on the controllers in the same server or on the controllers on different servers. The RAID controller will see only the physical drives that are zoned for it in the disk zoning.

Follow these steps to configure disk zoning.

1. From the IMC home screen, choose Inventory > Dynamic Storage.
2. Select the Zoning tab.
3. Select the drives PD1 to PD24 and PD49 to PD52 and click Assign to Server 1. Then click Save Change.

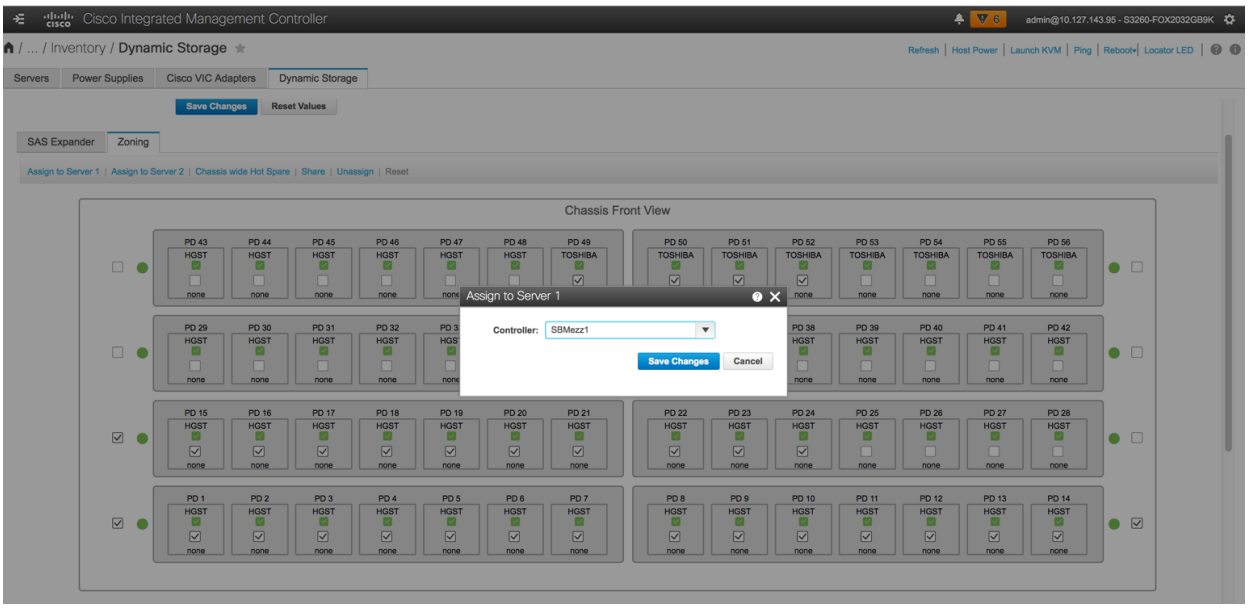

4. Select the drives PD25 to PD48 and PD53 to PD56 and click Assign to Server 2. Then click Save Changes.

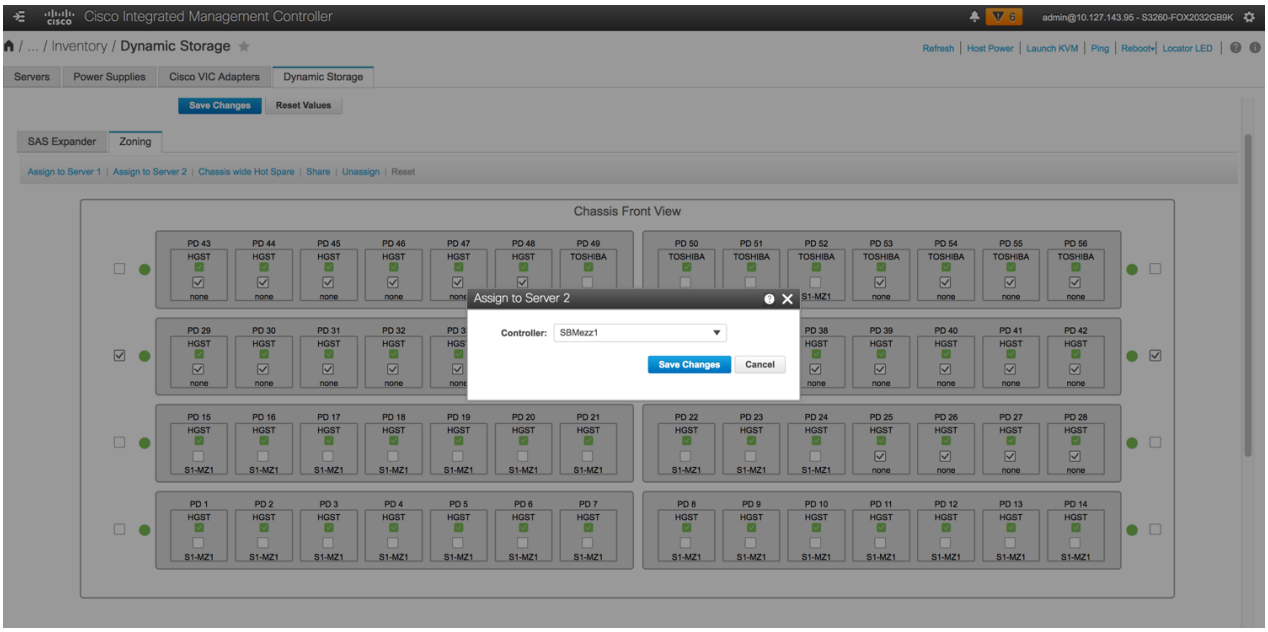

Configure BIOS policy

Configure the BIOS policy for the vSAN environment as shown here. This policy is configured to achieve high performance using virtualization best practices (refer to M4 BIOS paper).

1. From the S3260 IMC home screen, select the Navigation icon and choose Compute.
2. Select the server and click choose BIOS > Advanced.
3. Change the BIOS settings as shown here and save the changes.

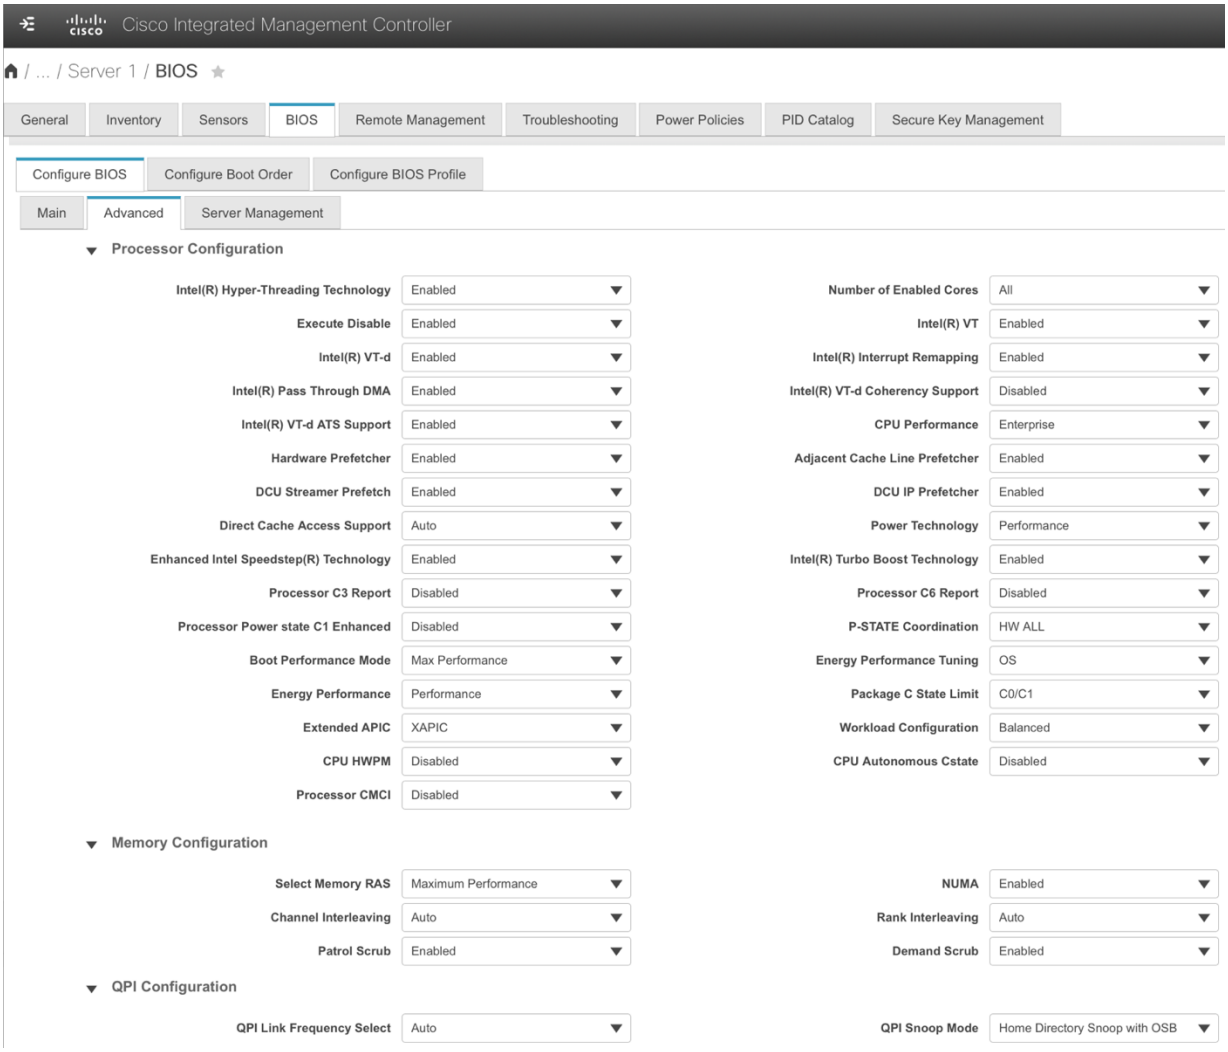

Configure boot policy

To create the boot policy, follow these steps.

1. From the S3260 IMC home screen, select the Navigation icon and choose Compute.
2. Select the server and choose BIOS > Configure Boot Order.
3. For Configured Boot Mode, choose Uefi. Then click Save Changes.
4. Click Configure Boot Order, select EFI and CDROM, click Right arrow button to configure the Boot order
5. Save the changes.

For more information, see "Setting Up Booting in UEFI Mode in the Cisco IMC GUI"

VMware vSphere 6.5U1 configuration

This section provides guidance for configuring vSphere.

Install VMware ESXi 6.5U1

Cisco recommends the use of the Cisco customized ESXi ISO image for installing the ESXi hypervisor, available at the following link: VMware vSphere Hypervisor (ESXi) 6.5U1.

Follow these steps to install the ESXi hypervisor on Cisco UCS S3260 nodes.

1. Log in to the Cisco UCS S3260 IMC.
2. Click Launch KVM, select the server node, and click Launch.
3. In the KVM window, click Virtual Media. Then click Activate Virtual Devices.
4. If you are prompted to accept an unencrypted KVM session, accept the session.
5. Click Virtual Media and select Map CD/DVD.
6. Browse to the ESXi installer ISO image file and click Open. Then click Map Device.
7. Click the KVM tab to monitor the server boot.
8. Boot the server by choosing the Power > Power on System option and clicking OK.
9. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that appears.
10. After the installer is finished loading, press Enter to continue the installation process.
11. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
12. Select the NVMe disk for installing ESXi and press Enter to continue with the installation.

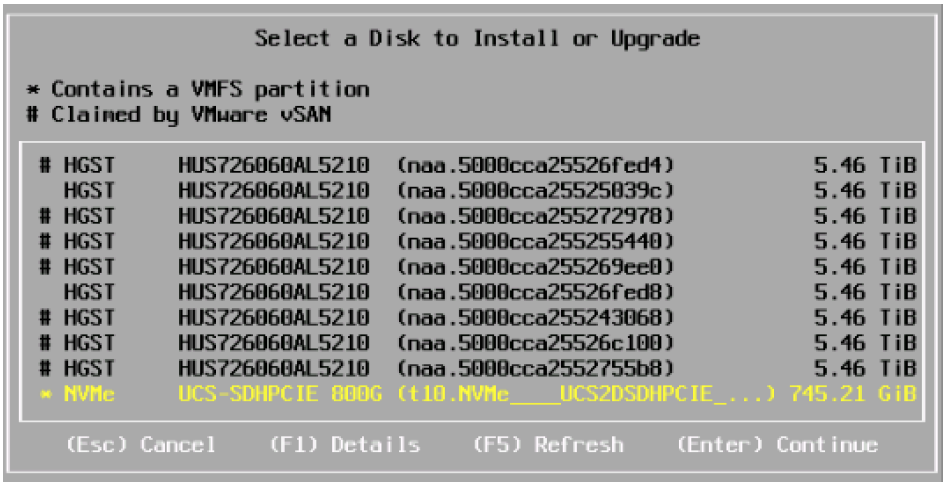

13. Select the appropriate keyboard layout and press Enter.
14. Enter and confirm the root password and press Enter.
15. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue the installation process.
16. After the installation is complete, press Enter to reboot the server.
17. Repeat the preceding steps on all the server nodes to install ESXi

Note: All drivers are Inbox only.

Deploy VMware vCenter Server

VMware vCenter Server provides a centralized platform for managing your VMware vSphere environments, allowing you to automate and deliver a virtual infrastructure with confidence.

You can deploy the vCenter Server appliance with an embedded or external platform services controller to manage your vSphere environment. For more information about installing the vCenter Server appliance, see Deploying the vCenter Server Appliance.

Configuring vCenter Server

Follow these steps to set up vCenter Server.

1. Using a web browser, navigate to https://<vcenter-ip>/vsphere-client.
2. Click Download Enhanced Authentication Plugin. Install the plug-in by double-clicking the downloaded file.
3. Log in using the single sign-on (SSO) username and password created during the vCenter installation.

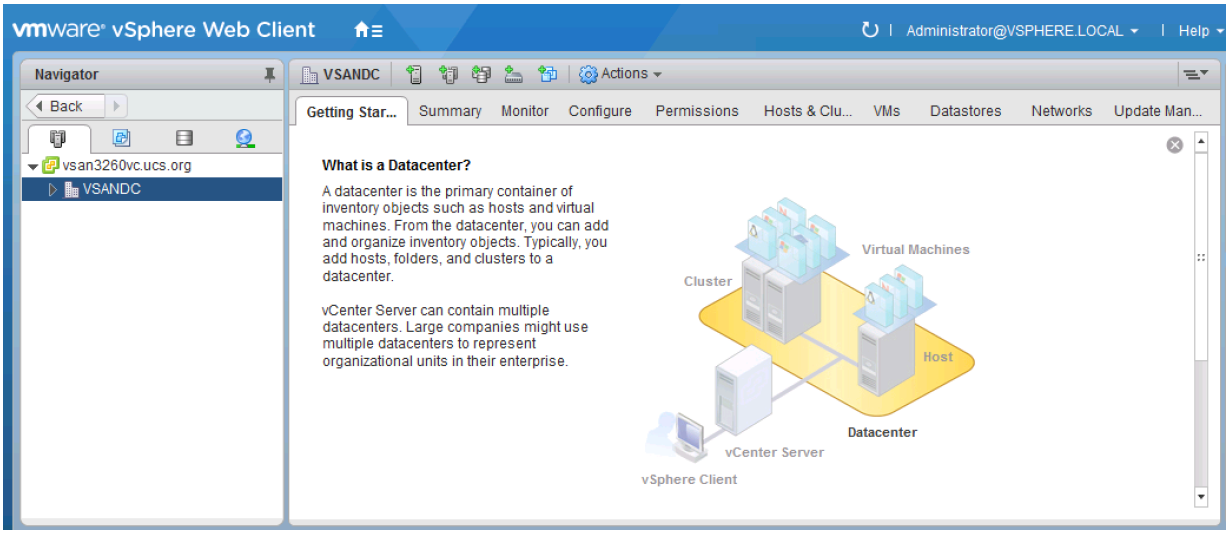

4. Click Create Datacenter in the center pane.
5. Enter **VSAN-DC** as the data center name.
6. Click OK.
7. Right-click the VSAN-DC data center in the list in the center pane. Choose New Cluster.
8. Enter the cluster name.
9. Check the box to turn on DRS. Leave the default values.
10. Check the box to turn on vSphere HA. Leave the default values.
11. Click OK to create the new cluster.
12. On the left, right-click Cluster and choose Add Host.
13. In the Host field, enter either the IP address or the fully qualified domain name (FQDN) of one of the ESXi hosts. Click Next.
14. Type **root** as the user name and the **root** password. Click Next to continue.
15. Click Yes to accept the certificate.
16. Review the host details and click Next to continue.
17. Assign a license or leave the evaluation mode set. Then click Next to continue.
18. Click Next to continue.

19. Click Next to continue.
20. Review the configuration parameters. Then click Finish to add the host.
21. Repeat the preceding steps to add all the hosts to the cluster.

VMware vSAN network design

vSAN supports the use of the VMware vSphere standard switch or vSphere distributed switch. This document covers the use of the vSphere standard switch.

VMware recommends isolating vSAN traffic by dedicating a separate VLAN to this traffic. When a design uses multiple vSAN clusters, each cluster should use a dedicated VLAN or segment for its traffic. This approach prevents interference between clusters and helps with troubleshooting cluster configuration. For more information on see Networking Requirements for vSAN.

Figure 5 illustrates the logical design of the network.

Figure 5. VMware vSAN logical design

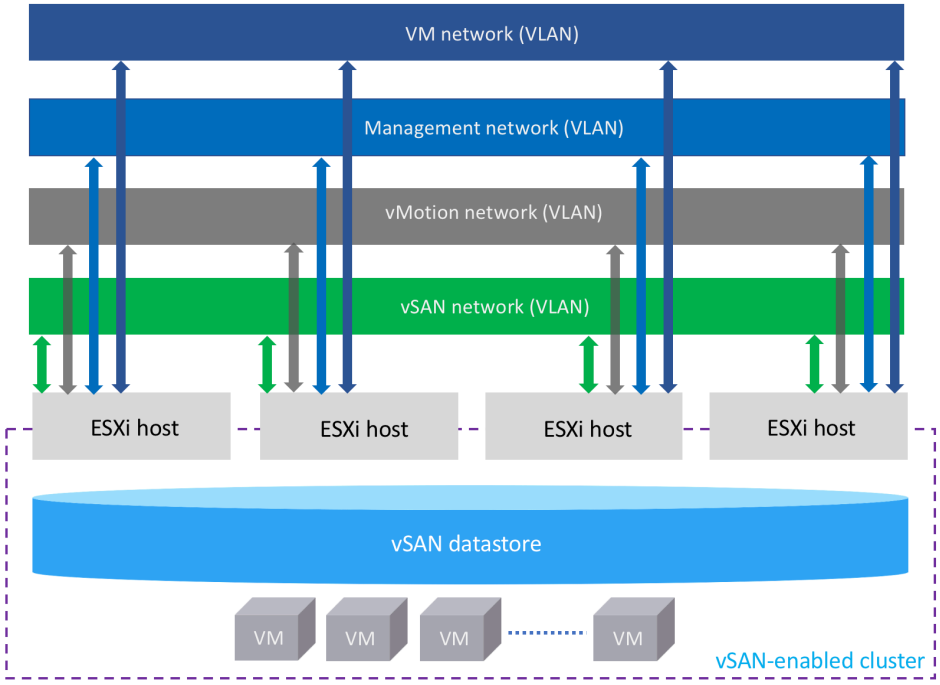

Set up VMkernel ports and virtual switches

Each host is configured with four virtual switches (vSwitches). Tables 6 and 7 show the vSwitch configuration.

Table 6. vSwitch configuration

Switch	VLAN and traffic type	Virtual machine NIC (vmnic)	MTU
vSwitch0	VLAN 20: Management traffic	vmnic0 and vmnic1	1500
vSwitch1	VLAN 100: vSAN traffic	vmnic2 and vmnic3	9000
vSwitch2	VLAN 200: vMotion traffic	vmnic4 and vmnic5	9000
vSwitch3	VLAN 10: Virtual machine data traffic	vmnic6 and vmnic7	9000

For information about creating vSwitch and vmnics, see the section "Create a vSphere Standard Switch"

Table 7. VMkernel IP address details

VMware ESXi host	VMkernel IP address details	
	VMware vSAN	VMware vMotion
ESXi01	192.168.100.75/24	192.168.200.175/24
ESXi02	192.168.100.176/24	192.168.200.176/24
ESXi03	192.168.100.177/24	192.168.200.177/24
ESXi04	192.168.100.178/24	192.168.200.178/24

Configure the vmnics as shown here.

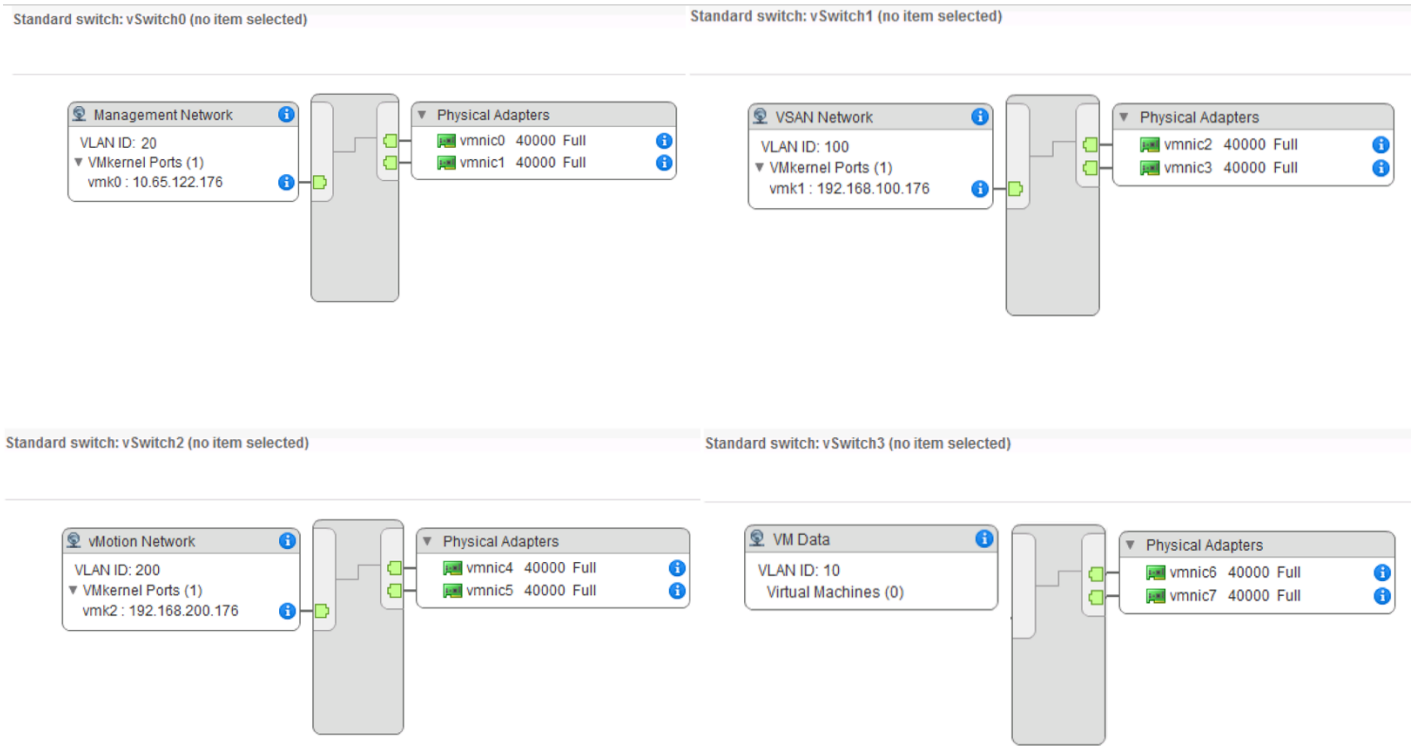

Configure VMware vSAN disk groups

Each Cisco UCS S3260 server node is configured with one 800-GB NVMe disk for the ESXi OS, four 1.6-TB SSDs for the cache tier, and twenty-four 5.6-TB HDDs for the capacity tier. Four disk groups (4 x (1 x 1.6-TB SSD + 6 x 5.6-TB HDDs) will be created on each S3260 server node and contribute to the vSAN data store. For the purposes of this deployment guide, the disk group and drive choices listed here are used; however, vSAN architecture permits other configurations. Please consult the VMware documentation for additional information about disk groups.

Enable VMware vSAN

vSAN is an ESXi cluster-level feature that is configured using the vSphere Web Client.

vSAN 6.6 and later releases have a uniform workflow for claiming disks across all scenarios. All available disks are grouped by model and size or by host. You must select which devices to use for the cache and which to use for capacity.

Follow these steps to enable and configure the vSAN on the cluster.

1. Navigate to an existing host cluster in the vSphere Web Client.
2. Click the Configure tab.
3. Under vSAN, select General and click the Configure button.
4. Select the vSAN capabilities:
   - Deduplication and Compression: vSAN can perform block-level deduplication and compression to save storage space. When you enable deduplication and compression on a vSAN all-flash cluster, redundant data within each disk group is reduced.

Note: Deduplication and compression are available only on all-flash disk groups.

   - Encryption: vSAN encryption is the industry's first native hyperconverged infrastructure (HCI) encryption solution. It is built into the vSAN software. Select the check box if you want to enable data-at-rest encryption, and select a key management server (KMS).

Note: To use vSAN encryption, just as for virtual machine encryption, a KMS is required.

5. Under Fault Domains and Stretched Cluster, select "Configure fault domains." Then click Next.

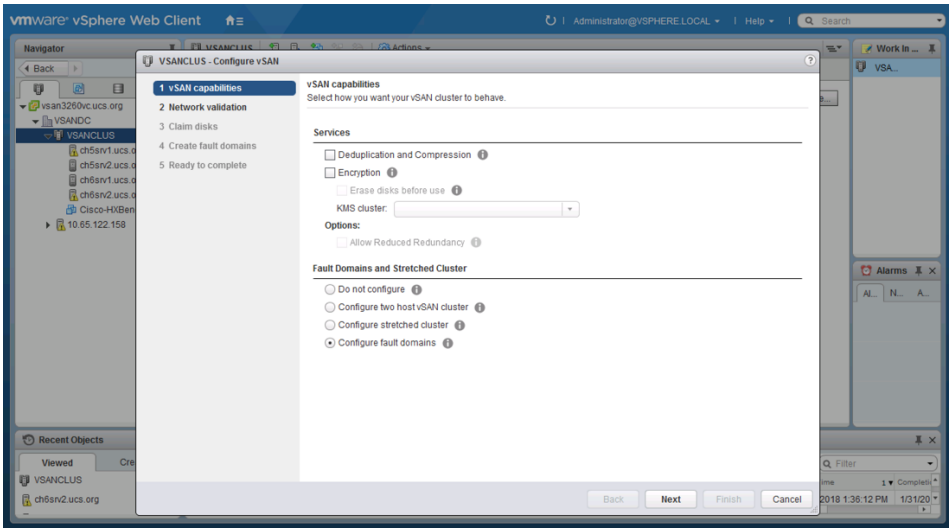

6. On the “Network validation” page, check the settings for vSAN VMkernel adapters. Then click Next.

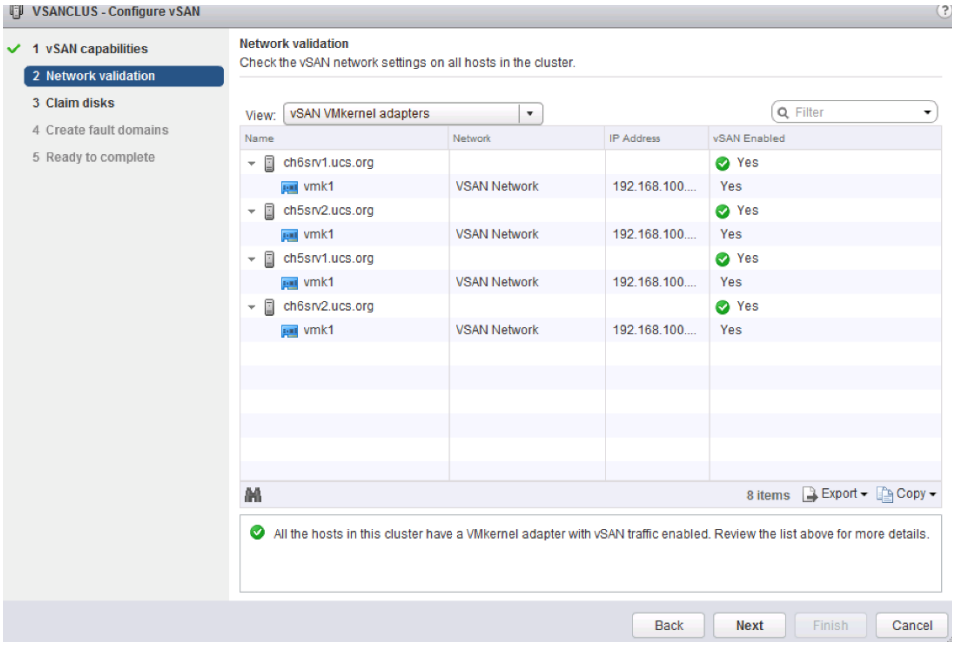

7. On the “Claim disks” page, select the cache and capacity disks for use by the cluster. Then click Next.

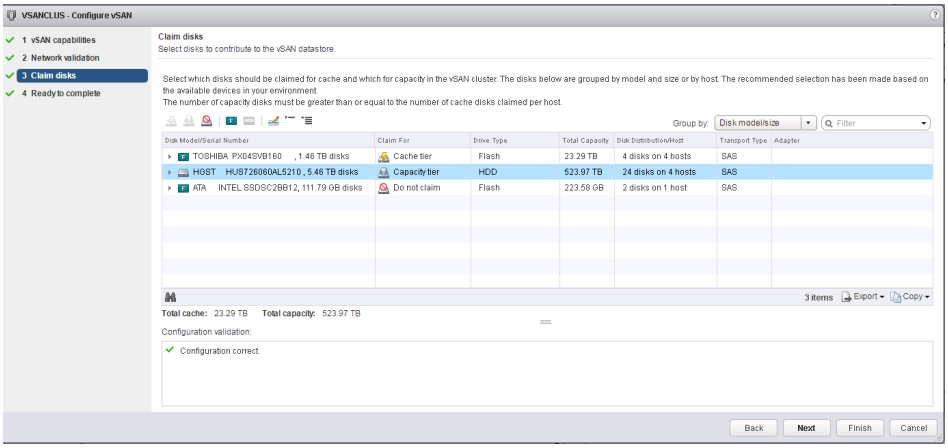

8. On the “Create fault domains” page, define fault domains for the cluster. Then click Next.

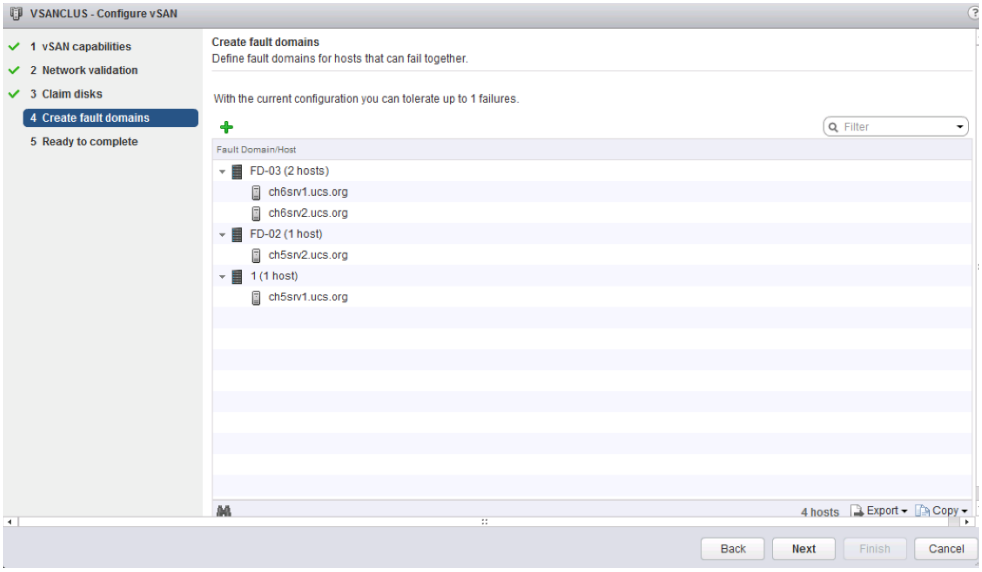

9. On the “Ready to complete” page, review the configuration. Then click Finish.

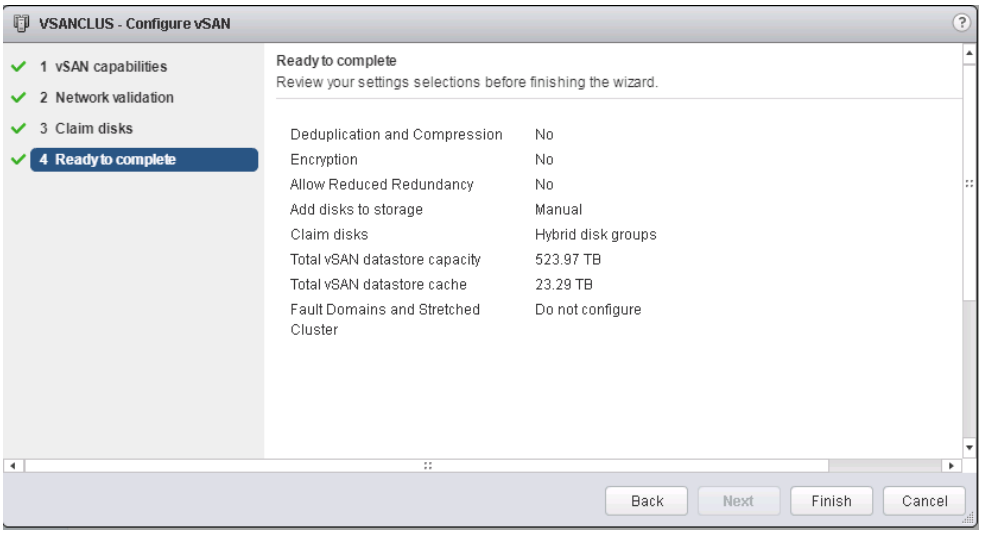

Create storage policy for VMware vSAN

Virtual machine storage policies form the basis of VMware’s software-defined storage vision. Rather than deploying virtual machines directly to a data store, you choose a virtual machine storage policy during the initial deployment process. The policy defines the characteristics and capabilities of the storage required by the virtual machine. On the basis of the policy, the correct underlying storage is chosen for the virtual machine. These policies are used to provide different levels of availability and performance for virtual machines based on the application requirements.

Follow the steps here to create the virtual machine storage policy. In this example, the policy defines a space-saving high-availability approach. However, you should evaluate the demands of your application before choosing these settings. (For additional information, refer to the vSAN document at https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-C8E919D0-9D80-4AE1-826B-D180632775F3.html.)

1. On the web client homepage, select VM Storage Policies.
2. Select the default storage policy vSAN Default Storage Policy and click Edit.
3. Edit the following settings:
   - Primary level of failures to tolerate: 1
   - Number of disk stipes per object: 1
   - Disable object checksum: No
   - Failure tolerate method: RAID-1 (Mirroring) – Performance

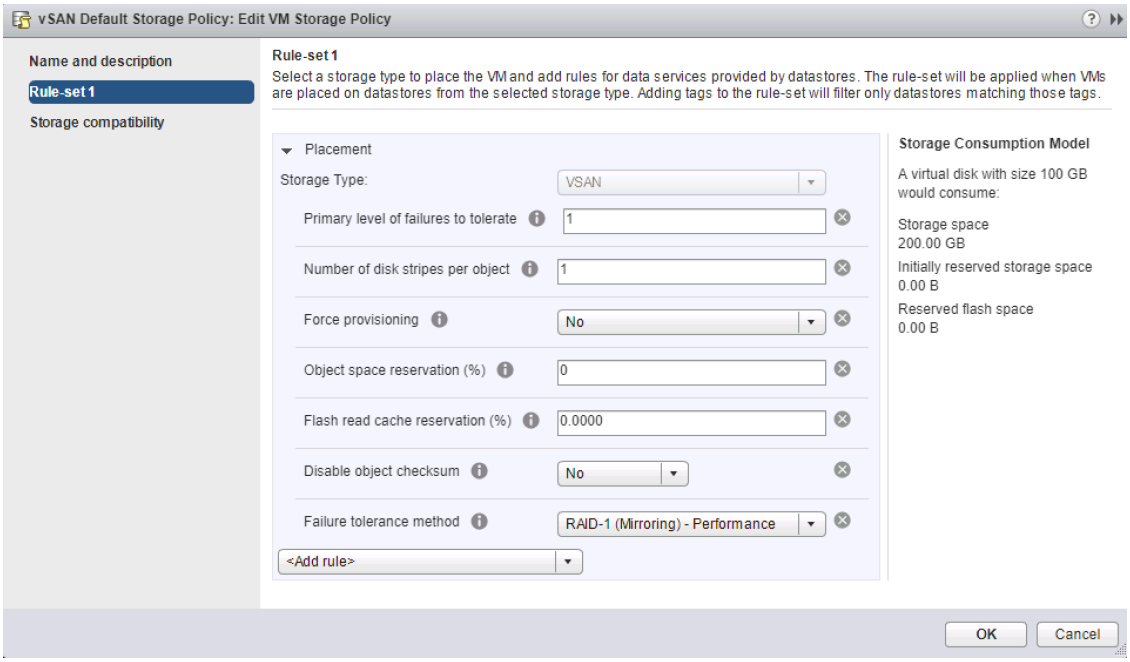

Table 8 lists the virtual machine storage policy requirements for vSAN.

Table 8. Virtual machine storage policy requirements

Policy	Definition
Number of disk stripes per object	This attribute defines the minimum number of capacity devices across which each replica of a virtual machine object is striped. A value higher than 1 may result in better performance, but also results in higher use of system resources. Default value is 1. Maximum value is 12. Do not change the default striping value. In a hybrid environment, the disk stripes are spread across magnetic disks. In the case of an all-flash configuration, the striping is across flash devices that make up the capacity layer. Make sure that your vSAN environment has sufficient capacity devices present to accommodate the request.
Flash read-cache reservation	This attribute defines the flash memory capacity reserved as read cache for the virtual machine object. It is specified as a percentage of the logical size of the virtual machine disk (VMDK) object. Reserved flash capacity cannot be used by other objects. Unreserved flash capacity is shared fairly among all objects. This option should be used only to address specific performance issues. You do not have to set a reservation to have cache space. Setting a read-cache reservation may cause a problem when you move the virtual machine object because the cache reservation settings are always included with the object. The flash read-cache reservation storage policy attribute is not supported for an all-flash cluster, and you must not use this attribute when defining a virtual machine storage policy. This attribute is supported only for hybrid configurations. Default value is 0%. Maximum value is 100%.
Number of failures to tolerate	This attribute defines the number of host and device failures that a virtual machine object can tolerate. For n failures tolerated, each piece of data written is stored in n+1 places, including parity copies if you are using RAID 5 or RAID 6. When provisioning a virtual machine, if you do not choose a storage policy, vSAN assigns this policy as the default virtual machine storage policy. Default value is 1. Maximum value is 3. If fault domains are configured, 2n+1 fault domains with hosts contributing capacity are required. A host, which is not part of any fault domain, is considered its own single-host fault domain. Default value is 1. Maximum value is 3.
Forced provisioning	If this option is set to Yes, the object is provisioned even if the number of failures to tolerate, number of disk stripes per object, and flash read-cache reservation policies specified in the storage policy cannot be satisfied by the data store. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is not possible. The default setting of No is acceptable for most production environments. vSAN fails to provision a virtual machine when the policy requirements are not met, but it successfully creates the user-defined storage policy.
Object-space reservation	This attribute defines the percentage of the logical size of the VMDK object that must be reserved, or thick provisioned, when deploying virtual machines. Default value is 0%. Maximum value is 100%.
Disable object checksum	If the option is set to No, the object calculates checksum information to help ensure the integrity of its data. If this option is set to Yes, the object does not calculate checksum information. vSAN uses end-to-end checksum to help ensure the integrity of data by confirming that each copy of a file is exactly the same as the source file. The system checks the validity of the data during read-write operations, and if an error is detected, vSAN repairs the data or reports the error. If a checksum mismatch is detected, vSAN automatically repairs the data by overwriting the incorrect data with the correct data. Checksum calculation and error-correction are performed as background operations. The default setting for all objects in the cluster is No, which means that checksum is enabled.
Failure-tolerance method	This attribute specifies whether the data replication method optimizes for performance or capacity. If you select RAID-1 (Mirroring) - Performance, vSAN uses more disk space to place the components of objects, but provides better performance for accessing the objects. If you select RAID-5/6 (Erasure Coding) - Capacity, vSAN uses less disk space, but the performance is reduced. You can use RAID 5 by applying the RAID-5/6 (Erasure Coding) - Capacity attribute to clusters with four or more fault domains, and set the number of failures to tolerate to 1. You can use RAID 6 by applying the RAID-5/6 (Erasure Coding) - Capacity attribute to clusters with six or more fault domains, and set the number of failures to tolerate to 2. For more information about RAID 5 or RAID 6, see Using RAID 5 or RAID 6 Erasure Coding.
IOPS limit for object	This attribute defines the IOPS limit for an object, such as a VMDK. IOPS is calculated as the number of I/O operations, using a weighted size. If the system uses the default base size of 32 KB, a 64-KB I/O value represents two I/O operations. When calculating IOPS, read and write are considered equivalent, but the cache hit ratio and sequence are not considered. If a disk's IOPS exceeds the limit, I/O operations are throttled. If the IOPS limit for an object is set to 0, IOPS limits are not enforced. vSAN allows the object to double the rate of the IOPS limit during the first second of operation or after a period of inactivity.

Conclusion

Implementing VMware vSAN on Cisco UCS allows customers to achieve excellent performance with a simplified management experience. Customers can continue to achieve the performance benefits of a Cisco UCS solution for applications hosted on their virtualized environments with VMware vSphere with vSAN as the hypervisor-converged storage solution.

For more information

1. https://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-s3260-storage-server/index.html
2. https://docs.vmware.com/en/VMware-vSphere/6.5/com.vmware.vsphere.virtualsan.doc/GUID-AEF15062-1ED9-4E2B-BA12-A5CE0932B976.html

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Printed in USA C11-740752-00 05/18