Administering VMware Virtual SAN

VMware vSphere 6.5

vSAN 6.6

vmware®

You can find the most up-to-date technical documentation on the VMware Web site at:

http://www.vmware.com/support/

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

Contents

About VMware Virtual SAN

Administering VMware Virtual SAN describes how to configure, manage, and monitor a VMware Virtual SAN (vSAN) cluster in a VMware vSphere® environment. In addition, *Administering VMware Virtual SAN* explains how to organize the local physical storage resources that serve as storage capacity devices in a Virtual SAN cluster, define storage policies for virtual machines deployed to Virtual SAN datastores, and manage failures in a Virtual SAN cluster.

Intended Audience

This information is for experienced virtualization administrators who are familiar with virtualization technology, day-to-day data center operations, and Virtual SAN concepts.

vSphere Client HTML5 for Virtual SAN

The vSphere Client is a new HTML5-based client that ships with vCenter Server alongside the vSphere Web Client. The new vSphere Client uses many of the same interface terminologies, topologies, and workflows as the vSphere Web Client. However, the vSphere Client does not support Virtual SAN. Users of Virtual SAN should continue to use the vSphere Web Client for those processes.

NOTE Not all functionality in the vSphere Web Client has been implemented for the vSphere Client in the vSphere 6.5 release. For an up-to-date list of unsupported functionality, see *Functionality Updates for the vSphere Client Guide* at http://www.vmware.com/info?id=1413.

Updated Information 1

Administering VMware Virtual SAN is updated with each release of the product or when necessary.

This table provides the update history of *Administering VMware Virtual SAN*.

Revision	Description
EN-002503-01	▪ The topic "Configure Network Interface for Witness Traffic," on page 67 was updated to correct the syntax of a command in the procedure. ▪ Additional minor revisions.
EN-002503-00	Initial release.

Introduction to Virtual SAN

2

VMware Virtual SAN (vSAN) is a distributed layer of software that runs natively as a part of the ESXi hypervisor. Virtual SAN aggregates local or direct-attached capacity devices of a host cluster and creates a single storage pool shared across all hosts in the Virtual SAN cluster.

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

This chapter includes the following topics:

- "Virtual SAN Concepts," on page 11
- "Virtual SAN Terms and Definitions," on page 13
- "Virtual SAN and Traditional Storage," on page 16
- "Building a Virtual SAN Cluster," on page 17
- "Integrating with Other VMware Software," on page 17
- "Limitations of Virtual SAN," on page 18

Virtual SAN Concepts

VMware Virtual SAN uses a software-defined approach that creates shared storage for virtual machines. It virtualizes the local physical storage resources of ESXi hosts and turns them into pools of storage that can be divided and assigned to virtual machines and applications according to their quality of service requirements. Virtual SAN is implemented directly in the ESXi hypervisor.

You can configure Virtual SAN to work as either a hybrid or all-flash cluster. In hybrid clusters, flash devices are used for the cache layer and magnetic disks are used for the storage capacity layer. In all-flash clusters, flash devices are used for both cache and capacity.

You can activate Virtual SAN on your existing host clusters and when you create new clusters. Virtual SAN aggregates all local capacity devices into a single datastore shared by all hosts in the Virtual SAN cluster. You can expand the datastore by adding capacity devices or hosts with capacity devices to the cluster. VMware recommends that the ESXi hosts in the cluster share similar or identical configurations across all cluster members, including similar or identical storage configurations. This ensures balanced virtual machine storage components across all devices and hosts in the cluster. Hosts without any local devices also can participate and run their virtual machines on the Virtual SAN datastore.

If a host contributes its local storage devices to the Virtual SAN datastore, it must provide at least one device for flash cache and at least one device for capacity, also called a data disk.

The devices on the contributing host form one or more disk groups. Each disk group contains one flash cache device, and one or multiple capacity devices for persistent storage. Each host can be configured to use multiple disk groups.

For best practices, capacity considerations, and general recommendations about designing and sizing a Virtual SAN cluster, see the *VMware Virtual SAN Design and Sizing Guide*.

Characteristics of Virtual SAN

This topic summarizes characteristics that apply to Virtual SAN, as well as its clusters and datastores.

Virtual SAN provides numerous benefits to your environment.

Table 2-1. Virtual SAN Features

Supported Features	Description
Shared storage support	Virtual SAN supports VMware features that require shared storage, such as HA, vMotion, and DRS. For example, if a host becomes overloaded, DRS can migrate virtual machines to other hosts in the cluster.
Just a Bunch Of Disks (JBOD)	Virtual SAN supports JBOD for use in a blade server environment. If your cluster contains blade servers, you can extend the capacity of the datastore with JBOD storage that is connected to the blade servers.
On-disk format	Virtual SAN 6.6 supports on-disk virtual file format 5.0, that provides highly scalable snapshot and clone management support per Virtual SAN cluster. For information about the number of virtual machine snapshots and clones supported per Virtual SAN cluster, see the *Configuration Maximums* documentation.
All-flash and hybrid configurations	Virtual SAN can be configured for all-flash or hybrid cluster.
Fault domains	Virtual SAN supports configuring fault domains to protect hosts from rack or chassis failure when the Virtual SAN cluster spans across multiple racks or blade server chassis in a data center.
Stretched cluster	Virtual SAN supports stretched clusters that span across two geographic locations.
Virtual SAN health service	Virtual SAN health service includes preconfigured health check tests to monitor, troubleshoot, diagnose the cause of cluster component problems, and identify any potential risk.
Virtual SAN performance service	Virtual SAN performance service includes statistical charts used to monitor IOPS, throughput, latency, and congestion. You can monitor performance of a Virtual SAN cluster, host, disk group, disk, and VMs.
Integration with vSphere storage features	Virtual SAN integrates with vSphere data management features traditionally used with VMFS and NFS storage. These features include snapshots, linked clones, vSphere Replication, and vSphere APIs for Data Protection.
Virtual Machine Storage Policies	Virtual SAN works with VM storage policies to support a VM-centric approach to storage management. If you do not assign a storage policy to the virtual machine during deployment, the Virtual SAN Default Storage Policy is automatically assigned to the VM.
Rapid provisioning	Virtual SAN enables rapid provisioning of storage in the vCenter Server® during virtual machine creation and deployment operations.

Virtual SAN Terms and Definitions

Virtual SAN introduces specific terms and definitions that are important to understand.

Before you get started with Virtual SAN, review the key Virtual SAN terms and definitions.

Disk group

A disk group is a unit of physical storage capacity on a host and a group of physical devices that provide performance and capacity to the Virtual SAN cluster. On each ESXi host that contributes its local devices to a Virtual SAN cluster, devices are organized into disk groups.

Each disk group must have one flash cache device and one or multiple capacity devices. The devices used for caching cannot be shared across disk groups, and cannot be used for other purposes. A single caching device must be dedicated to a single disk group. In hybrid clusters, flash devices are used for the cache layer and magnetic disks are used for the storage capacity layer. In an all-flash cluster, flash devices are used for both cache and capacity. For information about creating and managing disk groups, see Chapter 11, "Device Management in a Virtual SAN Cluster," on page 99.

Consumed capacity

Consumed capacity is the amount of physical capacity consumed by one or more virtual machines at any point. Consumed capacity is determined by many factors, including the consumed size of your VMDKs, protection replicas, and so on. When calculating for cache sizing, do not consider the capacity used for protection replicas.

Object-based storage

Virtual SAN stores and manages data in the form of flexible data containers called objects. An object is a logical volume that has its data and metadata distributed across the cluster. For example, every VMDK is an object, as is every snapshot. When you provision a virtual machine on a Virtual SAN datastore, Virtual SAN creates a set of objects comprised of multiple components for each virtual disk. It also creates the VM home namespace, which is a container object that stores all metadata files of your virtual machine. Based on the assigned virtual machine storage policy, Virtual SAN provisions and manages each object individually, which might also involve creating a RAID configuration for every object.

When Virtual SAN creates an object for a virtual disk and determines how to distribute the object in the cluster, it considers the following factors:

■ Virtual SAN verifies that the virtual disk requirements are applied according to the specified virtual machine storage policy settings.

■ Virtual SAN verifies that the correct cluster resources are utilized at the time of provisioning. For example, based on the protection policy, Virtual SAN determines how many replicas to create. The performance policy determines the amount of flash read cache allocated for each replica and how many stripes to create for each replica and where to place them in the cluster.

■ Virtual SAN continually monitors and reports the policy compliance status of the virtual disk. If you find any noncompliant policy status, you must troubleshoot and resolve the underlying problem.

NOTE   When required, you can edit VM storage policy settings. Changing the storage policy settings does not affect virtual machine access. Virtual SAN actively throttles the storage and network resources used for reconfiguration to minimize the impact of object reconfiguration to normal workloads. When you change VM storage policy settings, Virtual SAN might initiate an object recreation process and subsequent resynchronization. See "About Virtual SAN Cluster Resynchronization," on page 133.

■ Virtual SAN verifies that the required protection components, such as mirrors and witnesses, are placed on separate hosts or fault domains. For example, to rebuild components during failure, Virtual SAN looks for ESXi hosts that satisfy the placement rules where protection components of virtual machine objects must be placed on two different hosts (not on the same host), or across different fault domains.

Virtual SAN datastore

After you enable Virtual SAN on a cluster, a single Virtual SAN datastore is created. It appears as another type of datastore in the list of datastores that might be available, including Virtual Volume, VMFS, and NFS. A single Virtual SAN datastore can provide different service levels for each virtual machine or each virtual disk. In vCenter Server®, storage characteristics of the Virtual SAN datastore appear as a set of capabilities. You can reference these capabilities when defining a storage policy for virtual machines. When you later deploy virtual machines, Virtual SAN uses this policy to place virtual machines in the optimal manner based on the requirements of each virtual machine. For general information about using storage policies, see the *vSphere Storage* documentation.

A Virtual SAN datastore has specific characteristics to consider.

■ Virtual SAN provides a single Virtual SAN datastore accessible to all hosts in the cluster, whether or not they contribute storage to the cluster. Each host can also mount any other datastores, including Virtual Volumes, VMFS, or NFS.

■ You can use Storage vMotion to move virtual machines between the Virtual SAN datastores, NFS and VMFS datastores.

■ Only magnetic disks and flash devices used for capacity can contribute to the datastore capacity. The devices used for flash cache are not counted as part of the datastore.

Objects and components

Each object is composed of a set of components, determined by capabilities that are in use in the VM Storage Policy. For example, when the **Primary level of failures to tolerate** policy configured to one, Virtual SAN ensures that the protection components, such as replicas and witnesses of the object are placed on separate hosts in the Virtual SAN cluster, where each replica is an object component. In addition, in the same policy, if the **Number of disk stripes per object** configured to two or more, Virtual SAN also stripes the object across multiple capacity devices and each stripe is considered a component of the specified object. When needed, Virtual SAN might also break large objects into multiple components.

A Virtual SAN datastore contains the following object types:

VM Home Namespace	The virtual machine home directory where all virtual machine configuration files are stored, such as `.vmx`, log files, vmdks, snapshot delta description files, and so on.
VMDK	A virtual machine disk or `.vmdk` file that stores the contents of the virtual machine's hard disk drive.
VM Swap Object	Created when a virtual machine is powered on.
Snapshot Delta VMDKs	Created when virtual machine snapshots are taken.
Memory object	Created when the snapshot memory option is selected when creating or suspending a virtual machine.

Virtual Machine Compliance Status: Compliant and Noncompliant

A virtual machine is considered noncompliant when one or more of its objects fail to meet the requirements of its assigned storage policy. For example, the status might become noncompliant when one of the mirror copies is inaccessible. If your virtual machines are in compliance with the requirements defined in the storage policy, the status of your virtual machines is compliant. From the **Physical Disk Placement** tab on the Virtual Disks page, you can verify the virtual machine object compliance status. For information about troubleshooting a Virtual SAN cluster, see "Handling Failures in Virtual SAN," on page 152.

Component State: Degraded and Absent states

Virtual SAN acknowledges the following failure states for components:

- Degraded. A component is Degraded when Virtual SAN detects a permanent component failure and determines that the failed component will never recover to its original working state. As a result, Virtual SAN starts to rebuild the degraded components immediately. This state might occur when a component is on a failed device.

- Absent. A component is Absent when Virtual SAN detects a temporary component failure where components, including all its data, might recover and return Virtual SAN to its original state. This state might occur when you are restarting hosts or if you unplug a device from a Virtual SAN host. Virtual SAN starts to rebuild the components in absent status after waiting for 60 minutes.

Object State: Healthy and Unhealthy

Depending on the type and number of failures in the cluster, an object might be in one of the following states:

- Healthy. When at least one full RAID 1 mirror is available, or the minimum required number of data segments are available, the object is considered healthy.

- Unhealthy. When no full mirror is available or the minimum required number of data segments are unavailable for RAID 5 or RAID 6 objects; or fewer than 50 percent of an object's votes are available. This may be due to multiple failures in the cluster. When the operational status of an object is considered unhealthy, it impacts the availability of the associated VM.

Witness

A witness is a component that contains only metadata and does not contain any actual application data. It serves as a tiebreaker when a decision needs to be made regarding the availability of the surviving datastore components, after a potential failure. A witness consumes approximately 2 MB of space for metadata on the Virtual SAN datastore when using on-disk format 1.0 and 4 MB for on-disk format for version 2.0 and later.

Virtual SAN 6.0 and later maintains quorum with an asymmetrical voting system where each component might have more than one vote to decide the availability of objects. Greater than 50 percent of the votes that make up a VM's storage object must be accessible at all times for the object to be considered available. When 50 percent or fewer votes are accessible to all hosts, the object is no longer available to the Virtual SAN datastore. This impacts the availability of the associated VM.

Storage Policy-Based Management (SPBM)

When you use Virtual SAN, you can define virtual machine storage requirements, such as performance and availability, in the form of a policy. Virtual SAN ensures that the virtual machines deployed to Virtual SAN datastores are assigned at least one virtual machine storage policy. When you know the storage requirements of your virtual machines, you can define storage policies and assign the policies to your virtual machines. If you do not apply a storage policy when deploying virtual machines, Virtual SAN automatically assigns a default Virtual SAN policy with **Primary level of failures to tolerate** configured to one, a single

disk stripe for each object, and thin provisioned virtual disk. For best results, you should define your own virtual machine storage policies, even if the requirements of your policies are the same as those defined in the default storage policy. For information about working with Virtual SAN storage policies, see Chapter 13, "Using Virtual SAN Policies," on page 123.

Ruby vSphere Console (RVC)

The Ruby vSphere Console (RVC) provides a command-line interface used for managing and troubleshooting the Virtual SAN cluster. RVC gives you a cluster-wide view, instead of the host-centric view offered by `esxcli`. RVC is bundled with vCenter Server Appliance and vCenter Server for Windows, so you do not need to install it separately. For information about the RVC commands, see the *RVC Command Reference Guide*.

vSphere PowerCLI

VMware vSphere PowerCLI adds command-line scripting support for Virtual SAN, to help you automate configuration and management tasks. vSphere PowerCLI provides a Windows PowerShell interface to the vSphere API. PowerCLI includes cmdlets for administering Virtual SAN components. For information about using vSphere PowerCLI, see *vSphere PowerCLI Documentation*.

Virtual SAN Observer

The VMware Virtual SAN Observer is a Web-based tool that runs on RVC and is used for in-depth performance analysis and monitoring of the Virtual SAN cluster. Use Virtual SAN Observer for information about the performance statistics of the capacity layer, detailed statistical information about physical disk groups, current CPU usage, consumption of Virtual SAN memory pools, and physical and in-memory object distribution across Virtual SAN clusters.

For information about configuring, launching, and using RVC and the Virtual SAN Observer, see the *Virtual SAN Troubleshooting Reference Manual*.

Virtual SAN and Traditional Storage

Although Virtual SAN shares many characteristics with traditional storage arrays, the overall behavior and function of Virtual SAN is different. For example, Virtual SAN can manage and work only with ESXi hosts and a single Virtual SAN instance can support only one cluster.

Virtual SAN and traditional storage also differ in the following key ways:

- Virtual SAN does not require external networked storage for storing virtual machine files remotely, such as on a Fibre Channel (FC) or Storage Area Network (SAN).

- Using traditional storage, the storage administrator preallocates storage space on different storage systems. Virtual SAN automatically turns the local physical storage resources of the ESXi hosts into a single pool of storage. These pools can be divided and assigned to virtual machines and applications according to their quality of service requirements.

- Virtual SAN has no concept of traditional storage volumes based on LUNs or NFS shares, although the iSCSI target service uses LUNs to enable an initiator on a remote host to transport block-level data to a storage device in the Virtual SAN cluster.

- Some standard storage protocols, such as FCP, do not apply to Virtual SAN.

- Virtual SAN is highly integrated with vSphere. You do not need dedicated plug-ins or a storage console for Virtual SAN, compared to traditional storage. You can deploy, manage, and monitor Virtual SAN by using the vSphere Web Client.

- A dedicated storage administrator does not need to manage Virtual SAN. Instead a vSphere administrator can manage a Virtual SAN environment.

■   With Virtual SAN usage, VM storage policies are automatically assigned when you deploy new VMs. The storage policies can be changed dynamically as needed.

Building a Virtual SAN Cluster

If you are considering Virtual SAN, you can choose from more than one configuration solution for deploying a Virtual SAN cluster.

Depending on your requirement, you can deploy Virtual SAN in one of the following ways.

Virtual SAN Ready Node

The Virtual SAN Ready Node is a preconfigured solution of the Virtual SAN software provided by VMware partners, such as Cisco, Dell, Fujitsu, IBM, and Supermicro. This solution includes validated server configuration in a tested, certified hardware form factor for Virtual SAN deployment that is recommended by the server OEM and VMware. For information about the Virtual SAN Ready Node solution for a specific partner, visit the VMware Partner web site.

User-Defined Virtual SAN Cluster

You can build a Virtual SAN cluster by selecting individual software and hardware components, such as drivers, firmware, and storage I/O controllers that are listed in the Virtual SAN Compatibility Guide (VCG) web site at http://www.vmware.com/resources/compatibility/search.php. You can choose any servers, storage I/O controllers, capacity and flash cache devices, memory, any number of cores you must have per CPU, and so on that are certified and listed on the VCG Web site. Review the compatibility information on the VCG Web site before choosing software and hardware components, drivers, firmware, and storage I/O controllers that are supported by Virtual SAN. When designing a Virtual SAN cluster, use only devices, firmware, and drivers that are listed on the VCG Web site. Using software and hardware versions that are not listed in the VCG might cause cluster failure or unexpected data loss. For information about designing a Virtual SAN cluster, see Chapter 4, "Designing and Sizing a Virtual SAN Cluster," on page 23.

Integrating with Other VMware Software

After you have Virtual SAN up and running, it is integrated with the rest of the VMware software stack. You can do most of what you can do with traditional storage by using vSphere components and features including vSphere vMotion, snapshots, clones, Distributed Resource Scheduler (DRS), vSphere High Availability, vCenter Site Recovery Manager, and more.

Integrating with vSphere HA

You can enable vSphere HA and Virtual SAN on the same cluster. As with traditional datastores, vSphere HA provides the same level of protection for virtual machines on Virtual SAN datastores. This level of protection imposes specific restrictions when vSphere HA and Virtual SAN interact. For specific considerations about integrating vSphere HA and Virtual SAN, see "Using Virtual SAN and vSphere HA," on page 54.

Integrating with VMware Horizon View

You can integrate Virtual SAN with VMware Horizon View. When integrated, Virtual SAN provides the following benefits to virtual desktop environments:

■   High-performance storage with automatic caching

■   Storage policy-based management, for automatic remediation

For information about integrating Virtual SAN with VMware Horizon, see the *VMware Horizon with View* documentation. For designing and sizing VMware Horizon View for Virtual SAN, see the *Designing and Sizing Guide for Horizon View*.

Limitations of Virtual SAN

This topic discusses the limitations of Virtual SAN.

When working with Virtual SAN, consider the following limitations:

- Virtual SAN does not support hosts participating in multiple Virtual SAN clusters. However, a Virtual SAN host can access other external storage resources that are shared across clusters.

- Virtual SAN does not support vSphere DPM and Storage I/O Control.

- Virtual SAN does not support SCSI reservations.

- Virtual SAN does not support RDM, VMFS, diagnostic partition, and other device access features.

Requirements for Enabling Virtual SAN

<div align="right">

3

</div>

Before you activate Virtual SAN, verify that your environment meets all requirements.

This chapter includes the following topics:

Hardware Requirements for vSAN

Verify that the ESXi hosts in your organization meet the vSAN hardware requirements.

Storage Device Requirements

All capacity devices, drivers, and firmware versions in your Virtual SAN configuration must be certified and listed in the Virtual SAN section of the *VMware Compatibility Guide*.

Table 3-1. Storage Device Requirements for vSAN Hosts

Storage Component	Requirements
Cache	■ One SAS or SATA solid-state disk (SSD) or PCIe flash device. ■ Before calculating the **Primary level of failures to tolerate**, check the size of the flash caching device in each disk group. Verify that it provides at least 10 percent of the anticipated storage consumed on the capacity devices, not including replicas such as mirrors. ■ vSphere Flash Read Cache must not use any of the flash devices reserved for vSAN cache. ■ The cache flash devices must not be formatted with VMFS or another file system.
Virtual machine data storage	■ For hybrid group configuration, make sure that at least one SAS, NL-SAS, or SATA magnetic disk is available. ■ For all-flash disk group configuration, make sure at least one SAS, or SATA solid-state disk (SSD), or PCIe flash device.
Storage controllers	One SAS or SATA host bus adapter (HBA), or a RAID controller that is in passthrough mode or RAID 0 mode.

Memory

The memory requirements for vSAN depend on the number of disk groups and devices that the ESXi hypervisor must manage. Each host must contain a minimum of 32 GB of memory to accommodate the maximum number of disk groups (5) and maximum number of capacity devices per disk group (7).

Flash Boot Devices

During installation, the ESXi installer creates a coredump partition on the boot device. The default size of the coredump partition satisfies most installation requirements.

- If the memory of the ESXi host has 512 GB of memory or less, you can boot the host from a USB, SD, or SATADOM device. When you boot a vSAN host from a USB device or SD card, the size of the boot device must be at least 4 GB.

- If the memory of the ESXi host has more than 512 GB, you must boot the host from a SATADOM or disk device. When you boot a vSAN host from a SATADOM device, you must use single-level cell (SLC) device. The size of the boot device must be at least 16 GB.

NOTE vSAN 6.5 and later enables you to resize an existing coredump partition on an ESXi host in a vSAN cluster, so you can boot from USB/SD devices. For more information, see the VMware knowledge base article at http://kb.vmware.com/kb/2147881.

When you boot an ESXi 6.0 or later host from USB device or from SD card, vSAN trace logs are written to RAMDisk. These logs are automatically offloaded to persistent media during shutdown or system crash (panic). This is the only support method for handling vSAN traces when booting an ESXi from a USB stick or SD card. If a power failure occurs, vSAN trace logs are not preserved.

When you boot an ESXi 6.0 or later host from a SATADOM device, vSAN trace logs are written directly to the SATADOM device. Therefore it is important that the SATADOM device meets the specifications outlined in this guide.

Cluster Requirements for Virtual SAN

Verify that a host cluster meets the requirements for enabling Virtual SAN.

- All capacity devices, drivers, and firmware versions in your Virtual SAN configuration must be certified and listed in the Virtual SAN section of the *VMware Compatibility Guide*.

- A Virtual SAN cluster must contain a minimum of three hosts that contribute capacity to the cluster. For information about the considerations for a three-host cluster, see "Design Considerations for a Virtual SAN Cluster," on page 31.

- A host that resides in a Virtual SAN cluster must not participate in other clusters.

Software Requirements for Virtual SAN

Verify that the vSphere components in your environment meet the software version requirements for using Virtual SAN.

To use the full set of Virtual SAN capabilities, the ESXi hosts that participate in Virtual SAN clusters must be version 6.5 or later. During the Virtual SAN upgrade from previous versions, you can keep the current on-disk format version, but you cannot use many of the new features. Virtual SAN 6.6 and later software supports all on-disk formats.

Networking Requirements for Virtual SAN

Verify that the network infrastructure and the networking configuration on the ESXi hosts meet the minimum networking requirements for Virtual SAN.

Table 3-2. Networking Requirements for Virtual SAN

Networking Component	Requirement
Host Bandwidth	Each host must have minimum bandwidth dedicated to Virtual SAN. ■ Dedicated 1 Gbps for hybrid configurations ■ Dedicated or shared 10 Gbps for all-flash configurations For information about networking considerations in Virtual SAN, see "Designing the Virtual SAN Network," on page 32.
Connection between hosts	Each host in the Virtual SAN cluster, regardless of whether it contributes capacity, must have a VMkernel network adapter for Virtual SAN traffic. See "Set Up a VMkernel Network for Virtual SAN," on page 47.
Host network	All hosts in your Virtual SAN cluster must be connected to a Virtual SAN Layer 2 or Layer 3 network.
IPv4 and IPv6 support	The Virtual SAN network supports both IPv4 and IPv6.

License Requirements

Verify that you have a valid license for Virtual SAN.

Using Virtual SAN in production environments requires a special license that you assign to the Virtual SAN clusters.

You can assign a standard Virtual SAN license to the cluster, or a license that covers advanced functions. Advanced features include RAID 5/6 erasure coding, and deduplication and compression. An enterprise license is required for IOPS limits and stretched clusters. For information about assigning licenses, see "Configure License Settings for a Virtual SAN Cluster," on page 52.

The capacity of the license must cover the total number of CPUs in the cluster.

Designing and Sizing a Virtual SAN Cluster

4

For best performance and use, plan the capabilities and configuration of your hosts and their storage devices before you deploy Virtual SAN in a vSphere environment. Carefully consider certain host and networking configurations within the Virtual SAN cluster.

The *Administering VMware vSAN* documentation examines the key points about designing and sizing a Virtual SAN cluster. For detailed instructions about designing and sizing a Virtual SAN cluster, see *VMware Virtual SAN Design and Sizing Guide*.

This chapter includes the following topics:

■ "Designing and Sizing Virtual SAN Storage Components," on page 23

■ "Designing and Sizing vSAN Hosts," on page 29

■ "Design Considerations for a Virtual SAN Cluster," on page 31

■ "Designing the Virtual SAN Network," on page 32

■ "Best Practices for Virtual SAN Networking," on page 34

■ "Designing and Sizing Virtual SAN Fault Domains," on page 34

■ "Using Boot Devices and vSAN," on page 35

■ "Persistent Logging in a Virtual SAN Cluster," on page 35

Designing and Sizing Virtual SAN Storage Components

Plan capacity and cache based on the expected consumption. Consider the requirements for availability and endurance.

■ Planning Capacity in Virtual SAN on page 24

You can size the capacity of a Virtual SAN datastore to accommodate the virtual machines (VMs) files in the cluster and to handle failures and maintenance operations.

■ Design Considerations for Flash Caching Devices in Virtual SAN on page 26

Plan the configuration of flash devices for Virtual SAN cache and all-flash capacity to provide high performance and required storage space, and to accommodate future growth.

■ Design Considerations for Flash Capacity Devices in Virtual SAN on page 27

Plan the configuration of flash capacity devices for Virtual SAN all-flash configurations to provide high performance and required storage space, and to accommodate future growth.

■ Design Considerations for Magnetic Disks in Virtual SAN on page 28

Plan the size and number of magnetic disks for capacity in hybrid configurations by following the requirements for storage space and performance.

■ Design Considerations for Storage Controllers in Virtual SAN on page 29

Include storage controllers on the hosts of a Virtual SAN cluster that best satisfy the requirements for performance and availability.

Planning Capacity in Virtual SAN

You can size the capacity of a Virtual SAN datastore to accommodate the virtual machines (VMs) files in the cluster and to handle failures and maintenance operations.

Raw Capacity

To determine the raw capacity of a Virtual SAN datastore, multiply the total number of disk groups in the cluster by the size of the capacity devices in those disk groups, and subtract the overhead required by the Virtual SAN on-disk format.

Primary level of Failures to Tolerate

When you plan the capacity of the Virtual SAN datastore, not including the number of virtual machines and the size of their VMDK files, you must consider the **Primary level of failures to tolerate** and the **Failure tolerance method** attributes of the virtual machine storage policies for the cluster.

The **Primary level of failures to tolerate** has an important role when you plan and size storage capacity for Virtual SAN. Based on the availability requirements of a virtual machine, the setting might result in doubled consumption or more, compared with the consumption of a virtual machine and its individual devices.

For example, if the **Failure tolerance method** is set to **RAID-1 (Mirroring) - Performance** and the **Primary level of failures to tolerate** (PFTT) is set to 1, virtual machines can use about 50 percent of the raw capacity. If the PFTT is set to 2, the usable capacity is about 33 percent. If the PFTT is set to 3, the usable capacity is about 25 percent.

But if the **Failure tolerance method** is set to **RAID-5/6 (Erasure Coding) - Capacity** and the PFTT is set to 1, virtual machines can use about 75 percent of the raw capacity. If the PFTT is set to 2, the usable capacity is about 67 percent. For more information about RAID 5/6, see "Using RAID 5 or RAID 6 Erasure Coding," on page 76.

For information about the attributes in a Virtual SAN storage policy, see Chapter 13, "Using Virtual SAN Policies," on page 123.

Calculating Required Capacity

Plan the capacity required for the virtual machines in a cluster with RAID 1 mirroring based on the following criteria:

1. Calculate the storage space that the virtual machines in the Virtual SAN cluster are expected to consume.

   ```
   expected overall consumption = number of VMs in the cluster * expected percentage of
   consumption per VMDK
   ```

2. Consider the **Primary level of failures to tolerate** attribute configured in the storage policies for the virtual machines in the cluster. This attribute directly impacts the number of replicas of a VMDK file on hosts in the cluster.

   ```
   datastore capacity = expected overall consumption * (PFTT + 1)
   ```

3. Estimate the overhead requirement of the Virtual SAN on-disk format.

   ■ On-disk format version 3.0 and later adds an additional overhead, typically no more than 1-2 percent capacity per device. Deduplication and compression with software checksum enabled require additional overhead of approximately 6.2 percent capacity per device.

- On-disk format version 2.0 adds an additional overhead, typically no more than 1-2 percent capacity per device.

- On-disk format version 1.0 adds an additional overhead of approximately 1 GB per capacity device.

Capacity Sizing Guidelines

- Keep at least 30 percent unused space to prevent Virtual SAN from rebalancing the storage load. Virtual SAN rebalances the components across the cluster whenever the consumption on a single capacity device reaches 80 percent or more. The rebalance operation might impact the performance of applications. To avoid these issues, keep storage consumption to less than 70 percent.

- Plan extra capacity to handle potential failure or replacement of capacity devices, disk groups, and hosts. When a capacity device is not reachable, Virtual SAN recovers the components from another device in the cluster. When a flash cache device fails or is removed, Virtual SAN recovers the components from the entire disk group.

- Reserve extra capacity to make sure that Virtual SAN recovers components after a host failure or when a host enters maintenance mode. For example, provision hosts with enough capacity so that you have sufficient free capacity left for components to successfully rebuild after a host failure or during maintenance. This is important when you have more than three hosts, so you have sufficient free capacity to rebuild the failed components. If a host fails, the rebuilding takes place on the storage available on another host, so that another failure can be tolerated. However, in a three-host cluster, Virtual SAN will not perform the rebuild operation if the **Primary level of failures to tolerate** is set to 1 because when one host fails, only two hosts remain in the cluster. To tolerate a rebuild after a failure, you must have at least three hosts.

- Provide enough temporary storage space for changes in the Virtual SAN VM storage policy. When you dynamically change a VM storage policy, Virtual SAN might create a layout of the replicas that make up an object. When Virtual SAN instantiates and synchronizes those replicas with the original replica, the cluster must temporarily provide additional space.

- If you plan to use advanced features, such as software checksum or deduplication and compression, reserve additional capacity to handle the operational overhead.

Considerations for Virtual Machine Objects

When you plan the storage capacity in the Virtual SAN datastore, consider the space required in the datastore for the VM home namespace objects, snapshots, and swap files.

- VM Home Namespace. You can assign a storage policy specifically to the home namespace object for a virtual machine. To prevent unnecessary allocation of capacity and cache storage, Virtual SAN applies only the **Primary level of failures to tolerate** and the **Force provisioning** settings from the policy on the VM home namespace. Plan storage space to meet the requirements for a storage policy assigned to a VM Home Namespace whose **Primary level of failures to tolerate** is greater than 0.

- Snapshots. Delta devices inherit the policy of the base VMDK file. Plan additional space according to the expected size and number of snapshots, and to the settings in the Virtual SAN storage policies.

  The space that is required might be different. Its size depends on how often the virtual machine changes data and how long a snapshot is attached to the virtual machine.

- Swap files. Virtual SAN uses an individual storage policy for the swap files of virtual machines. The policy tolerates a single failure, defines no striping and read cache reservation, and enables force provisioning.

Design Considerations for Flash Caching Devices in Virtual SAN

Plan the configuration of flash devices for Virtual SAN cache and all-flash capacity to provide high performance and required storage space, and to accommodate future growth.

Choosing Between PCIe or SSD Flash Devices

Choose PCIe or SSD flash devices according to the requirements for performance, capacity, write endurance, and cost of the Virtual SAN storage.

- Compatibility. The model of the PCIe or SSD devices must be listed in the Virtual SAN section of the *VMware Compatibility Guide*.

- Performance. PCIe devices generally have faster performance than SSD devices.

- Capacity. The maximum capacity that is available for PCIe devices is generally greater than the maximum capacity that is currently listed for SSD devices for Virtual SAN in the *VMware Compatibility Guide*.

- Write endurance. The write endurance of the PCIe or SSD devices must meet the requirements for capacity or for cache in all-flash configurations, and for cache in hybrid configurations.

  For information about the write endurance requirements for all-flash and hybrid configurations, see the *VMware Virtual SAN Design and Sizing Guide*. For information about the write endurance class of PCIe and SSD devices, see the Virtual SAN section of the *VMware Compatibility Guide*.

- Cost. PCIe devices generally have higher cost than SSD devices.

Flash Devices as Virtual SAN Cache

Design the configuration of flash cache for Virtual SAN for write endurance, performance, and potential growth based on these considerations.

Table 4-1. Sizing Virtual SAN Cache

Storage Configuration	Considerations
All-flash and hybrid configurations	■ The flash caching device must provide at least 10 percent of the anticipated storage that virtual machines are expected to consume, not including replicas such as mirrors.
	The **Primary level of failures to tolerate** attribute from the VM storage policy does not impact the size of the cache.
	■ A higher cache-to-capacity ratio eases future capacity growth. Oversizing cache enables you to easily add more capacity to an existing disk group without the need to increase the size of the cache.
	■ Flash caching devices must have high write endurance.
	■ When a flash caching device is at the end of its life, replacing it is more complicated than replacing a capacity device because such an operation impacts the entire disk group.
	■ If you add more flash devices to increase the size of the cache, you must create more disk groups. The ratio between flash cache devices and disk groups is always 1:1.
	A configuration of multiple disk groups provides the following advantages:
	■ Reduced risk of failure because fewer capacity devices are affected if a single caching device fails
	■ Potentially improved performance if you deploy multiple disk groups that contain smaller flash caching devices.
	However, when you configure multiple disk groups, the memory consumption of the hosts increases.
All-flash configurations	In all-flash configurations, Virtual SAN uses the cache layer for write caching only. The write cache must be able to handle very high write activities. This approach extends the life of capacity flash that might be less expensive and might have lower write endurance.
Hybrid configurations	If the read cache reservation is configured in the active VM storage policy for performance reasons, the hosts in the Virtual SAN cluster must have sufficient cache to satisfy the reservation during a post-failure rebuild or maintenance operation.
	If the available read cache is not sufficient to satisfy the reservation, the rebuild or maintenance operation fails. Use read cache reservation only if you must meet a specific, known performance requirement for a particular workload.
	The use of snapshots consumes cache resources. If you plan to use several snapshots, consider dedicating more cache than the conventional 10 percent cache-to-consumed-capacity ratio.

Design Considerations for Flash Capacity Devices in Virtual SAN

Plan the configuration of flash capacity devices for Virtual SAN all-flash configurations to provide high performance and required storage space, and to accommodate future growth.

Choosing Between PCIe or SSD Flash Devices

Choose PCIe or SSD flash devices according to the requirements for performance, capacity, write endurance, and cost of the Virtual SAN storage.

■ Compatibility. The model of the PCIe or SSD devices must be listed in the Virtual SAN section of the *VMware Compatibility Guide*.

■ Performance. PCIe devices generally have faster performance than SSD devices.

- Capacity. The maximum capacity that is available for PCIe devices is generally greater than the maximum capacity that is currently listed for SSD devices for Virtual SAN in the *VMware Compatibility Guide*.

- Write endurance. The write endurance of the PCIe or SSD devices must meet the requirements for capacity or for cache in all-flash configurations, and for cache in hybrid configurations.

  For information about the write endurance requirements for all-flash and hybrid configurations, see the *VMware Virtual SAN Design and Sizing Guide*. For information about the write endurance class of PCIe and SSD devices, see the Virtual SAN section of the *VMware Compatibility Guide*.

- Cost. PCIe devices generally have higher cost than SSD devices.

Flash Devices as Virtual SAN Capacity

In all-flash configurations, Virtual SAN does not use cache for read operations and does not apply the read-cache reservation setting from the VM storage policy. For cache, you can use a small amount of more expensive flash that has high write endurance. For capacity, you can use flash that is less expensive and has lower write endurance.

Plan a configuration of flash capacity devices by following these guidelines:

- For better performance of Virtual SAN, use more disk groups of smaller flash capacity devices.

- For balanced performance and predictable behavior, use the same type and model of flash capacity devices.

Design Considerations for Magnetic Disks in Virtual SAN

Plan the size and number of magnetic disks for capacity in hybrid configurations by following the requirements for storage space and performance.

SAS, NL-SAS, and SATA Magnetic Devices

Use SAS, NL-SAS, or SATA magnetic devices by following the requirements for performance, capacity, and cost of the Virtual SAN storage.

- Compatibility. The model of the magnetic disk must be certified and listed in the Virtual SAN section of the *VMware Compatibility Guide*.

- Performance. SAS and NL-SAS devices have faster performance than SATA disks.

- Capacity. The capacity of SAS, NL-SAS, and SATA magnetic disks for Virtual SAN is available in the Virtual SAN section of the *VMware Compatibility Guide*. Consider using a larger number of smaller devices instead of a smaller number of larger devices.

- Cost. SAS and NL-SAS devices are more expensive than SATA disks.

Using SATA disks instead of SAS and NL-SAS devices is justifiable in environments where capacity and reduced cost have higher priority than performance.

Magnetic Disks as Virtual SAN Capacity

Plan a magnetic disk configuration by following these guidelines:

- For better performance of Virtual SAN, use many magnetic disks that have smaller capacity.

  You must have enough magnetic disks that provide adequate aggregated performance for transferring data between cache and capacity. Using more small devices provides better performance than using fewer large devices. Using multiple magnetic disk spindles can speed up the destaging process.

In environments that contain many virtual machines, the number of magnetic disks is also important for read operations when data is not available in the read cache and Virtual SAN reads it from the magnetic disk. In environments that contain a small number of virtual machines, the disk number impacts read operations if the **Number of disk stripes per object** in the active VM storage policy is greater than one.

■ For balanced performance and predictable behavior, use the same type and model of magnetic disks in a Virtual SAN datastore.

■ Dedicate a high enough number of magnetic disks to satisfy the value of the **Primary level of failures to tolerate** and the **Number of disk stripes per object** attributes in the defined storage policies. For information about the VM storage policies for Virtual SAN, see Chapter 13, "Using Virtual SAN Policies," on page 123.

Design Considerations for Storage Controllers in Virtual SAN

Include storage controllers on the hosts of a Virtual SAN cluster that best satisfy the requirements for performance and availability.

■ Use storage controller models, and driver and firmware versions that are listed in the *VMware Compatibility Guide*. Search for Virtual SAN in the *VMware Compatibility Guide*.

■ Use multiple storage controllers, if possible, to improve performance and to isolate a potential controller failure to only a subset of disk groups.

■ Use storage controllers that have the highest queue depths in the *VMware Compatibility Guide*. Using controllers with high queue depth improves performance. For example, when Virtual SAN is rebuilding components after a failure or when a host enters maintenance mode.

■ Use storage controllers in passthrough mode for best performance of Virtual SAN. Storage controllers in RAID 0 mode require higher configuration and maintenance efforts compared to storage controllers in passthrough mode.

Designing and Sizing vSAN Hosts

Plan the configuration of the hosts in the vSAN cluster for best performance and availability.

Memory and CPU

Size the memory and the CPU of the hosts in the vSAN cluster based on the following considerations.

Table 4-2. Sizing Memory and CPU of vSAN Hosts

Compute Resource	Considerations
Memory	- Memory per virtual machine - Memory per host, based on the expected number of virtual machines - At least 32-GB memory for fully operational vSAN with 5 disk groups per host and 7 capacity devices per disk group Hosts that have 512-GB memory or less can boot from a USB, SD, or SATADOM device. If the memory of the host is greater than 512 GB, boot the host from a SATADOM or disk device.
CPU	- Sockets per host - Cores per socket - Number of vCPUs based on the expected number of virtual machines - vCPU-to-core ratio 10% CPU overhead for vSAN

Host Networking

Provide more bandwidth for vSAN traffic to improve performance.

- If you plan to use hosts that have 1-GbE adapters, dedicate adapters for vSAN only. For all-flash configurations, plan hosts that have dedicated or shared 10-GbE adapters.

- If you plan to use 10-GbE adapters, they can be shared with other traffic types for both hybrid and all-flash configurations.

- If a 10-GbE adapter is shared with other traffic types, use a vSphere Distributed Switch for vSAN traffic to isolate the traffic by using Network I/O Control and VLANs.

- Create a team of physical adapters for vSAN traffic for redundancy.

Multiple Disk Groups

If the flash cache or storage controller stops responding, an entire disk group can fail. As a result, vSAN rebuilds all components for the failed disk group from another location in the cluster.

Use of multiple disk groups, with each disk group providing less capacity, provides the following benefits and disadvantages:

- Benefits

  - Performance is improved because the datastore has more aggregated cache, and I/O operations are faster.

  - Risk of failure is spread among multiple disk groups.

  - If a disk group fails, vSAN rebuilds fewer components, so performance is improved.

- Disadvantages

  - Costs are increased because two or more caching devices are required.

  - More memory is required to handle more disk groups.

  - Multiple storage controllers are required to reduce the risk of a single point of failure.

Drive Bays

For easy maintenance, consider hosts whose drive bays and PCIe slots are at the front of the server body.

Blade Servers and External Storage

The capacity of blade servers usually does not scale in a vSAN datastore because they have a limited number of disk slots. To extend the planned capacity of blade servers, use external storage enclosures. For information about the supported models of external storage enclosures, see *VMware Compatibility Guide*.

Hot Plug and Swap of Devices

Consider the storage controller passthrough mode support for easy hot plugging or replacement of magnetic disks and flash capacity devices on a host. If a controller works in RAID 0 mode, you must perform additional steps before the host can discover the new drive.

Design Considerations for a Virtual SAN Cluster

Design the configuration of hosts and management nodes for best availability and tolerance to consumption growth.

Sizing the Virtual SAN Cluster for Failures to Tolerate

You configure the **Primary level of failures to tolerate** (PFTT) attribute in the VM storage policies to handle host failures. The number of hosts required for the cluster is calculated as follows: `2 * PFTT + 1`. The more failures the cluster is configured to tolerate, the more capacity hosts are required.

If the cluster hosts are connected in rack servers, you can organize the hosts into fault domains to improve failure management. See "Designing and Sizing Virtual SAN Fault Domains," on page 34.

Limitations of a Two-Host or Three-Host Cluster Configuration

In a two-host or three-host configuration, you can tolerate only one host failure by setting the **V of failures to tolerate** to 1. Virtual SAN saves each of the two required replicas of virtual machine data on separate hosts. The witness object is on a third host. Because of the small number of hosts in the cluster, the following limitations exist:

■ When a host fails, Virtual SAN cannot rebuild data on another host to protect against another failure.

■ If a host must enter maintenance mode, Virtual SAN cannot reprotect evacuated data. Data is exposed to a potential failure while the host is in maintenance mode.

You can use only the **Ensure data accessibility** data evacuation option. The **Evacuate all data** option is not available because the cluster does not have a spare host that it can use for evacuating data.

As a result, virtual machines are at risk because they become inaccessible if another failure occurs.

Balanced and Unbalanced Cluster Configuration

Virtual SAN works best on hosts with uniform configurations.

Using hosts with different configurations has the following disadvantages in a Virtual SAN cluster:

■ Reduced predictability of storage performance because Virtual SAN does not store the same number of components on each host.

■ Different maintenance procedures.

■ Reduced performance on hosts in the cluster that have smaller or different types of cache devices.

Deploying vCenter Server on Virtual SAN

If you deploy vCenter Server on the Virtual SAN datastore, you might not be able to use vCenter Server for troubleshooting, if a problem occurs in the Virtual SAN cluster.

Designing the Virtual SAN Network

Consider networking features that can provide availability, security, and bandwidth guarantee in a Virtual SAN cluster.

For details about the Virtual SAN network configuration, see the *VMware Virtual SAN Design and Sizing Guide* and *Virtual SAN Network Design Guide*.

Networking Failover and Load Balancing

Virtual SAN uses the teaming and failover policy that is configured on the backing virtual switch for network redundancy only. Virtual SAN does not use NIC teaming for load balancing.

If you plan to configure a NIC team for availability, consider these failover configurations.

Teaming Algorithm	Failover Configuration of the Adapters in the Team
Route based on originating virtual port	Active/Passive
Route based on IP hash	Active/Active with static EtherChannel for standard switch and LACP port channel for distributed switch
Route based on physical network adapter load	Active/Active

Virtual SAN supports IP-hash load balancing, but cannot guarantee improvement in performance for all configurations. You can benefit from IP hash when Virtual SAN is among its many consumers. In this case, IP hash performs load balancing. If Virtual SAN is the only consumer, you might notice no improvement. This behavior specifically applies to 1-GbE environments. For example, if you use four 1-GbE physical adapters with IP hash for Virtual SAN, you might not be able to use more than 1 Gbps. This behavior also applies to all NIC teaming policies that VMware supports.

Virtual SAN does not support multiple VMkernel adapters on the same subnet. You can use multiple VMkernel adapters on different subnets, such as another VLAN or separate physical fabric. Providing availability by using several VMkernel adapters has configuration costs including vSphere and the network infrastructure. Network availability by teaming physical network adapters is easier to achieve with less setup.

Using Unicast in Virtual SAN Network

In Virtual SAN 6.6 and later releases, multicast is not required on the physical switches that support the Virtual SAN cluster. You can design a simple unicast network for Virtual SAN. Earlier releases of Virtual SAN rely on multicast to enable heartbeat and to exchange metadata between hosts in the cluster. If some hosts in your Virtual SAN cluster are running earlier versions of software, a multicast network is still required. For more information about using multicast in a Virtual SAN cluster, refer to an earlier version of *Administering VMware Virtual SAN*.

NOTE The following configuration is not supported: vCenter Server deployed on a Virtual SAN 6.6 cluster that is using IP addresses from DHCP without reservations. You can use DHCP with reservations, because the assigned IP addresses are bound to the MAC addresses of VMkernel ports.

Allocating Bandwidth for Virtual SAN by Using Network I/O Control

If Virtual SAN traffic uses 10-GbE physical network adapters that are shared with other system traffic types, such as vSphere vMotion traffic, vSphere HA traffic, virtual machine traffic, and so on, you can use the vSphere Network I/O Control in vSphere Distributed Switch to guarantee the amount of bandwidth that is required for Virtual SAN.

In vSphere Network I/O Control, you can configure reservation and shares for the Virtual SAN outgoing traffic.

- Set a reservation so that Network I/O Control guarantees that minimum bandwidth is available on the physical adapter for Virtual SAN.

- Set shares so that when the physical adapter assigned for Virtual SAN becomes saturated, certain bandwidth is available to Virtual SAN and to prevent Virtual SAN from consuming the entire capacity of the physical adapter during rebuild and synchronization operations. For example, the physical adapter might become saturated when another physical adapter in the team fails and all traffic in the port group is transferred to the other adapters in the team.

For example, on a 10-GbE physical adapter that handles traffic for Virtual SAN, vSphere vMotion, and virtual machines, you can configure certain bandwidth and shares.

Table 4-3. Example Network I/O Control Configuration for a Physical Adapter That Handles Virtual SAN

Traffic Type	Reservation, Gbps	Shares
Virtual SAN	1	100
vSphere vMotion	0.5	70
Virtual machine	0.5	30

If the 10-GbE adapter becomes saturated, Network I/O Control allocates 5 Gbps to Virtual SAN on the physical adapter.

For information about using vSphere Network I/O Control to configure bandwidth allocation for Virtual SAN traffic, see the *vSphere Networking* documentation.

Marking Virtual SAN Traffic

Priority tagging is a mechanism to indicate to the connected network devices that Virtual SAN traffic has high Quality of Service (QoS) demands. You can assign Virtual SAN traffic to a certain class and mark the traffic accordingly with a Class of Service (CoS) value from 0 (low priority) to 7 (high priority), using the traffic filtering and marking policy of vSphere Distributed Switch.

Segmenting Virtual SAN Traffic in a VLAN

Consider isolating Virtual SAN traffic in a VLAN for enhanced security and performance, especially if you share the capacity of the backing physical adapter among several traffic types.

Jumbo Frames

If you plan to use jumbo frames with Virtual SAN to improve CPU performance, verify that jumbo frames are enabled on all network devices and hosts in the cluster.

By default, the TCP segmentation offload (TSO) and large receive offload (LRO) features are enabled on ESXi. Consider whether using jumbo frames improves the performance enough to justify the cost of enabling them on all nodes on the network.

Creating Static Routes for Virtual SAN Networking

You might need to create static routes in your Virtual SAN environment.

In traditional configurations, where vSphere uses a single default gateway, all routed traffic attempts to reach its destination through this gateway.

However, certain Virtual SAN deployments might require static routing. For example, deployments where the witness is on a different network, or the stretched cluster deployment, where both the data sites and the witness host are on different sites.

To configure static routing on your ESXi hosts, use the esxcli command:

```
esxcli network ip route ipv4 add -n remote-network -g gateway-to-use
```

remote-network is the remote network that your host must access, and *gateway-to-use* is the interface to use when traffic is sent to the remote network.

For more information, see

Best Practices for Virtual SAN Networking

Consider networking best practices for Virtual SAN to improve performance and throughput.

- For hybrid configurations, dedicate at least 1-GbE physical network adapter. Place Virtual SAN traffic on a dedicated or shared 10-GbE physical adapter for best networking performance.

- For all-flash configurations, use a dedicated or shared 10-GbE physical network adapter.

- Provision one additional physical NIC as a failover NIC.

- If you use a shared 10-GbE network adapter, place the Virtual SAN traffic on a distributed switch and configure Network I/O Control to guarantee bandwidth to Virtual SAN.

Designing and Sizing Virtual SAN Fault Domains

The Virtual SAN fault domains feature instructs Virtual SAN to spread redundancy components across the servers in separate computing racks. In this way, you can protect the environment from a rack-level failure such as loss of power or connectivity.

Fault Domain Constructs

Virtual SAN requires at least two fault domains, each of which consists of one or more hosts. Fault domain definitions must acknowledge physical hardware constructs that might represent a potential zone of failure, for example, an individual computing rack enclosure.

If possible, use at least four fault domains. Three fault domains do not support certain data evacuation modes, and Virtual SAN is unable to reprotect data after a failure. In this case, you need an additional fault domain with capacity for rebuilding, which you cannot provide with only three fault domains.

If fault domains are enabled, Virtual SAN applies the active virtual machine storage policy to the fault domains instead of the individual hosts.

Calculate the number of fault domains in a cluster based on the **Primary level of failures to tolerate** (PFTT) attribute from the storage policies that you plan to assign to virtual machines.

```
number of fault domains = 2 * PFTT + 1
```

If a host is not a member of a fault domain, Virtual SAN interprets it as a stand-alone fault domain.

Using Fault Domains Against Failures of Several Hosts

Consider a cluster that contains four server racks, each with two hosts. If the **Primary level of failures to tolerate** is set to one and fault domains are not enabled, Virtual SAN might store both replicas of an object with hosts in the same rack enclosure. In this way, applications might be exposed to a potential data loss on a rack-level failure. When you configure hosts that could potentially fail together into separate fault domains, Virtual SAN ensures that each protection component (replicas and witnesses) is placed in a separate fault domain.

If you add hosts and capacity, you can use the existing fault domain configuration or you can define fault domains.

For balanced storage load and fault tolerance when using fault domains, consider the following guidelines:

- Provide enough fault domains to satisfy the **Primary level of failures to tolerate** that are configured in the storage policies.

  Define at least three fault domains. Define a minimum of four domains for best protection.

- Assign the same number of hosts to each fault domain.

- Use hosts that have uniform configurations.

- Dedicate one fault domain of free capacity for rebuilding data after a failure, if possible.

Using Boot Devices and vSAN

Starting an ESXi installation that is a part of a vSAN cluster from a flash device imposes certain restrictions.

When you boot a vSAN host from a USB/SD device, you must use a high-quality USB or SD flash drive of 4 GB or larger.

When you boot a vSAN host from a SATADOM device, you must use single-level cell (SLC) device. The size of the boot device must be at least 16 GB.

During installation, the ESXi installer creates a coredump partition on the boot device. The default size of the coredump partition satisfies most installation requirements.

If the memory of the ESXi host has 512 GB of memory or less, you can boot the host from a USB, SD, or SATADOM device. If the memory of the ESXi host has more than 512 GB, you must boot the host from a SATADOM or disk device.

NOTE   vSAN 6.5 and later enables you to resize an existing coredump partition on an ESXi host in a vSAN cluster, and enables you to boot from USB/SD devices. For more information, see the VMware knowledge base article at http://kb.vmware.com/kb/2147881.

Hosts that boot from a disk have a local VMFS. If you have a disk with VMFS that runs VMs, you must separate the disk for an ESXi boot that is not for vSAN. In this case you need separate controllers.

Log Information and Boot Devices in vSAN

When you boot ESXi from a USB or SD device, log information and stack traces are lost on host reboot. They are lost because the scratch partition is on a RAM drive. Use persistent storage for logs, stack traces, and memory dumps.

Do not store log information on the vSAN datastore. This configuration is not supported because a failure in the vSAN cluster could impact the accessibility of log information.

Consider the following options for persistent log storage:

- Use a storage device that is not used for vSAN and is formatted with VMFS or NFS.

- Configure the ESXi Dump Collector and vSphere Syslog Collector on the host to send memory dumps and system logs to vCenter Server.

For information about setting up the scratch partition with a persistent location, see the *vSphere Installation and Setup* documentation.

Persistent Logging in a Virtual SAN Cluster

Provide storage for persistence of the logs from the hosts in the Virtual SAN cluster.

If you install ESXi on a USB or SD device and you allocate local storage to Virtual SAN, you might not have enough local storage or datastore space left for persistent logging.

To avoid potential loss of log information, configure the ESXi Dump Collector and vSphere Syslog Collector to redirect ESXi memory dumps and system logs to a network server. See the *vSphere Installation and Setup* documentation.

Preparing a New or Existing Cluster for Virtual SAN

5

Before you enable Virtual SAN on a cluster and start using it as virtual machine storage, provide the infrastructure that is required for correct operation of Virtual SAN.

This chapter includes the following topics:

Selecting or Verifying the Compatibility of Storage Devices

An important step before you deploy Virtual SAN is to verify that your storage devices, drivers, and firmware are compatible with Virtual SAN by consulting the *VMware Compatibility Guide*.

You can choose from several options for Virtual SAN compatibility.

- Use a Virtual SAN Ready Node server, a physical server that OEM vendors and VMware validate for Virtual SAN compatibility.

- Assemble a node by selecting individual components from validated device models.

VMware Compatibility Guide Section	Component Type for Verification
Systems	Physical server that runs ESXi.
Virtual SAN	- Magnetic disk SAS or SATA model for hybrid configurations. - Flash device model that is listed in the *VMware Compatibility Guide*. Certain models of PCIe flash devices can also work with Virtual SAN. Consider also write endurance and performance class. - Storage controller model that supports passthrough. Virtual SAN can work with storage controllers that are configured for RAID 0 mode if each storage device is represented as an individual RAID 0 group.

Preparing Storage

Provide enough disk space for Virtual SAN and for the virtualized workloads that use the Virtual SAN datastore.

Preparing Storage Devices

Use flash devices and magnetic disks based on the requirements for Virtual SAN.

Verify that the cluster has the capacity to accommodate anticipated virtual machine consumption and the **Primary level of failures to tolerate** in the storage policy for the virtual machines.

The storage devices must meet the following requirements so that Virtual SAN can claim them:

■   The storage devices are local to the ESXi hosts. Virtual SAN cannot claim remote devices.

■   The storage devices do not have any preexisting partition information.

■   On the same host, you cannot have both all-flash and hybrid disk groups.

Prepare Devices for Disk Groups

Each disk group provides one flash caching device and at least one magnetic disk or one flash capacity device. The capacity of the flash caching device must be at least 10 percent of the anticipated consumed storage on the capacity device, without the protection copies.

Virtual SAN requires at least one disk group on a host that contributes storage to a cluster that consists of at least three hosts. Use hosts that have uniform configuration for best performance of Virtual SAN.

Raw and Usable Capacity

Provide raw storage capacity that is greater than the capacity for virtual machines to handle certain cases.

■   Do not include the size of the flash caching devices as capacity. These devices do not contribute storage and are used as cache unless you have added flash devices for storage.

■   Provide enough space to handle the **Primary level of failures to tolerate** (PFTT) value in a virtual machine storage policy. A PFTT that is greater than 0 extends the device footprint. If the PFTT is set to 1, the footprint is double. If the PFTT is set to 2, the footprint is triple, and so on.

■   Verify whether the Virtual SAN datastore has enough space for an operation by examining the space on the individual hosts rather than on the consolidated Virtual SAN datastore object. For example, when you evacuate a host, all free space in the datastore might be on the host that you are evacuating and the cluster is not able to accommodate the evacuation to another host.

■   Provide enough space to prevent the datastore from running out of capacity, if workloads that have thinly provisioned storage start consuming a large amount of storage.

■   Verify that the physical storage can accommodate the reprotection and maintenance mode of the hosts in the Virtual SAN cluster.

■   Consider the Virtual SAN overhead to the usable storage space.
    ■   On-disk format version 1.0 adds an additional overhead of approximately 1 GB per capacity device.

    ■   On-disk format version 2.0 adds an additional overhead, typically no more than 1-2 percent capacity per device.

    ■   On-disk format version 3.0 and later adds an additional overhead, typically no more than 1-2 percent capacity per device. Deduplication and compression with software checksum enabled require additional overhead of approximately 6.2 percent capacity per device.

For more information about planning the capacity of Virtual SAN datastores, see the *VMware Virtual SAN Design and Sizing Guide*.

Virtual SAN Policy Impact on Capacity

The Virtual SAN storage policy for virtual machines affects the capacity devices in several ways.

Table 5-1. Virtual SAN VM Policy and Raw Capacity

Aspects of Policy Influence	Description
Policy changes	■ The **Primary level of failures to tolerate** (PFTT) influences the physical storage space that you must supply for virtual machines. The greater the PFTT is for higher availability, the more space you must provide. When PFTT is set to 1, it imposes two replicas of the VMDK file of a virtual machine. With PFTT set to 1, a VMDK file that is 50 GB requires 100 GB space on different hosts. If the PFTT is changed to 2, you must have enough space to support three replicas of the VMDK across the hosts in the cluster, or 150 GB. ■ Some policy changes, such as a new number of disk stripes per object, require temporary resources. Virtual SAN recreates the new objects that are affected by the change and for a certain time, the physical storage must accommodate the old and new objects.
Available space for reprotecting or maintenance mode	When you place a host in maintenance mode or you clone a virtual machine, although the Virtual SAN datastore indicates that enough space is available, the datastore might not be able to evacuate the virtual machine objects because the free space is on the host that is being placed in maintenance mode.

Mark Flash Devices as Capacity Using ESXCLI

You can manually mark the flash devices on each host as capacity devices using esxcli.

Prerequisites

Verify that you are using Virtual SAN 6.5 or later.

Procedure

1 To learn the name of the flash device that you want to mark as capacity, run the following command on each host.

   a In the ESXi Shell, run the `esxcli storage core device list` command.

   b Locate the device name at the top of the command output and write the name down.

   The command takes the following options:

Table 5-2. Command Options

Options	Description
–d|––disk=str	The name of the device that you want to tag as a capacity device. For example, `mpx.vmhba1:C0:T4:L0`
–t|––tag=str	Specify the tag that you want to add or remove. For example, the `capacityFlash` tag is used for marking a flash device for capacity.

   The command lists all device information identified by ESXi.

2 In the output, verify that the `Is SSD` attribute for the device is `true`.

3   To tag a flash device as capacity, run the `esxcli vsan storage tag add –d <device name> –t capacityFlash` command.

For example, the `esxcli vsan storage tag add –t capacityFlash –d mpx.vmhba1:C0:T4:L0` command, where `mpx.vmhba1:C0:T4:L0` is the device name.

4   Verify whether the flash device is marked as capacity.

a   In the output, identify whether the `IsCapacityFlash` attribute for the device is set to `1`.

Example: Command Output

You can run the `vdq –q –d <device name>` command to verify the `IsCapacityFlash` attribute. For example, running the `vdq –q –d mpx.vmhba1:C0:T4:L0` command, returns the following output.

```
\{
"Name"     : "mpx.vmhba1:C0:T4:L0",
"VSANUUID" : "",
"State"    : "Eligible for use by VSAN",
"ChecksumSupport": "0",
"Reason"   : "None",
"IsSSD"    : "1",
"IsCapacityFlash": "1",
"IsPDL"    : "0",
    \},
```

Untag Flash Devices Used as Capacity Using ESXCLI

You can untag flash devices that are used as capacity devices, so that they are available for caching.

Procedure

1   To untag a flash device marked as capacity, run the `esxcli vsan storage tag remove –d <device name> –t capacityFlash` command. For example, the `esxcli vsan storage tag remove –t capacityFlash –d mpx.vmhba1:C0:T4:L0` command, where `mpx.vmhba1:C0:T4:L0` is the device name.

2   Verify whether the flash device is untagged.

a   In the output, identify whether the `IsCapacityFlash` attribute for the device is set to `0`.

Example: Command Output

You can run the `vdq –q –d <device name>` command to verify the `IsCapacityFlash` attribute. For example, running the `vdq –q –d mpx.vmhba1:C0:T4:L0` command, returns the following output.

```
[
    \{
"Name"     : "mpx.vmhba1:C0:T4:L0",
"VSANUUID" : "",
"State"    : "Eligible for use by VSAN",
"ChecksumSupport": "0",
"Reason"   : "None",
"IsSSD"    : "1",
"IsCapacityFlash": "0",
"IsPDL"    : "0",
    \},
```

Mark Flash Devices as Capacity using RVC

Run the `vsan.host_claim_disks_differently` RVC command to mark storage devices as flash, capacity flash, or magnetic disk (HDD).

You can use the RVC tool to tag flash devices as capacity devices either individually, or in a batch by specifying the model of the device. When you want to tag flash devices as capacity devices you can include them in all-flash disk groups.

NOTE The `vsan.host_claim_disks_differently` command does not check the device type before tagging them. The command tags any device that you append with the `capacity_flash` command option, including the magnetic disks and devices that are already in use. Make sure you verify the device status before tagging.

For information about the RVC commands for Virtual SAN management, see the *RVC Command Reference Guide*.

Prerequisites

- Verify that you are using Virtual SAN version 6.5 or later.

- Verify that SSH is enabled on the vCenter Server Appliance.

Procedure

1 Open an SSH connection to the vCenter Server Appliance.

2 Log into the appliance by using a local account that has administrator privilege.

3 Start the RVC by running the following command.

   `rvc local_user_name@target_vCenter_Server`

   For example, to use the same vCenter Server Appliance to mark flash devices for capacity as a user root, run the following command:

   `rvc root@localhost`

4 Enter the password for the user name.

5 Navigate to the *vcenter_server*/*data_center*/computers/*cluster*/hosts directory in the vSphere infrastructure.

6 Run the `vsan.host_claim_disks_differently` command with the `--claim-type capacity_flash --model` *model_name* options to mark all flash devices of the same model as capacity on all hosts in the cluster.

   `vsan.host_claim_disks_differently --claim-type capacity_flash --model model_name *`

What to do next

Enable Virtual SAN on the cluster and claim capacity devices.

Providing Memory for Virtual SAN

You must provision hosts with memory according to the maximum number of devices and disk groups that you intend to map to Virtual SAN.

To satisfy the case of the maximum number of devices and disk groups, you must provision hosts with 32 GB of memory for system operations. For information about the maximum device configuration, see the *vSphere Configuration Maximums* documentation.

Preparing Your Hosts for Virtual SAN

As a part of the preparation for enabling Virtual SAN, review the requirements and recommendations about the configuration of hosts for the cluster.

- Verify that the storage devices on the hosts, and the driver and firmware versions for them, are listed in the Virtual SAN section of the *VMware Compatibility Guide*.

- Make sure that a minimum of three hosts contribute storage to the Virtual SAN datastore.

- For maintenance and remediation operations on failure, add at least four hosts to the cluster.

- Designate hosts that have uniform configuration for best storage balance in the cluster.

- Do not add hosts that have only compute resources to the cluster to avoid unbalanced distribution of storage components on the hosts that contribute storage. Virtual machines that require a lot of storage space and run on compute-only hosts might store a great number of components on individual capacity hosts. As a result, the storage performance in the cluster might be lower.

- Do not configure aggressive CPU power management policies on the hosts for saving power. Certain applications that are sensitive to CPU speed latency might have very low performance. For information about CPU power management policies, see the *vSphere Resource Management* documentation.

- If your cluster contains blade servers, consider extending the capacity of the datastore with an external storage enclose that is connected to the blade servers and is listed in the Virtual SAN section of the *VMware Compatibility Guide*.

- Consider the configuration of the workloads that you place on a hybrid or all-flash disk configuration.

  - For high levels of predictable performance, provide a cluster of all-flash disk groups.

  - For balance between performance and cost, provide a cluster of hybrid disk groups.

Virtual SAN and vCenter Server Compatibility

Synchronize the versions of vCenter Server and of ESXi to avoid potential faults because of differences in the Virtual SAN support in vCenter Server and ESXi.

For best integration between Virtual SAN components on vCenter Server and ESXi, deploy the latest version of the two vSphere components. See the *vSphere Installation and Setup* and *vSphere Upgrade* documentation.

Preparing Storage Controllers

Configure the storage controller on a host according to the requirements of Virtual SAN.

Verify that the storage controllers on the Virtual SAN hosts satisfy certain requirements for mode, driver, and firmware version, queue depth, caching and advanced features.

Table 5-3. Examining Storage Controller Configuration for Virtual SAN

Storage Controller Feature	Storage Controller Requirement
Required mode	■ Review the Virtual SAN requirements in the *VMware Compatibility Guide* for the required mode, passthrough or RAID 0, of the controller. ■ If both passthrough and RAID 0 modes are supported, configure passthrough mode instead of RAID0. RAID 0 introduces complexity for disk replacement.
RAID mode	■ In the case of RAID 0, create one RAID volume per physical disk device. ■ Do not enable a RAID mode other than the mode listed in the *VMware Compatibility Guide*. ■ Do not enable controller spanning.
Driver and firmware version	■ Use the latest driver and firmware version for the controller according to *VMware Compatibility Guide*. ■ If you use the in-box controller driver, verify that the driver is certified for Virtual SAN. OEM ESXi releases might contain drivers that are not certified and listed in the *VMware Compatibility Guide*.
Queue depth	Verify that the queue depth of the controller is 256 or higher. Higher queue depth provides improved performance.
Cache	Disable the storage controller cache, or set it to 100 percent read if disabling cache is not possible.
Advanced features	Disable advanced features, for example, HP SSD Smart Path.

Configuring Virtual SAN Network

Before you enable Virtual SAN on a cluster and ESXi hosts, you must construct the necessary network to carry the Virtual SAN communication.

Virtual SAN provides a distributed storage solution, which implies exchanging data across the ESXi hosts that participate in the cluster. Preparing the network for installing Virtual SAN includes certain configuration aspects.

For information about network design guidelines, see

Placing Hosts in the Same Subnet

Hosts must be connected in the same subnet for best networking performance. In Virtual SAN 6.0 and later, you can also connect hosts in the same Layer 3 network if required.

Dedicating Network Bandwidth on a Physical Adapter

Allocate at least 1 Gbps bandwidth for Virtual SAN. You might use one of the following configuration options:

■ Dedicate 1-GbE physical adapters for a hybrid host configuration.

■ Use dedicated or shared 10-GbE physical adapters for all-flash configurations.

■ Use dedicated or shared 10-GbE physical adapters for hybrid configurations if possible.

■ Direct Virtual SAN traffic on a 10-GbE physical adapter that handles other system traffic and use vSphere Network I/O Control on a distributed switch to reserve bandwidth for Virtual SAN.

Configuring a Port Group on a Virtual Switch

Configure a port group on a virtual switch for Virtual SAN.

- Assign the physical adapter for Virtual SAN to the port group as an active uplink.

    In the case of a NIC team for network availability, select a teaming algorithm based on the connection of the physical adapters to the switch.

- If designed, assign Virtual SAN traffic to a VLAN by enabling tagging in the virtual switch.

Examining the Firewall on a Host for Virtual SAN

Virtual SAN sends messages on certain ports on each host in the cluster. Verify that the host firewalls allow traffic on these ports.

Table 5-4. Ports on the Hosts in Virtual SAN

Virtual SAN Service	Traffic Direction	Communicating Nodes	Transport Protocol	Port
Virtual SAN Vendor Provider (vsanvp)	Incoming and outgoing	vCenter Server and ESXi	TCP	8080
Virtual SAN Clustering Service		ESXi	UDP	12345, 23451
Virtual SAN Transport		ESXi	TCP	2233
Unicast agent		ESXi	UDP	12321

Considerations about the Virtual SAN License

When you prepare your cluster for Virtual SAN, review the requirements of the Virtual SAN license.

- Make sure that you obtained a valid license for full host configuration control in the cluster. The license should be different from the one that you used for evaluation purposes.

    After the license or the evaluation period of a Virtual SAN expires, you can continue to use the current configuration of Virtual SAN resources. However, you cannot add capacity to a disk group or create disk groups.

- If the cluster consists of all-flash disk groups, verify that the all-flash feature is available under your license.

- If the Virtual SAN cluster uses advanced features such as deduplication and compression or stretched cluster, verify that the feature is available under your license.

- Consider the CPU capacity of the Virtual SAN license across the cluster when adding and removing hosts to the cluster.

    Virtual SAN licenses have per CPU capacity. When you assign a Virtual SAN license to a cluster, the amount of license capacity that is used equals the total number of CPUs on the hosts that participate in the cluster.

Creating a Virtual SAN Cluster

6

You can activate Virtual SAN when you create a cluster or enable Virtual SAN on your existing clusters.

This chapter includes the following topics:

Characteristics of a Virtual SAN Cluster

Before working on a Virtual SAN environment, you should be aware of the characteristics of a Virtual SAN cluster.

A Virtual SAN cluster includes the following characteristics:

- You can have multiple Virtual SAN clusters for each vCenter Server instance. You can use a single vCenter Server to manage more than one Virtual SAN cluster.

- Virtual SAN consumes all devices, including flash cache and capacity devices, and does not share devices with other features.

- Virtual SAN clusters can include hosts with or without capacity devices. The minimum requirement is three hosts with capacity devices. For best result, create a Virtual SAN cluster with uniformly configured hosts.

- If a host contributes capacity, it must have at least one flash cache device and one capacity device.

- In hybrid clusters, the magnetic disks are used for capacity and flash devices for read and write cache. Virtual SAN allocates 70 percent of all available cache for read cache and 30 percent of available cache for write buffer. In this configurations, the flash devices serve as a read cache and a write buffer.

- In all-flash cluster, one designated flash device is used as a write cache, additional flash devices are used for capacity. In all-flash clusters, all read requests come directly from the flash pool capacity.

- Only local or direct-attached capacity devices can participate in a Virtual SAN cluster. Virtual SAN cannot consume other external storage, such as SAN or NAS, attached to cluster.

For best practices about designing and sizing a Virtual SAN cluster, see Chapter 4, "Designing and Sizing a Virtual SAN Cluster," on page 23.

Before Creating a Virtual SAN Cluster

This topic provides a checklist of software and hardware requirements for creating a Virtual SAN cluster. You can also use the checklist to verify that the cluster meets the guidelines and basic requirements.

Requirements for Virtual SAN Cluster

Before you get started, verify specific models of hardware devices, and specific versions of drivers and firmware in the VMware Compatibility Guide Web site at http://www.vmware.com/resources/compatibility/search.php. The following table lists the key software and hardware requirements supported by Virtual SAN.

⚠️ **Caution**   Using uncertified software and hardware components, drivers, controllers, and firmware might cause unexpected data loss and performance issues.

Table 6-1. Virtual SAN Cluster Requirements

Requirements	Description
ESXi Hosts	■ Verify that you are using the latest version of ESXi on your hosts. ■ Verify that there are at least three ESXi hosts with supported storage configurations available to be assigned to the Virtual SAN cluster. For best results, configure the Virtual SAN cluster with four or more hosts.
Memory	■ Verify that each host has a minimum of 8 GB of memory. ■ For larger configurations and better performance, you must have a minimum of 32 GB of memory in the cluster. See "Designing and Sizing vSAN Hosts," on page 29.
Storage I/O controllers, drivers, firmware	■ Verify that the storage I/O controllers, drivers, and firmware versions are certified and listed in the VCG Web site at http://www.vmware.com/resources/compatibility/search.php. ■ Verify that the controller is configured for passthrough or RAID 0 mode. ■ Verify that the controller cache and advanced features are disabled. If you cannot disable the cache, you must set the read cache to 100 percent. ■ Verify that you are using controllers with higher queue depths. Using controllers with queue depths less than 256 can significantly impact the performance of your virtual machines during maintenance and failure.
Cache and capacity	■ Verify that Virtual SAN hosts contributing storage to the cluster must have at least one cache and one capacity device. Virtual SAN requires exclusive access to the local cache and capacity devices of the hosts in the Virtual SAN cluster. They cannot share these devices with other uses, such as Virtual Flash File System (VFFS), VMFS partitions, or an ESXi boot partition. ■ For best results, create a Virtual SAN cluster with uniformly configured hosts.
Network connectivity	■ Verify that each host is configured with at least one network adapter. ■ For hybrid configurations, verify that Virtual SAN hosts have a minimum dedicated bandwidth of 1 GbE. ■ For all-flash configurations, verify that Virtual SAN hosts have a minimum bandwidth of 10 GbE. For best practices and considerations about designing the Virtual SAN network, see "Designing the Virtual SAN Network," on page 32 and "Networking Requirements for Virtual SAN," on page 21.

Table 6-1. Virtual SAN Cluster Requirements (Continued)

Requirements	Description
Virtual SAN and vCenter Server Compatibility	Verify that you are using the latest version of the vCenter Server.
License key	■ Verify that you have a valid Virtual SAN license key. ■ To use the all-flash feature, your license must support that capability. ■ To use advanced features, such as stretched clusters or deduplication and compression, your license must support those features. ■ Verify that the amount of license capacity that you plan on using equals the total number of CPUs in the hosts participating in the Virtual SAN cluster. Do not provide license capacity only for hosts providing capacity to the cluster. For information about licensing for Virtual SAN, see the *vCenter Server and Host Management* documentation.

For detailed information about Virtual SAN Cluster requirements, see Chapter 3, "Requirements for Enabling Virtual SAN," on page 19.

For in-depth information about designing and sizing the Virtual SAN cluster, see the *VMware Virtual SAN Design and Sizing Guide*.

Enabling Virtual SAN

To use Virtual SAN, you must create a host cluster and enable Virtual SAN on the cluster.

A Virtual SAN cluster can include hosts with capacity and hosts without capacity. Follow these guidelines when you create a Virtual SAN cluster.

■ A Virtual SAN cluster must include a minimum of three ESXi hosts. For a Virtual SAN cluster to tolerate host and device failures, at least three hosts that join the Virtual SAN cluster must contribute capacity to the cluster. For best results, consider adding four or more hosts contributing capacity to the cluster.

■ Only ESXi 5.5 Update 1 or later hosts can join the Virtual SAN cluster.

■ All hosts in the Virtual SAN cluster must have the same on-disk format.

■ Before you move a host from a Virtual SAN cluster to another cluster, make sure that the destination cluster is Virtual SAN enabled.

■ To be able to access the Virtual SAN datastore, an ESXi host must be a member of the Virtual SAN cluster.

After you enable Virtual SAN, the Virtual SAN storage provider is automatically registered with vCenter Server and the Virtual SAN datastore is created. For information about storage providers, see the *vSphere Storage* documentation.

Set Up a VMkernel Network for Virtual SAN

To enable the exchange of data in the Virtual SAN cluster, you must provide a VMkernel network adapter for Virtual SAN traffic on each ESXi host.

Procedure

1   In the vSphere Web Client, navigate to the host.

2   Click the **Configure** tab.

3   Under **Networking**, select **VMkernel adapters**.

4   Click the **Add host networking** icon ( ) to open the Add Networking wizard.

5   On the Select connection type page, select **VMkernel Network Adapter** and click **Next**.

6   Configure the target switching device.

7   On the Port properties page, select **vSAN traffic**.

8   Complete the VMkernel adapter configuration.

9   On the Ready to complete page, verify that Virtual SAN is Enabled in the status for the VMkernel adapter, and click **Finish**.

Virtual SAN network is enabled for the host.

What to do next

You can enable Virtual SAN on the host cluster.

Create a Virtual SAN Cluster

You can enable Virtual SAN when you create a cluster.

Procedure

1   Right-click a data center in the vSphere Web Client and select **New Cluster**.

2   Type a name for the cluster in the **Name** text box.

    This name appears in the vSphere Web Client navigator.

3   Select the vSAN **Turn ON** check box and click **OK**.

    The cluster appears in the inventory.

4   Add hosts to the Virtual SAN cluster. See "Add a Host to the Virtual SAN Cluster," on page 108.

    Virtual SAN clusters can include hosts with or without capacity devices. For best results, add hosts with capacity.

Enabling Virtual SAN creates a Virtual SAN datastore and registers the Virtual SAN storage provider. Virtual SAN storage providers are built-in software components that communicate the storage capabilities of the datastore to vCenter Server.

What to do next

Verify that the Virtual SAN datastore has been created. See "View Virtual SAN Datastore," on page 53.

Verify that the Virtual SAN storage provider is registered. See "View Virtual SAN Storage Providers," on page 126.

Claim the storage devices or create disk groups. See Chapter 11, "Device Management in a Virtual SAN Cluster," on page 99.

Configure the Virtual SAN cluster. See "Configure a Cluster for Virtual SAN," on page 49.

Configure a Cluster for Virtual SAN

You can use the Configure Virtual SAN wizard to complete the basic configuration of your Virtual SAN cluster.
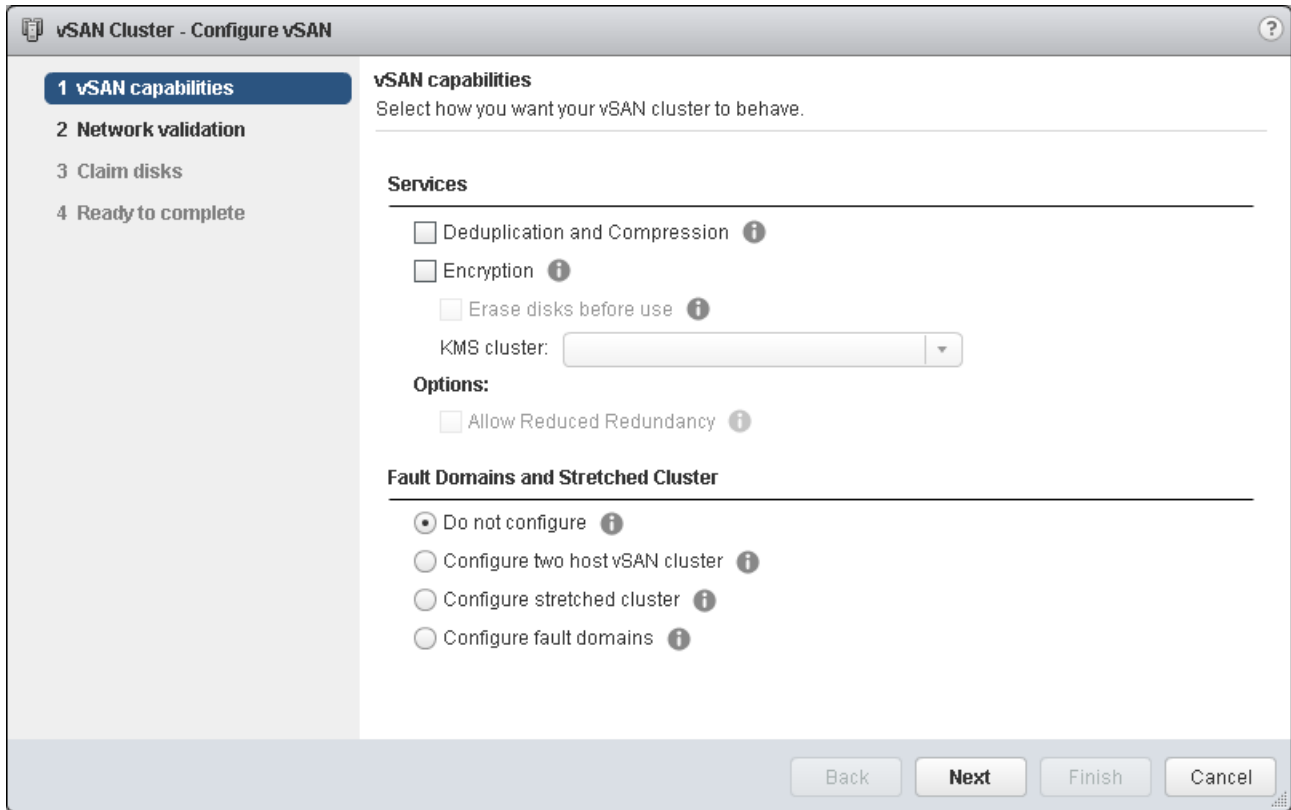
Prerequisites

You must create a cluster and add hosts to the cluster before using the Configure Virtual SAN wizard to complete the basic configuration.

Procedure

1　Navigate to an existing cluster in the vSphere Web Client.

2　Click the **Configure** tab.

3　Under vSAN, select **General** and click the **Configure** button.

4　Select **vSAN capabilities**.

　a　(Optional) Select the **Deduplication and Compression** check box if you want to enable deduplication and compression on the cluster.

　　You can select the **Allow Reduced Redundancy** check box to enable deduplication and compression on a Virtual SAN cluster that has limited resources, such as a three-host cluster with the **Primary level of failures to tolerate** set to 1. If you allow reduced redundancy, your data might be at risk during the disk reformat operation.

　b　(Optional) Select the **Encryption** check box if you want to enable data at rest encryption, and select a KMS.

c   Select the fault tolerance mode for the cluster.

Option	Description
Do not configure	Default setting used for a single-site Virtual SAN cluster.
2 host virtual SAN cluster	Provides fault tolerance for a cluster that has two hosts at a remote office, with a witness host in the main office. Set the **Primary level of failures to tolerate** policy to 1.
Stretched cluster	Supports two active sites, each with an even number of hosts and storage devices, and a witness host at a third site.
Configure fault domains	Supports fault domains that you can use to group Virtual SAN hosts that might fail together. Assign one or more hosts to each fault domain.

d   You can select the **Allow Reduced Redundancy** check box to enable encryption or deduplication and compression on a Virtual SAN cluster that has limited resources. For example, if you have a three-host cluster with the **Primary level of failures to tolerate** set to 1. If you allow reduced redundancy, your data might be at risk during the disk reformat operation.

5   Click **Next**.

6   On the Network validation page, check the settings for Virtual SAN VMkernel adapters, and click **Next**.

7   On the Claim disks page, select the disks for use by the cluster and click **Next**.

For each host that contributes storage, select one flash device for the cache tier, and one or more devices for the capacity tier.

8   Follow the wizard to complete the configuration of the cluster, based on the fault tolerance mode.

a   If you selected **Configure two host vSAN cluster**, choose a witness host for the cluster, and claim disks for the witness host.

b   If you selected **Configure stretched cluster**, define fault domains for the cluster, choose a witness host, and claim disks for the witness host.

c   If you selected **Configure fault domains**, define fault domains for the cluster.

For more information about fault domains, see "Managing Fault Domains in Virtual SAN Clusters," on page 113.

For more information about stretched clusters, see Chapter 7, "Extending a Datastore Across Two Sites with Stretched Clusters," on page 61.

9   On the Ready to complete page, review the configuration, and click **Finish**.

Edit Virtual SAN Settings

You can edit the settings of your Virtual SAN cluster to change the method for claiming disks and to enable deduplication and compression.

Edit the settings of an existing Virtual SAN cluster if you want to enable deduplication and compression, or to enable encryption. If you enable deduplication and compression, or if you enable encryption, the on-disk format of the cluster is automatically upgraded to the latest version.

Procedure

1   Navigate to the Virtual SAN host cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General**.

4   Click the Virtual SAN is turned ON **Edit** button.

5   (Optional) If you want to enable deduplication and compression on the cluster, select the Deduplication and compression check box.

Virtual SAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.

6   (Optional) If you want to enable encryption on the cluster, select theEncryption check box, and select a KMS server.

Virtual SAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.

7   Click **OK**.

Enable Virtual SAN on an Existing Cluster

You can edit cluster properties to enable Virtual SAN for an existing cluster.

After enabling Virtual SAN on your cluster, you cannot move Virtual SAN hosts from a Virtual SAN enabled cluster to a non-Virtual SAN cluster.

Prerequisites

Verify that your environment meets all requirements. See Chapter 3, "Requirements for Enabling Virtual SAN," on page 19.

Procedure

1   Navigate to an existing host cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General** and click **Edit** to edit the cluster settings.

4   If you want to enable deduplication and compression on the cluster, select the Deduplication and compression check box.

Virtual SAN automatically upgrades the on-disk format, causing a rolling reformat of every disk group in the cluster.

5   (Optional) If you want to enable encryption on the cluster, select theEncryption check box, and select a KMS server.

Virtual SAN automatically upgrades the on-disk format, causing a rolling reformat of every disk group in the cluster.

6   Click **OK**.

What to do next

Claim the storage devices or create disk groups. See Chapter 11, "Device Management in a Virtual SAN Cluster," on page 99.

Disable Virtual SAN

You can turn off Virtual SAN for a host cluster.

When you disable the Virtual SAN cluster, all virtual machines located on the shared Virtual SAN datastore become inaccessible. If you intend to use virtual machine while Virtual SAN is disabled, make sure you migrate virtual machines from Virtual SAN datastore to another datastore before disabling the Virtual SAN cluster.

Prerequisites

Verify that the hosts are in maintenance mode.

Procedure

1   Navigate to the host cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General** and click **Edit** to edit Virtual SAN settings.

4   Deselect the vSAN **Turn On** check box.

Configure License Settings for a Virtual SAN Cluster

You must assign a license to a Virtual SAN cluster before its evaluation period expires or its currently assigned license expires.

If you upgrade, combine, or divide Virtual SAN licenses, you must assign the new licenses to Virtual SAN clusters. When you assign a Virtual SAN license to a cluster, the amount of license capacity that is used equals the total number of CPUs in the hosts participating in the cluster. The license usage of the Virtual SAN cluster is recalculated and updated every time you add or remove a host from the cluster. For information about managing licenses and licensing terminology and definitions, see the *vCenter Server and Host Management* documentation.

When you enable Virtual SAN on a cluster, you can use Virtual SAN in evaluation mode to explore its features. The evaluation period starts when Virtual SAN is enabled, and expires after 60 days. To use Virtual SAN, you must license the cluster before the evaluation period expires. Just like vSphere licenses, Virtual SAN licenses have per CPU capacity. Some advanced features, such as all-flash configuration and stretched clusters, require a license that supports the feature.
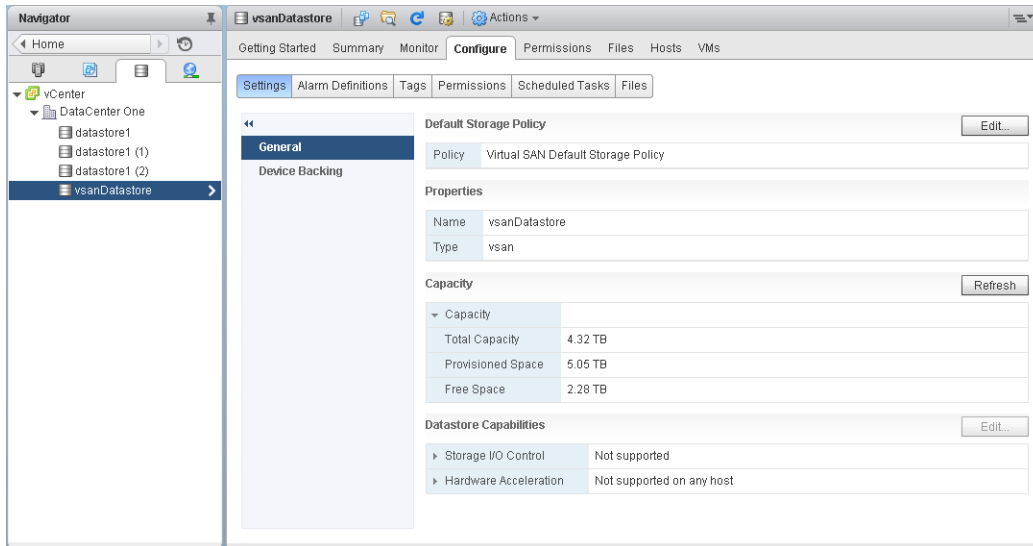
Prerequisites

■   To view and manage Virtual SAN licenses, you must have the **Global.Licenses** privilege on the vCenter Server systems, where the vSphere Web Client runs.

Procedure

1   In the vSphere Web Client, navigate to a cluster where you have enabled Virtual SAN.

2   Click the **Configure** tab.

3   Under **Configuration**, select **Licensing**, and click **Assign License**.

4   Select a licensing option.

   ■   Select an existing license and click **OK**.

   ■   Create a new Virtual SAN license.

      a   Click the **Create New License** (✚) icon.

      b   In the New Licenses dialog box, type or copy and paste a Virtual SAN license key and click **Next**.

      c   On the Edit license names page, rename the new license as appropriate and click **Next**.

      d   Click **Finish**.

      e   In the Assign License dialog, select the newly created license and click **OK**.

View Virtual SAN Datastore

After you enable Virtual SAN, a single datastore is created. You can review the capacity of the Virtual SAN datastore.

Prerequisites

Activate Virtual SAN and configure disk groups.

Procedure

1 Navigate to Storage in the vSphere Web Client.

2 Select the Virtual SAN datastore.

3 Click the **Configure** tab.

4 Review the Virtual SAN datastore capacity.

The size of the Virtual SAN datastore depends on the number of capacity devices per ESXi host and the number of ESXi hosts in the cluster. For example, if a host has seven 2 TB for capacity devices, and the cluster includes eight hosts, the approximate storage capacity would be 7 x 2 TB x 8 = 112 TB. Note that when using the all-flash configuration, flash devices are used for capacity. For hybrid configuration, magnetic disks are used for capacity.

Some capacity is allocated for metadata.

■ On-disk format version 1.0 adds approximately 1 GB per capacity device.

■ On-disk format version 2.0 adds capacity overhead, typically no more than 1-2 percent capacity per device.

■ On-disk format version 3.0 and later adds capacity overhead, typically no more than 1-2 percent capacity per device. Deduplication and compression with software checksum enabled require additional overhead of approximately 6.2 percent capacity per device.

What to do next

Use the storage capabilities of the Virtual SAN datastore to create a storage policy for virtual machines. For information, see the *vSphere Storage* documentation.

Using Virtual SAN and vSphere HA

You can enable vSphere HA and Virtual SAN on the same cluster. As with traditional datastores, vSphere HA provides the same level of protection for virtual machines on Virtual SAN datastores. This level of protection imposes specific restrictions when vSphere HA and Virtual SAN interact.

ESXi Host Requirements

You can use Virtual SAN with a vSphere HA cluster only if the following conditions are met:

■ The cluster's ESXi hosts all must be version 5.5 Update 1 or later.

■ The cluster must have a minimum of three ESXi hosts. For best results, configure the Virtual SAN cluster with four or more hosts.

Networking Differences

Virtual SAN uses its own logical network. When Virtual SAN and vSphere HA are enabled for the same cluster, the HA interagent traffic flows over this storage network rather than the management network. vSphere HA uses the management network only when Virtual SAN is disabled.vCenter Server chooses the appropriate network when vSphere HA is configured on a host.

NOTE You must disable vSphere HA before you enable Virtual SAN on the cluster. Then you can reenable vSphere HA.

When a virtual machine is only partially accessible in all network partitions, you cannot power on the virtual machine or fully access it in any partition. For example, if you partition a cluster into P1 and P2, the VM namespace object is accessible to the partition named P1 and not to P2. The VMDK is accessible to the partition named P2 and not to P1. In such cases, the virtual machine cannot be powered on and it is not fully accessible in any partition.

The following table shows the differences in vSphere HA networking whether or not Virtual SAN is used.

Table 6-2. vSphere HA Networking Differences

	Virtual SAN Enabled	Virtual SAN Disabled
Network used by vSphere HA	Virtual SAN storage network	Management network
Heartbeat datastores	Any datastore mounted to more than one host, but not Virtual SAN datastores	Any datastore mounted to more than one host
Host declared isolated	Isolation addresses not pingable and Virtual SAN storage network inaccessible	Isolation addresses not pingable and management network inaccessible

If you change the Virtual SAN network configuration, the vSphere HA agents do not automatically acquire the new network settings. To make changes to the Virtual SAN network, you must reenable host monitoring for the vSphere HA cluster by using vSphere Web Client:

1 Disable Host Monitoring for the vSphere HA cluster.

2 Make the Virtual SAN network changes.

3 Right-click all hosts in the cluster and select **Reconfigure HA**.

4 Reenable Host Monitoring for the vSphere HA cluster.

Capacity Reservation Settings

When you reserve capacity for your vSphere HA cluster with an admission control policy, this setting must be coordinated with the corresponding **Primary level of failures to tolerate** policy setting in the Virtual SAN rule set and must not be lower than the capacity reserved by the vSphere HA admission control setting. For example, if the Virtual SAN rule set allows for only two failures, the vSphere HA admission control policy must reserve capacity that is equivalent to only one or two host failures. If you are using the Percentage of Cluster Resources Reserved policy for a cluster that has eight hosts, you must not reserve more than 25 percent of the cluster resources. In the same cluster, with the **Primary level of failures to tolerate** policy, the setting must not be higher than two hosts. If vSphere HA reserves less capacity, failover activity might be unpredictable. Reserving too much capacity overly constrains the powering on of virtual machines and intercluster vSphere vMotion migrations. For information about the Percentage of Cluster Resources Reserved policy, see the *vSphere Availability* documentation.

Virtual SAN and vSphere HA Behavior in a Multiple Host Failure

After a Virtual SAN cluster fails with a loss of failover quorum for a virtual machine object, vSphere HA might not be able to restart the virtual machine even when the cluster quorum has been restored. vSphere HA guarantees the restart only when it has a cluster quorum and can access the most recent copy of the virtual machine object. The most recent copy is the last copy to be written.

Consider an example where a Virtual SAN virtual machine is provisioned to tolerate one host failure. The virtual machine runs on a Virtual SAN cluster that includes three hosts, H1, H2, and H3. All three hosts fail in a sequence, with H3 being the last host to fail.

After H1 and H2 recover, the cluster has a quorum (one host failure tolerated). Despite this quorum, vSphere HA is unable to restart the virtual machine because the last host that failed (H3) contains the most recent copy of the virtual machine object and is still inaccessible.

In this example, either all three hosts must recover at the same time, or the two-host quorum must include H3. If neither condition is met, HA attempts to restart the virtual machine when host H3 is online again.

Deploying Virtual SAN with vCenter Server Appliance

You can create a Virtual SAN cluster as you deploy a vCenter Server Appliance, and host the appliance on that cluster.

The vCenter Server Appliance is a preconfigured Linux virtual machine, which is used for running VMware vCenter Server on Linux systems. This feature enables you to configure a Virtual SAN cluster on new ESXi hosts without using vCenter Server.

When you use the vCenter Server Appliance Installer to deploy a vCenter Server Appliance, you can create a single-host Virtual SAN cluster, and host the vCenter Server Appliance on the cluster. During Stage 1 of the deployment, when you select a datastore, click **Install on a new Virtual SAN cluster containing the target host**. Follow the steps in the Installer wizard to complete the deployment.

The vCenter Server Appliance Installer creates a one-host Virtual SAN cluster, with disks claimed from the host. vCenter Server Appliance is deployed on the Virtual SAN cluster.

After you complete the deployment, you can manage the single-host Virtual SAN cluster with the vCenter Server Appliance. You must complete the configuration of the Virtual SAN cluster.

You can deploy a Platform Services Controller and vCenter Server on the same Virtual SAN cluster or on separate clusters.

- You can deploy a Platform Services Controller and vCenter Server to the same Virtual SAN cluster. Deploy the PSC and vCenter Server to the same single-host virtual SAN datastore. After you complete the deployment, the Platform Services Controller and vCenter Server both run on the same cluster.

■ You can deploy a Platform Services Controller and vCenter Server to different Virtual SAN clusters. Deploy the Platform Services Controller and vCenter Server to separate single-host Virtual SAN clusters. After you complete the deployment, you must complete the configuration of each Virtual SAN cluster separately.

Using Virtual SAN Configuration Assist and Updates

You can use Configuration Assist to check the configuration of your Virtual SAN cluster, and resolve any issues.

Virtual SAN Configuration Assist enables you to verify the configuration of cluster components, resolve issues, and troubleshoot problems. The configuration checks cover hardware compatibility, network, and Virtual SAN configuration options.
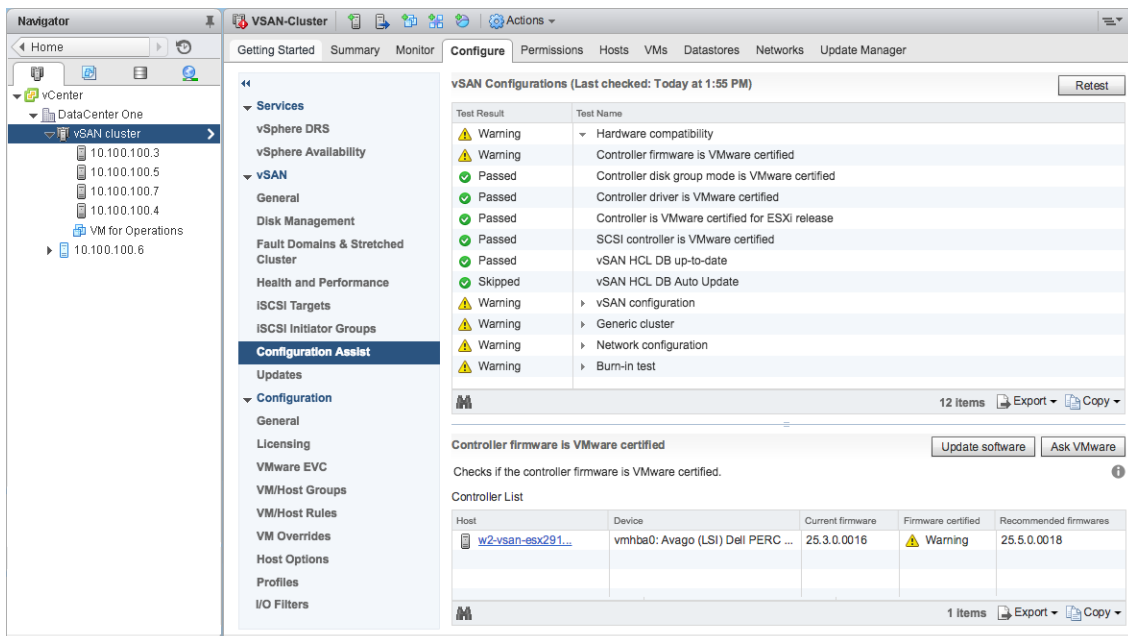
The Configuration Assist checks are divided into categories. Each category contains individual configuration checks.

Table 6-3. Configuration Assist Categories

Configuration Category	Description
Hardware compatibility	Checks the hardware components for the Virtual SAN cluster, to ensure that they are using supported hardware, software, and drivers.
vSAN configuration	Checks Virtual SAN configuration options.
Generic cluster	Checks basic cluster configuration options.
Network configuration	Checks Virtual SAN network configuration.
Burn-in test	Checks burn-in test operations.

If storage controller firmware or drivers do not meet the requirements listed in the *VMware Compatibility Guide*, you can use the Updates page to update the controllers.

Check Virtual SAN Configuration

You can view the configuration status of your Virtual SAN cluster, and resolve issues that affect the operation of your cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configuration** tab.

3   Under **vSAN**, click **Configuration Assist** to review the Virtual SAN configuration categories.

    If the Test Result column displays a warning icon, expand the category to review the results of individual configuration checks.

4   Select an individual configuration check and review the detailed information at the bottom of the page.

    You can click the **Ask VMware** button to open a knowledge base article that describes the check and provides information about how to resolve the issue.

    Some configuration checks provide additional buttons that help you complete the configuration.

Configure Distributed Switch for Virtual SAN

You can use the Configure New Distributed Switch for Virtual SAN wizard to configure a vSphere Distributed Switch to support Virtual SAN traffic.

If your cluster does not have a vSphere Distributed Switch configured to support Virtual SAN traffic, the Configuration Assist page issues a warning for **Network configuration > Use vDS for vSAN**.

Procedure

1   Navigate to your Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **Configuration Assist** and click to expand the **Network configuration** category.

4   Click **Use vDS for vSAN**. In the lower half of the page, click **Create vDS**.

5   In Name and type, enter a name for the new distributed switch, and choose whether to create a new switch or migrate an existing standard switch.

6   Select the unused adapters you want to migrate to the new distributed switch, and click **Next**.

7   (Optional) In Migrate infrastructure VMs, select the VM to treat as an infrastructure VM during the migration for existing standard switch, and click **Next**.

    This step is not necessary if you are creating a new distributed switch.

8   In Ready to complete, review the configuration, and click **Finish**.

Create VMkernel Network Adapter for Virtual SAN

You can use the New VMkernel Network Adapters for vSAN wizard to configure vmknics to support Virtual SAN traffic.

If ESXi hosts in your cluster do not have vmknics configured to support Virtual SAN traffic, the Configuration Assist page issues a warning for **Network configuration > All hosts have a vSAN vmknic configured**.

Procedure

1   Navigate to your Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **Configuration Assist** and click to expand the **Network configuration** category.

4   Click **All hosts have a vSAN vmknic configured**. In the lower half of the page, click **Create VMkernel Network Adapter**.

5   In Select hosts, select the check box for each host that does not have a vmknic configured for Virtual SAN, and click **Next**.

Hosts without a Virtual SAN vmknic are listed in the Configuration Assist page.

6   In Location and services, select a distributed switch and select the **vSAN traffic** check box. Click **Next**.

7   In vSAN adapter settings, select a port group, IP settings and configuration, and click **Next**.

8   In Ready to complete, review the configuration, and click **Finish**.

Install Controller Management Tools for Driver and Firmware Updates

Storage controller vendors provide a software management tool that Virtual SAN can use to update controller drivers and firmware. If the management tool is not present on ESXi hosts, you can download the tool.

The Updates page only supports specific storage controller models from selected vendors.

Prerequisites

■   Verify hardware compatibility on the Configuration Assist page.

■   DRS must be enabled if you must keep VMs running during the software updates.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configuration** tab.

3   Under vSAN, click **Updates** to review the components that are missing or ready to install.

4   Select the Management (Mgmt) tool for your controller, and click the **Download** icon.

The Management tool is downloaded from the Internet to your vCenter Server.

5   Click the **Update All** icon to install the management tool on the ESXi hosts in your cluster.

Confirm whether you want to update all hosts at once, or if you want to use a rolling update.

6   Click the **Refresh** icon.

The Updates page displays controller components that require an update.

What to do next

When the storage controller Management tool is available, the Updates page lists any missing drivers or firmware. You can update those missing components.

Update Storage Controller Drivers and Firmware

You can use Virtual SAN to update old or incorrect drivers and firmware on storage controllers.
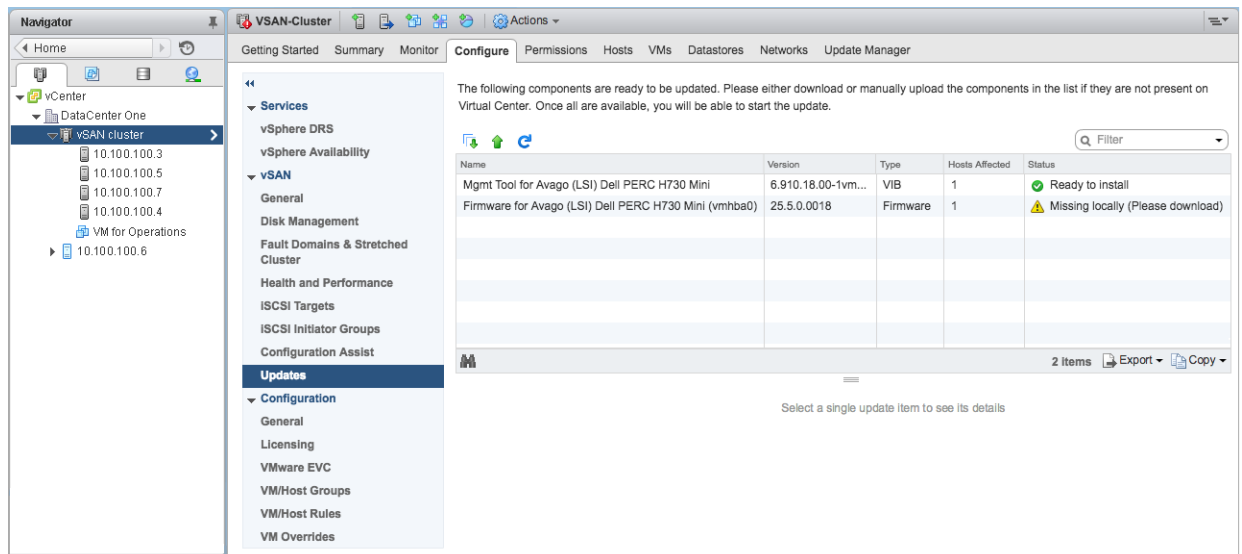
Configuration Assist verifies that your storage controllers use the latest driver and firmware version according to the *VMware Compatibility Guide*. If controller drivers or firmware do not meet the requirements, use the Updates page to perform driver and firmware updates.

Prerequisites

The controller Management tools for your storage devices must be present on the ESXi host.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configuration** tab.

3 Under vSAN, click **Updates** to review the components that are missing or ready to install.

The Updates page lists any missing firmware and driver components.

NOTE If the controller Management (Mgmt) tool is not available, you are prompted to download and install the Management tool. When the management tool is available, any missing drivers or firmware are listed.

---

4 Select the component you want to update, and click the **Update** icon to update the component on the ESXi hosts in your cluster. Or you can click the **Update All** icon to update all missing components.

Confirm whether you want to update all hosts at once, or if you want to use a rolling update.

NOTE For some management tools and drivers, the update process bypasses maintenance mode and performs a reboot based on the installation result. In these cases, the **MM Required** and **Reboot Required** fields are empty.

---

5 Click the **Refresh** icon.

The updated components are removed from the display.

Extending a Datastore Across Two Sites with Stretched Clusters

<div align="right">

7

</div>

You can create a stretched cluster that spans two geographic locations (or sites). Stretched clusters enable you to extend the Virtual SAN datastore across two sites to use it as stretched storage. The stretched cluster continues to function if a failure or scheduled maintenance occurs at one site.

This chapter includes the following topics:

■ "Introduction to Stretched Clusters," on page 61

■ "Stretched Cluster Design Considerations," on page 63

■ "Best Practices for Working with Stretched Clusters," on page 64

■ "Network Design for Stretched Clusters," on page 64

■ "Configure Virtual SAN Stretched Cluster," on page 65

■ "Change the Preferred Fault Domain," on page 66

■ "Change the Witness Host," on page 66

■ "Deploying a Virtual SAN Witness Appliance," on page 66

■ "Configure Network Interface for Witness Traffic," on page 67

■ "Convert a Stretched Cluster to a Standard Virtual SAN Cluster," on page 69

Introduction to Stretched Clusters

Stretched clusters extend the Virtual SAN cluster from a single site to two sites for a higher level of availability and intersite load balancing. Stretched clusters are typically deployed in environments where the distance between data centers is limited, such as metropolitan or campus environments.

You can use stretched clusters to manage planned maintenance and avoid disaster scenarios, because maintenance or loss of one site does not affect the overall operation of the cluster. In a stretched cluster configuration, both sites are active sites. If either site fails, Virtual SAN uses the storage on the other site. vSphere HA restarts any VM that must be restarted on the remaining active site.

You must designate one site as the preferred site. The other site becomes a secondary or nonpreferred site. The system uses the preferred site only in cases where there is a loss of network connection between the two active sites, so the one designated as preferred is the one that remains operational.

A Virtual SAN stretched cluster can tolerate one link failure at a time without data becoming unavailable. A link failure is a loss of network connection between the two sites or between one site and the witness host. During a site failure or loss of network connection, Virtual SAN automatically switches to fully functional sites.
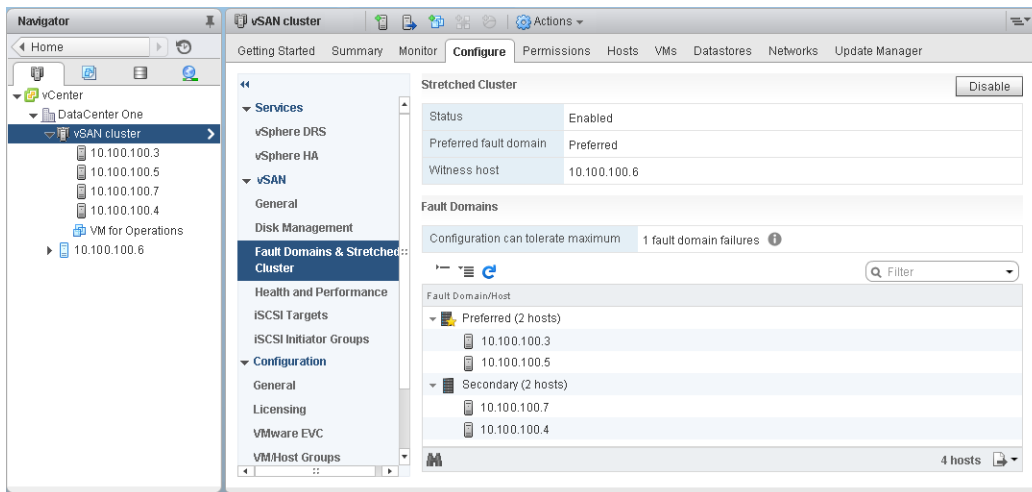
For more information about working with stretched clusters, see the *Virtual SAN Stretched Cluster Guide*.

Witness Host

Each stretched cluster consists of two sites and one witness host. The witness host resides at a third site and contains the witness components of virtual machine objects. It contains only metadata, and does not participate in storage operations.

The witness host serves as a tiebreaker when a decision must be made regarding availability of datastore components when the network connection between the two sites is lost. In this case, the witness host typically forms a Virtual SAN cluster with the preferred site. But if the preferred site becomes isolated from the secondary site and the witness, the witness host forms a cluster using the secondary site. When the preferred site is online again, data is resynchronized to ensure that both sites have the latest copies of all data.

If the witness host fails, all corresponding objects become noncompliant but are fully accessible.

The witness host has the following characteristics:

■ The witness host can use low bandwidth/high latency links.

■ The witness host cannot run VMs.

■ A single witness host can support only one Virtual SAN stretched cluster.

■ The witness host must have one VMkernel adapter with Virtual SAN traffic enabled, with connections to all hosts in the cluster. The witness host uses one VMkernel adapter for management and one VMkernel adapter for Virtual SAN data traffic. The witness host can have only one VMkernel adapter dedicated to Virtual SAN.

■ The witness host must be a standalone host dedicated to the stretched cluster. It cannot be added to any other cluster or moved in inventory through vCenter Server.

The witness host can be a physical host or an ESXi host running inside a VM. The VM witness host does not provide other types of functionality, such as storing or running VMs. Multiple witness hosts can run as VMs on a single physical server. For patching and basic networking and monitoring configuration, the VM witness host works in the same way as a typical ESXi host. You can manage it with vCenter Server, patch it and update it by using `esxcli` or vSphere Update Manager, and monitor it with standard tools that interact with ESXi hosts.

You can use a witness virtual appliance as the witness host in a stretched cluster. The witness virtual appliance is an ESXi host in a VM, packaged as an OVF or OVA. The appliance is available in different options, based on the size of the deployment.

Stretched Cluster Versus Fault Domains

Stretched clusters provide redundancy and failure protection across data centers in two geographical locations. Fault domains provide protection from rack-level failures within the same site. Each site in a stretched cluster resides in a separate fault domain.

A stretched cluster requires three fault domains: the preferred site, the secondary site, and a witness host.

In Virtual SAN 6.6 and later releases, you can provide an additional level of local fault protection for virtual machine objects in stretched clusters. When you configure a stretched cluster with four or more hosts in each site, the following policy rules are available for objects in the cluster:

- **Primary level of failures to tolerate (PFTT)**. This rule defines the number of host and device failures that a virtual machine object can tolerate across the two sites. The default value is 1, and the maximum value is 3.

- **Secondary level of failures to tolerate**. Defines the number of host and object failures that a virtual machine object can tolerate within a single site. The default value is 0, and the maximum value is 3.

- **Affinity**. This rule is available only if the **Primary level of failures to tolerate** is set to 0. You can set the Affinity rule to None, Preferred, or Secondary. This rule enables you to restrict virtual machine objects to a selected site in the stretched cluster. The default value is None.

NOTE   When you configure the **Secondary level of failures to tolerate** for the stretched cluster, the **Fault tolerance method** rule applies to the **Secondary level of failures to tolerate**. The failure tolerance method used for the **Primary level of failures to tolerate (PFTT)** defaults to RAID 1.

In a stretched cluster with local fault protection, even when one site is unavailable, the cluster can perform repairs on missing or broken components in the available site.

Stretched Cluster Design Considerations

Consider these guidelines when working with a Virtual SAN stretched cluster.

- Configure DRS settings for the stretched cluster.

  - DRS must be enabled on the cluster. If you place DRS in partially automated mode, you can control which VMs to migrate to each site.

  - Create two host groups, one for the preferred site and one for the secondary site.

  - Create two VM groups, one to hold the VMs on the preferred site and one to hold the VMs on the secondary site.

  - Create two VM-Host affinity rules that map VMs-to-host groups, and specify which VMs and hosts reside in the preferred site and which VMs and hosts reside in the secondary site.

  - Configure VM-Host affinity rules to perform the initial placement of VMs in the cluster.

- Configure HA settings for the stretched cluster.

  - HA must be enabled on the cluster.

  - HA rule settings should respect VM-Host affinity rules during failover.

  - Disable HA datastore heartbeats.

- Stretched clusters require on-disk format 2.0 or later. If necessary, upgrade the on-disk format before configuring a stretched cluster. See "Upgrade Virtual SAN Disk Format Using vSphere Web Client," on page 95.

- Configure the **Primary level of failures to tolerate** to 1 for stretched clusters.

- Virtual SAN stretched clusters do not support symmetric multiprocessing fault tolerance (SMP-FT).

- When a host is disconnected or not responding, you cannot add or remove the witness host. This limitation ensures that Virtual SAN collects enough information from all hosts before initiating reconfiguration operations.

- Using `esxcli` to add or remove hosts is not supported for stretched clusters.

Best Practices for Working with Stretched Clusters

When working with Virtual SAN stretched clusters, follow these recommendations for proper performance.

- If one of the sites (fault domains) in a stretched cluster is inaccessible, new VMs can still be provisioned in the sub-cluster containing the other two sites. These new VMs are implicitly force provisioned and will be non-compliant until the partitioned site rejoins the cluster. This implicit force provisioning is performed only when two of the three sites are available. A site here refers to either a data site or the witness host.

- If an entire site goes offline due to a power outage or loss of network connection, restart the site immediately, without much delay. Instead of restarting Virtual SAN hosts one by one, bring all hosts online approximately at the same time, ideally within a span of 10 minutes. By following this process, you avoid resynchronizing a large amount of data across the sites.

- If a host is permanently unavailable, remove the host from the cluster before you perform any reconfiguration tasks.

- If you want to clone a VM witness host to support multiple stretched clusters, do not configure the VM as a witness host before cloning it. First deploy the VM from OVF, then clone the VM, and configure each clone as a witness host for a different cluster. Or you can deploy as many VMs as you need from the OVF, and configure each one as a witness host for a different cluster.

Network Design for Stretched Clusters

All three sites in a stretched cluster communicate across the management network and across the Virtual SAN network. The VMs in both data sites communicate across a common virtual machine network.

A Virtual SAN stretched cluster must meet certain basic networking requirements.

- Management network requires connectivity across all three sites, using a Layer 2 stretched network or a Layer 3 network.

- Virtual SAN network requires connectivity across all three sites. VMware recommends using a Layer 2 stretched network between the two data sites and a Layer 3 network between the data sites and the witness host.

- VM network requires connectivity between the data sites, but not the witness host. VMware recommends using a Layer 2 stretched network between the data sites. In the event of a failure, the VMs do not require a new IP address to work on the remote site.

- vMotion network requires connectivity between the data sites, but not the witness host. VMware supports using a Layer 2 stretched or a Layer 3 network between data sites.

Using Static Routes on ESXi Hosts

If you use a single default gateway on ESXi hosts, note that each ESXi host contains a default TCP/IP stack that has a single default gateway. The default route is typically associated with the management network TCP/IP stack.

The management network and the Virtual SAN network might be isolated from one another. For example, the management network might use vmk0 on physical NIC 0, while the Virtual SAN network uses vmk2 on physical NIC 1 (separate network adapters for two distinct TCP/IP stacks). This configuration implies that the Virtual SAN network has no default gateway.

Consider a Virtual SAN network that is stretched over two data sites on a Layer 2 broadcast domain (for example, 172.10.0.0) and the witness host is on another broadcast domain (for example, 172.30.0.0). If the VMkernel adapters on a data site try to connect to the Virtual SAN network on the witness host, the connection will fail because the default gateway on the ESXi host is associated with the management network and there is no route from the management network to the Virtual SAN network.

You can use static routes to resolve this issue. Define a new routing entry that indicates which path to follow to reach a particular network. For a Virtual SAN network on a stretched cluster, you can add static routes to ensure proper communication across all hosts.

For example, you can add a static route to the hosts on each data site, so requests to reach the 172.30.0.0 witness network are routed through the 172.10.0.0 interface. Also add a static route to the witness host so that requests to reach the 172.10.0.0 network for the data sites are routed through the 172.30.0.0 interface.

Note If you use static routes, you must manually add the static routes for new ESXi hosts added to either site before those hosts can communicate across the cluster. If you replace the witness host, you must update the static route configuration.

Use the `esxcli network ip route` command to add static routes.

Configure Virtual SAN Stretched Cluster

Configure a Virtual SAN cluster that stretches across two geographic locations or sites.

Prerequisites

- Verify that you have a minimum of three hosts: one for the preferred site, one for the secondary site, and one host to act as a witness.

- Verify that you have configured one host to serve as the witness host for the stretched cluster. Verify that the witness host is not part of the Virtual SAN cluster, and that it has only one VMkernel adapter configured for Virtual SAN data traffic.

- Verify that the witness host is empty and does not contain any components. To configure an existing Virtual SAN host as a witness host, first evacuate all data from the host and delete the disk group.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3 Under vSAN, click **Fault Domains and Stretched Cluster**.

4 Click the Stretched Cluster **Configure** button to open the stretched cluster configuration wizard.

5 Select the fault domain that you want to assign to the secondary site and click **>>**.

   The hosts that are listed under the Preferred fault domain are in the preferred site.

6 Click **Next**.

7 Select a witness host that is not a member of the Virtual SAN stretched cluster and click **Next**.

8 Claim storage devices on the witness host and click **Next**.

   Claim storage devices on the witness host. Select one flash device for the cache tier, and one or more devices for the capacity tier.

9 On the Ready to complete page, review the configuration and click **Finish**.

Change the Preferred Fault Domain

You can configure the secondary site as the preferred site. The current preferred site becomes the secondary site.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select the secondary fault domain and click the **Mark Fault Domain as preferred for Stretched Cluster** icon (🌟).

5   Click **Yes** to confirm.

The selected fault domain is marked as the preferred fault domain.

Change the Witness Host

You can change the witness host for a Virtual SAN stretched cluster.

Change the ESXi host used as a witness host for your Virtual SAN stretched cluster.

Prerequisites

Verify that the witness host is not in use.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Click the **Change witness host** button.

5   Select a new host to use as a witness host, and click **Next**.

6   Claim disks on the new witness host, and click **Next**.

7   On the Ready to complete page, review the configuration, and click **Finish**.

Deploying a Virtual SAN Witness Appliance

Specific Virtual SAN configurations, such as a stretched cluster, require a witness host. Instead of using a dedicated physical ESXi host as a witness host, you can deploy the Virtual SAN witness appliance. The appliance is a preconfigured virtual machine that runs ESXi and is distributed as an OVA file.

Unlike a general-purpose ESXi host, the witness appliance does not run virtual machines. Its only purpose is to serve as a Virtual SAN witness.

The workflow to deploy and configure the Virtual SAN witness appliance includes this process.

1   Download the appliance from the VMware Web site.

2   Deploy the appliance to a Virtual SAN host or cluster. For more information, see Deploying OVF Templates in the *vSphere Virtual Machine Administration* documentation.

3   Configure the Virtual SAN network on the witness appliance.

4   Configure the management network on the witness appliance.

5   Add the appliance to vCenter Server as a witness ESXi host. Make sure to configure the Virtual SAN VMkernel interface on the host.

Set Up the Virtual SAN Network on the Witness Appliance

The Virtual SAN witness appliance includes two preconfigured network adapters. You must change the configuration of the second adapter so that the appliance can connect to the Virtual SAN network.

Procedure

1   In the vSphere Web Client, navigate to the virtual appliance that contains the witness host.

2   Right-click the appliance and select **Edit Settings**.

3   On the **Virtual Hardware** tab, expand the second Network adapter.

4   From the drop-down menu, select the vSAN port group and click **OK**.

Configure Management Network

Configure the witness appliance, so that it is reachable on the network.

By default, the appliance can automatically obtain networking parameters if your network includes a DHCP server. If not, you must configure appropriate settings.

Procedure

1   Power on your witness appliance and open its console.

Because your appliance is an ESXi host, you see the Direct Console User Interface (DCUI).

2   Press F2 and navigate to the Network Adapters page.

3   On the Network Adapters page, verify that at least one vmnic is selected for transport.

4   Configure the IPv4 parameters for the management network.

a   Navigate to the IPv4 Configuration section and change the default DHCP setting to static.

b   Enter the following settings:

- IP address

- Subnet mask

- Default gateway

5   Configure DNS parameters.

- Primary DNS server

- Alternate DNS server

- Hostname

Configure Network Interface for Witness Traffic

vSAN data traffic requires a low-latency, high-bandwidth link. Witness traffic can use a high-latency, low-bandwidth and routable link. To separate data traffic from witness traffic, you can configure a dedicated VMkernel network adapter for vSAN witness traffic.

You can separate data traffic from witness traffic in supported stretched cluster configurations. The VMkernel adapter used for vSAN data traffic and the VMkernel adapter used for witness traffic must be connected to the same physical switch.

You can add support for a direct network cross-connection to carry vSAN data traffic in a two-host vSAN stretched cluster. You can configure a separate network connection for witness traffic. On each data host in the cluster, configure the management VMkernel network adapter to also carry witness traffic. Do not configure the witness traffic type on the witness host.

Prerequisites

■ Verify that the data site to witness traffic connection has a minimum bandwidth of 100 MBps and a latency of less than 200 ms RTT.

■ Verify that vSAN traffic can be carried over a direct Ethernet cable connection with a speed of 10 GBps.

■ Verify that data traffic and witness traffic use the same IP version.

Procedure

1 Open an SSH connection to the ESXi host.

2 Use the `esxcli network ip interface list` command to determine which VMkernel network adapter is used for management traffic.

For example:

```
esxcli network  ip interface list
vmk0
   Name: vmk0
   MAC Address: e4:11:5b:11:8c:16
   Enabled: true
   Portset: vSwitch0
   Portgroup: Management Network
   Netstack Instance: defaultTcpipStack
   VDS Name: N/A
   VDS UUID: N/A
   VDS Port: N/A
   VDS Connection: −1
   Opaque Network ID: N/A
   Opaque Network Type: N/A
   External ID: N/A
   MTU: 1500
   TSO MSS: 65535
   Port ID: 33554437

vmk1
   Name: vmk1
   MAC Address: 00:50:56:6a:3a:74
   Enabled: true
   Portset: vSwitch1
   Portgroup: vsandata
   Netstack Instance: defaultTcpipStack
  VDS Name: N/A
   VDS UUID: N/A
   VDS Port: N/A
   VDS Connection: −1
   Opaque Network ID: N/A
   Opaque Network Type: N/A
```

```
External ID: N/A
MTU: 9000
TSO MSS: 65535
Port ID: 50331660
```

NOTE Multicast information is included for backward compatibility. vSAN 6.6 and later releases do not require multicast.

3 Use the `esxcli vsan network ip add` command to configure the management VMkernel network adapter to support witness traffic.

```
esxcli vsan network ip add –i vmkx –T=witness
```

4 Use the `esxcli vsan network list` command to verify the new network configuration.

For example:

```
esxcli vsan network list
Interface
   VmkNic Name: vmk0
   IP Protocol: IP
   Interface UUID: 8cf3ec57–c9ea–148b–56e1–a0369f56dcc0
   Agent Group Multicast Address: 224.2.3.4
   Agent Group IPv6 Multicast Address: ff19::2:3:4
   Agent Group Multicast Port: 23451
   Master Group Multicast Address: 224.1.2.3
   Master Group IPv6 Multicast Address: ff19::1:2:3
   Master Group Multicast Port: 12345
   Host Unicast Channel Bound Port: 12321
   Multicast TTL: 5
   Traffic Type: witness

Interface
   VmkNic Name: vmk1
   IP Protocol: IP
   Interface UUID: 6df3ec57–4fb6–5722–da3d–a0369f56dcc0
   Agent Group Multicast Address: 224.2.3.4
   Agent Group IPv6 Multicast Address: ff19::2:3:4
   Agent Group Multicast Port: 23451
   Master Group Multicast Address: 224.1.2.3
   Master Group IPv6 Multicast Address: ff19::1:2:3
   Master Group Multicast Port: 12345
   Host Unicast Channel Bound Port: 12321
   Multicast TTL: 5
   Traffic Type: vsan
```

In the vSphere Web Client, the management VMkernel network interface is not selected for vSAN traffic. Do not re-enable the interface in the vSphere Web Client.

Convert a Stretched Cluster to a Standard Virtual SAN Cluster

You can decommission a stretched cluster and convert it to a standard Virtual SAN cluster.

When you disable a stretched cluster, the witness host is removed, but the fault domain configuration remains. Because the witness host is not available, all witness components are missing for your virtual machines. To ensure full availability for your VMs, repair the cluster objects immediately.

Procedure

1 Navigate to the Virtual SAN stretched cluster in the vSphere Web Client.

2 Disable the stretched cluster.

    a Click the **Configure** tab.

    b Under vSAN, click **Fault Domains and Stretched Cluster**.

    c Click the Stretched Cluster **Configure** button.

       The stretched cluster configuration wizard is displayed.

    d Click **Disable**, and click **Yes** to confirm.

3 Remove the fault domain configuration.

    a Select a fault domain and click the **Remove selected fault domains** icon (✖). Click **Yes** to confirm.

    b Select the other fault domain and click the **Remove selected fault domains** icon (✖). Click **Yes** to confirm.

4 Repair the objects in the cluster.

    a Click the **Monitor** tab and select **Virtual SAN**.

    b Under Virtual SAN, click **Health** and click **vSAN object health**.

    c Click **Repair object immediately**.

Virtual SAN recreates the witness components within the cluster.

Increasing Space Efficiency in a Virtual SAN Cluster

8

You can use space efficiency techniques to reduce the amount of space for storing data. These techniques reduce the total storage space required to meet your needs.

This chapter includes the following topics:

Introduction to Virtual SAN Space Efficiency

You can use space efficiency techniques to reduce the amount of space for storing data. These techniques reduce the total storage capacity required to meet your needs.

You can enable deduplication and compression on a Virtual SAN cluster to eliminate duplicate data and reduce the amount of space needed to store data.

You can set the **Failure tolerance method** policy attribute on VMs to use RAID 5 or RAID 6 erasure coding. Erasure coding can protect your data while using less storage space than the default RAID 1 mirroring.

You can use deduplication and compression, and RAID 5 or RAID 6 erasure coding to increase storage space savings. RAID 5 or RAID 6 provide clearly defined space savings over RAID 1. Deduplication and compression can provide additional savings.

Using Deduplication and Compression

Virtual SAN can perform block-level deduplication and compression to save storage space. When you enable deduplication and compression on a Virtual SAN all-flash cluster, redundant data within each disk group is reduced.

Deduplication removes redundant data blocks, whereas compression removes additional redundant data within each data block. These techniques work together to reduce the amount of space required to store the data. Virtual SAN applies deduplication and then compression as it moves data from the cache tier to the capacity tier.

You can enable deduplication and compression as a cluster-wide setting, but they are applied on a disk group basis. When you enable deduplication and compression on a Virtual SAN cluster, redundant data within a particular disk group is reduced to a single copy.

You can enable deduplication and compression when you create a new Virtual SAN all-flash cluster or when you edit an existing Virtual SAN all-flash cluster. For more information about creating and editing Virtual SAN clusters, see "Enabling Virtual SAN," on page 47.

When you enable or disable deduplication and compression, Virtual SAN performs a rolling reformat of every disk group on every host. Depending on the data stored on the Virtual SAN datastore, this process might take a long time. It is recommended that you do not perform these operations frequently. If you plan to disable deduplication and compression, you must first verify that enough physical capacity is available to place your data.

NOTE Deduplication and compression might not be effective for encrypted VMs, because VM Encryption encrypts data on the host before it is written out to storage. Consider storage tradeoffs when using VM Encryption.

How to Manage Disks in a Cluster with Deduplication and Compression

Consider the following guidelines when managing disks in a cluster with deduplication and compression enabled.

■ Avoid adding disks to a disk group incrementally. For more efficient deduplication and compression, consider adding a new disk group to increase cluster storage capacity.

■ When you add a new disk group manually, add all of the capacity disks at the same time.

■ You cannot remove a single disk from a disk group. You must remove the entire disk group to make modifications.

■ A single disk failure causes the entire disk group to fail.

Verifying Space Savings from Deduplication and Compression

The amount of storage reduction from deduplication and compression depends on many factors, including the type of data stored and the number of duplicate blocks. Larger disk groups tend to provide a higher deduplication ratio. You can check the results of deduplication and compression by viewing the Deduplication and Compression Overview in the Virtual SAN Capacity monitor.
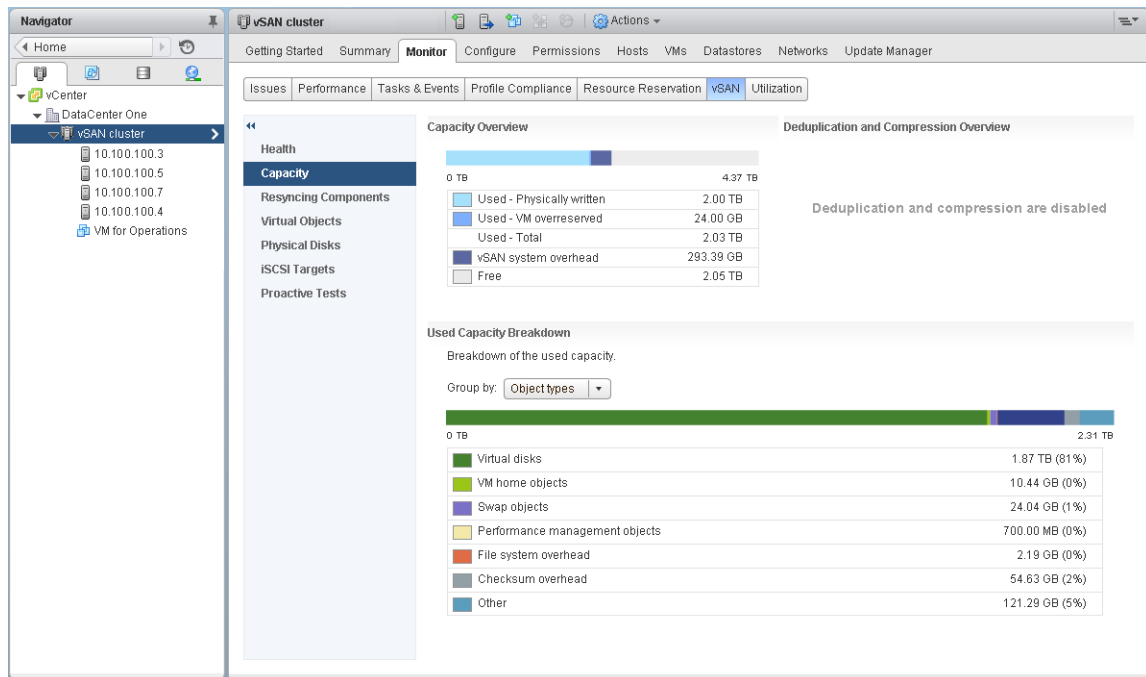
You can view the Deduplication and Compression Overview when you monitor Virtual SAN capacity in the vSphere Web Client. It displays information about the results of deduplication and compression. The Used Before space indicates the logical space required before applying deduplication and compression, while the Used After space indicates the physical space used after applying deduplication and compression. The Used After space also displays an overview of the amount of space saved, and the Deduplication and Compression ratio.

The Deduplication and Compression ratio is based on the logical (Used Before) space required to store data before applying deduplication and compression, in relation to the physical (Used After) space required after applying deduplication and compression. Specifically, the ratio is the Used Before space divided by the Used After space. For example, if the Used Before space is 3 GB, but the physical Used After space is 1 GB, the deduplication and compression ratio is 3x.

When deduplication and compression are enabled on the Virtual SAN cluster, it might take several minutes for capacity updates to be reflected in the Capacity monitor as disk space is reclaimed and reallocated.

Deduplication and Compression Design Considerations

Consider these guidelines when you configure deduplication and compression in a Virtual SAN cluster.

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

■ On-disk format version 3.0 or later is required to support deduplication and compression.

■ You must have a valid license to enable deduplication and compression on a cluster.

■ You can enable deduplication and compression only if the storage-claiming method is set to manual. You can change the storage-claiming method to automatic after deduplication and compression has been enabled.

■ When you enable deduplication and compression on a Virtual SAN cluster, all disk groups participate in data reduction through deduplication and compression.

■ Virtual SAN can eliminate duplicate data blocks within each disk group, but not across disk groups.

■ Capacity overhead for deduplication and compression is approximately five percent of total raw capacity.

■ Policies with 100 percent proportional capacity reservations are always honored. Using these policies can make deduplication and compression less efficient.

■ Policies with less than 100 percent proportional capacity are treated as if no proportional capacity reservation was requested. The object remains compliant with the policy, and no events are logged.

Enable Deduplication and Compression on a New Virtual SAN Cluster

You can enable deduplication and compression when you configure a new Virtual SAN all-flash cluster.

Procedure

1   Navigate to an existing cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under Virtual SAN, select **General** and click the **Configure vSAN** button.

4   Configure deduplication and compression on the cluster.

a   On the **vSAN capabilites** page, select the **Enable** check box under Deduplication and Compression.

b   (Optional) Enable reduced redundancy for your VMs.

See "Reducing VM Redundancy for Virtual SAN Cluster," on page 75.

5　On the **Claim disks** page, specify which disks to claim for the Virtual SAN cluster.

    a　Select a flash device to be used for capacity and click the **Claim for capacity tier** icon (⬛).

    b　Select a flash device to be used as cache and click the **Claim for cache tier** icon (⚡).

6　Complete your cluster configuration.

Enable Deduplication and Compression on Existing Virtual SAN Cluster

You can enable deduplication and compression by editing configuration parameters on an existing Virtual SAN cluster.

Prerequisites

Create a Virtual SAN cluster.

Procedure

1　Navigate to the Virtual SAN host cluster in the vSphere Web Client.

2　Click the **Configure** tab.

3　Under vSAN, select **General**.

4　In the vSAN is turned ON pane, click the **Edit** button.

5　Configure deduplication and compression.

    a　Set deduplication and compression to **Enabled**.

    b　(Optional) Enable reduced redundancy for your VMs.

      See "Reducing VM Redundancy for Virtual SAN Cluster," on page 75.

    c　Click **OK** to save your configuration changes.

While enabling deduplication and compression, Virtual SAN changes disk format on each disk group of the cluster. To accomplish this change, Virtual SAN evacuates data from the disk group, removes the disk group, and recreates it with a new format that supports deduplication and compression.

The enablement operation does not require virtual machine migration or DRS. The time required for this operation depends on the number of hosts in the cluster and amount of data. You can monitor the progress on the **Tasks and Events** tab.

Disable Deduplication and Compression

You can disable deduplication and compression on your Virtual SAN cluster.

When deduplication and compression are disabled on the Virtual SAN cluster, the size of the used capacity in the cluster can expand (based on the deduplication ratio). Before you disable deduplication and compression, verify that the cluster has enough capacity to handle the size of the expanded data.

Procedure

1　Navigate to the Virtual SAN host cluster in the vSphere Web Client.

2　Click the **Configure** tab.

3　Under vSAN, select **General**.

4　In the vSAN is turned ON pane, click the **Edit** button.

5   Disable deduplication and compression.

    a   Set the disk claiming mode to **Manual**.

    b   Set deduplication and compression to **Disabled**.

    c   Click **OK** to save your configuration changes.

While disabling deduplication and compression, Virtual SAN changes disk format on each disk group of the cluster. To accomplish this change, Virtual SAN evacuates data from the disk group, removes the disk group, and recreates it with a format that does not supports deduplication and compression.

The time required for this operation depends on the number of hosts in the cluster and amount of data. You can monitor the progress on the **Tasks and Events** tab.

Reducing VM Redundancy for Virtual SAN Cluster

When you enable deduplication and compression, in certain cases, you might need to reduce the level of protection for your virtual machines.

Enabling deduplication and compression requires a format change for disk groups. To accomplish this change, Virtual SAN evacuates data from the disk group, removes the disk group, and recreates it with a new format that supports deduplication and compression.

In certain environments, your Virtual SAN cluster might not have enough resources for the disk group to be fully evacuated. Examples for such deployments include a three-node cluster with no resources to evacuate the replica or witness while maintaining full protection. Or a four-node cluster with RAID-5 objects already deployed. In the latter case, you have no place to move part of the RAID-5 stripe, since RAID-5 objects require a minimum of four nodes.

You can still enable deduplication and compression and use the Allow Reduced Redundancy option. This option keeps the VMs running, but the VMs might be unable to tolerate the full level of failures defined in the VM storage policy. As a result, temporarily during the format change for deduplication and compression, your virtual machines might be at risk of experiencing data loss. Virtual SAN restores full compliance and redundancy after the format conversion is completed.

Adding or Removing Disks when Deduplication and Compression Is Enabled

When you add disks to a Virtual SAN cluster with enabled deduplication and compression, specific considerations apply.

■   You can add a capacity disk to a disk group with enabled deduplication and compression. However, for more efficient deduplication and compression, instead of adding capacity disks, create a new disk group to increase cluster storage capacity.

■   When you remove a disk form a cache tier, the entire disk group is removed. Removing a cache tier disk when deduplication and compression is enabled triggers data evacuation.

■   Deduplication and compression is implemented at a disk group level. You cannot remove a capacity disk from the cluster with enabled deduplication and compression. You must remove the entire disk group.

■   If a capacity disk fails, the entire disk group becomes unavailable. To resolve this issue, identify and replace the failing component immediately. When removing the failed disk group, use the No Data Migration option.

Using RAID 5 or RAID 6 Erasure Coding

You can use RAID 5 or RAID 6 erasure coding to protect against data loss and increase storage efficiency. Erasure coding can provide the same level of data protection as mirroring (RAID 1), while using less storage capacity.

RAID 5 or RAID 6 erasure coding enables Virtual SAN to tolerate the failure of up to two capacity devices in the datastore. You can configure RAID 5 on all-flash clusters with four or more fault domains. You can configure RAID 5 or RAID 6 on all-flash clusters with six or more fault domains.

RAID 5 or RAID 6 erasure coding requires less additional capacity to protect your data than RAID 1 mirroring. For example, a VM protected by a **Primary level of failures to tolerate** value of 1 with RAID 1 requires twice the virtual disk size, but with RAID 5 it requires 1.33 times the virtual disk size. The following table shows a general comparison between RAID 1 and RAID 5 or RAID 6.

Table 8-1. Capacity Required to Store and Protect Data at Different RAID Levels

RAID Configuration	Primary level of Failures to Tolerate	Data Size	Capacity Required
RAID 1 (mirroring)	1	100 GB	200 GB
RAID 5 or RAID 6 (erasure coding) with four fault domains	1	100 GB	133 GB
RAID 1 (mirroring)	2	100 GB	300 GB
RAID 5 or RAID 6 (erasure coding) with six fault domains	2	100 GB	150 GB

RAID 5 or RAID 6 erasure coding is a policy attribute that you can apply to virtual machine components. To use RAID 5, set **Failure tolerance method** to **RAID-5/6 (Erasure Coding) - Capacity** and **Primary level of failures to tolerate** to 1. To use RAID 6, set **Failure tolerance method** to **RAID-5/6 (Erasure Coding) - Capacity** and **Primary level of failures to tolerate** to 2. RAID 5 or RAID 6 erasure coding does not support a **Primary level of failures to tolerate** value of 3.

To use RAID 1, set **Failure tolerance method** to **RAID-1 (Mirroring) - Performance**. RAID 1 mirroring requires fewer I/O operations to the storage devices, so it can provide better performance. For example, a cluster resynchronization takes less time to complete with RAID 1.

NOTE In a Virtual SAN stretched cluster, the **Failure tolerance method** of **RAID-5/6 (Erasure Coding) - Capacity** applies only to the **Secondary level of failures to tolerate**.

For more information about configuring policies, see Chapter 13, "Using Virtual SAN Policies," on page 123.

RAID 5 or RAID 6 Design Considerations

Consider these guidelines when you configure RAID 5 or RAID 6 erasure coding in a Virtual SAN cluster.

- RAID 5 or RAID 6 erasure coding is available only on all-flash disk groups.

- On-disk format version 3.0 or later is required to support RAID 5 or RAID 6.

- You must have a valid license to enable RAID 5/6 on a cluster.

- You can achieve additional space savings by enabling deduplication and compression on the Virtual SAN cluster.

Using Encryption on a Virtual SAN Cluster

<div align="right">

9

</div>

You can use data at rest encryption to protect data in your Virtual SAN cluster.

Virtual SAN can perform data at rest encryption. Data is encrypted after all other processing, such as deduplication, is performed. Data at rest encryption protects data on storage devices, in case a device removed from the cluster.

Using encryption on your Virtual SAN cluster requires some preparation. After your environment is set up, you can enable encryption on your Virtual SAN cluster.

Virtual SAN encryption requires an external Key Management Server (KMS), the vCenter Server system, and your ESXi hosts. vCenter Server requests encryption keys from an external KMS. The KMS generates and stores the keys, and vCenter Server obtains the key IDs from the KMS and distributes them to the ESXi hosts.

vCenter Server does not store the KMS keys, but keeps a list of key IDs.

This chapter includes the following topics:

How Virtual SAN Encryption Works

When you enable encryption, Virtual SAN encrypts everything in the Virtual SAN datastore. All files are encrypted, so all virtual machines and their corresponding data are protected. Only administrators with encryption privileges can perform encryption and decryption tasks.

Virtual SAN uses encryption keys as follows:

■ vCenter Server requests an AES-256 Key Encryption Key (KEK) from the KMS. vCenter Server stores only the ID of the KEK, but not the key itself.

■ The ESXi host encrypts disk data using the industry standard AES-256 XTS mode. Each disk has a different randomly generated Data Encryption Key (DEK).

■ Each ESXi host uses the KEK to encrypt its DEKs, and stores the encrypted DEKs on disk. The host does not store the KEK on disk. If a host reboots, it requests the KEK with the corresponding ID from the KMS. The host can then decrypt its DEKs as needed.

■ A host key is used to encrypt core dumps, not data. All hosts in the same cluster use the same host key. When collecting support bundles, a random key is generated to re-encrypt the core dumps. Use a password when you encrypt the random key.

When a host reboots, it does not mount its disk groups until it receives the KEK. This process can take several minutes or longer to complete. You can monitor the status of the disk groups in the Virtual SAN health service, under **Physical disks > Software state health**.

Design Considerations for Virtual SAN Encryption

Consider these guidelines when working with Virtual SAN encryption.

■ Do not deploy your KMS server on the same Virtual SAN datastore that you plan to encrypt.

■ Encryption is CPU intensive. AES-NI significantly improves encryption performance. Enable AES-NI in your BIOS.

■ The witness host in a stretched cluster does not participate in Virtual SAN encryption. Only metadata is stored on the witness host.

■ Establish a policy regarding core dumps. Core dumps are encrypted because they can contain sensitive information such as keys. If you decrypt a core dump, carefully handle its sensitive information. ESXi core dumps might contain keys for the ESXi host and for the data on it.

   ■ Always use a password when you collect a `vm-support` bundle. You can specify the password when you generate the support bundle from the vSphere Web Client or using the `vm-support` command.

   The password recrypts core dumps that use internal keys to use keys that are based on the password. You can later use the password to decrypt any encrypted core dumps that might be included in the support bundle. Unencrypted core dumps or logs are not affected.

   ■ The password that you specify during `vm-support` bundle creation is not persisted in vSphere components. You are responsible for keeping track of passwords for support bundles.

Set Up the KMS Cluster

A Key Management Server (KMS) cluster provides the keys that you can use to encrypt the Virtual SAN datastore.

Before you can encrypt the Virtual SAN datastore, you must set up a KMS cluster to support encryption. That task includes adding the KMS to vCenter Server and establishing trust with the KMS. vCenter Server provisions encryption keys from the KMS cluster.

The KMS must support the Key Management Interoperability Protocol (KMIP) 1.1 standard.

Add a KMS to vCenter Server

You add a Key Management Server (KMS) to your vCenter Server system from the vSphere Web Client.

vCenter Server creates a KMS cluster when you add the first KMS instance. If you configure the KMS cluster on two or more vCenter Servers, make sure you use the same KMS cluster name.

NOTE Do not deploy your KMS servers on the Virtual SAN cluster you plan to encrypt. If a failure occurs, hosts in the Virtual SAN cluster must communicate with the KMS.

■ When you add the KMS, you are prompted to set this cluster as a default. You can later change the default cluster explicitly.

■ After vCenter Server creates the first cluster, you can add KMS instances from the same vendor to the cluster.

■ You can set up the cluster with only one KMS instance.

- If your environment supports KMS solutions from different vendors, you can add multiple KMS clusters.

Prerequisites

- Verify that the key server is in the *vSphere Compatibility Matrixes* and is KMIP 1.1 compliant.

- Verify that you have the required privileges: **Cryptographer.ManageKeyServers**

- Connecting to a KMS by using only an IPv6 address is not supported.

- Connecting to a KMS through a proxy server that requires user name or password is not supported.

Procedure

1 Log in to the vCenter Server system with the vSphere Web Client.

2 Browse the inventory list and select the vCenter Server instance.

3 Click **Configure** and click **Key Management Servers**.

4 Click **Add KMS**, specify the KMS information in the wizard, and click **OK**.

Option	Value
KMS cluster	Select **Create new cluster** for a new cluster. If a cluster exists, you can select that cluster.
Cluster name	Name for the KMS cluster. You can use this name to connect to the KMS if your vCenter Server instance becomes unavailable.
Server alias	Alias for the KMS. You can use this alias to connect to the KMS if your vCenter Server instance becomes unavailable.
Server address	IP address or FQDN of the KMS.
Server port	Port on which vCenter Server connects to the KMS.
Proxy address	Optional proxy address for connecting to the KMS.
Proxy port	Optional proxy port for connecting to the KMS.
User name	Some KMS vendors allow users to isolate encryption keys that are used by different users or groups by specifying a user name and password. Specify a user name only if your KMS supports this functionality, and if you intend to use it.
Password	Some KMS vendors allow users to isolate encryption keys that are used by different users or groups by specifying a user name and password. Specify a password only if your KMS supports this functionality, and if you intend to use it.

Establish a Trusted Connection by Exchanging Certificates

After you add the KMS to the vCenter Server system, you can establish a trusted connection. The exact process depends on the certificates that the KMS accepts, and on company policy.

Prerequisites

Add the KMS cluster.

Procedure

1 Log in to the vSphere Web Client, and select a vCenter Server system.

2 Click **Configure** and select **Key Management Servers**.

3 Select the KMS instance with which you want to establish a trusted connection.

4 Click **Establish trust with KMS**.

5   Select the option appropriate for your server and complete the steps.

Option	See
Root CA certificate	"Use the Root CA Certificate Option to Establish a Trusted Connection," on page 80.
Certificate	"Use the Certificate Option to Establish a Trusted Connection," on page 80.
New Certificate Signing Request	"Use the New Certificate Signing Request Option to Establish a Trusted Connection," on page 81.
Upload certificate and private key	"Use the Upload Certificate and Private Key Option to Establish a Trusted Connection," on page 81.

Use the Root CA Certificate Option to Establish a Trusted Connection

Some KMS vendors such as SafeNet require that you upload your root CA certificate to the KMS. All certificates that are signed by your root CA are then trusted by this KMS.

The root CA certificate that vSphere Virtual Machine Encryption uses is a self-signed certificate that is stored in a separate store in the VMware Endpoint Certificate Store (VECS) on the vCenter Server system.

NOTE   Generate a root CA certificate only if you want to replace existing certificates. If you do, other certificates that are signed by that root CA become invalid. You can generate a new root CA certificate as part of this workflow.

Procedure

1   Log in to the vSphere Web Client, and select a vCenter Server system.

2   Click **Configure** and select **Key Management Servers**.

3   Select the KMS instance with which you want to establish a trusted connection.

4   Select **Root CA Certificate** and click **OK**.

The Download Root CA Certificate dialog box is populated with the root certificate that vCenter Server uses for encryption. This certificate is stored in VECS.

5   Copy the certificate to the clipboard or download the certificate as a file.

6   Follow the instructions from your KMS vendor to upload the certificate to their system.

NOTE   Some KMS vendors, for example SafeNet, require that the KMS vendor restarts the KMS to pick up the root certificate that you upload.

What to do next

Finalize the certificate exchange. See "Complete the Trust Setup," on page 82.

Use the Certificate Option to Establish a Trusted Connection

Some KMS vendors such as Vormetric require that you upload the vCenter Server certificate to the KMS. After the upload, the KMS accepts traffic that comes from a system with that certificate.

vCenter Server generates a certificate to protect connections with the KMS. The certificate is stored in a separate key store in the VMware Endpoint Certificate Store (VECS) on the vCenter Server system.

Procedure

1   Log in to the vSphere Web Client, and select a vCenter Server system.

2   Click **Configure** and select **Key Management Servers**.

3   Select the KMS instance with which you want to establish a trusted connection.

4   Select **Certificate** and click **OK**.

The Download Certificate dialog box is populated with the root certificate that vCenter Server uses for encryption. This certificate is stored in VECS.

> **Note** Do not generate a new certificate unless you want to replace existing certificates.

5   Copy the certificate to the clipboard or download it as a file.

6   Follow the instructions from your KMS vendor to upload the certificate to the KMS.

What to do next

Finalize the trust relationship. See "Complete the Trust Setup," on page 82.

Use the New Certificate Signing Request Option to Establish a Trusted Connection

Some KMS vendors, for example Thales, require that vCenter Server generate a Certificate Signing Request (CSR) and send that CSR to the KMS. The KMS signs the CSR and returns the signed certificate. You can upload the signed certificate to vCenter Server.

Using the **New Certificate Signing Request** option is a two-step process. First you generate the CSR and send it to the KMS vendor. Then you upload the signed certificate that you receive from the KMS vendor to vCenter Server.

Procedure

1   Log in to the vSphere Web Client, and select a vCenter Server system.

2   Click **Configure** and select **Key Management Servers**.

3   Select the KMS instance with which you want to establish a trusted connection.

4   Select **New Certificate Signing Request** and click **OK**.

5   In the dialog box, copy the full certificate in the text box to the clipboard or download it as a file, and click **OK**.

Use the **Generate new CSR** button in the dialog box only if you explicitly want to generate a CSR. Using that option makes any signed certificates that are based on the old CSR invalid.

6   Follow the instructions from your KMS vendor to submit the CSR.

7   When you receive the signed certificate from the KMS vendor, click **Key Management Servers** again, and select **New Certificate Signing Request** again.

8   Paste the signed certificate into the bottom text box or click **Upload File** and upload the file, and click **OK**.

What to do next

Finalize the trust relationship. See "Complete the Trust Setup," on page 82.

Use the Upload Certificate and Private Key Option to Establish a Trusted Connection

Some KMS vendors such as HyTrust require that you upload the KMS server certificate and private key to the vCenter Server system.

Some KMS vendors generate a certificate and private key for the connection and make them available to you. After you upload the files, the KMS trusts your vCenter Server instance.

Prerequisites

■   Request a certificate and private key from the KMS vendor. The files are X509 files in PEM format.

Procedure

1 Log in to the vSphere Web Client, and select a vCenter Server system.

2 Click **Configure** and select **Key Management Servers**.

3 Select the KMS instance with which you want to establish a trusted connection.

4 Select **Upload certificate and private key** and click **OK**.

5 Paste the certificate that you received from the KMS vendor into the top text box or click **Upload File** to upload the certificate file.

6 Paste the key file into the bottom text box or click **Upload File** to upload the key file.

7 Click **OK**.

What to do next

Finalize the trust relationship. See "Complete the Trust Setup," on page 82.

Set the Default KMS Cluster

You must set the default KMS cluster if you do not make the first cluster the default cluster, or if your environment uses multiple clusters and you remove the default cluster.

Prerequisites

As a best practice, verify that the Connection Status in the Key Management Servers tab shows Normal and a green check mark.

Procedure

1 Log in to the vSphere Web Client and select a vCenter Server system.

2 Click the **Configure** tab and click **Key Management Servers** under **More**.

3 Select the cluster and click **Set KMS cluster as default**.

Do not select the server. The menu to set the default is available only for the cluster.

4 Click **Yes**.

The word `default` appears next to the cluster name.

Complete the Trust Setup

Unless the **Add Server** dialog box prompted you to trust the KMS, you must explicitly establish trust after certificate exchange is complete.

You can complete the trust setup, that is, make vCenter Server trust the KMS, either by trusting the KMS or by uploading a KMS certificate. You have two options:

■ Trust the certificate explicitly by using the **Refresh KMS certificate** option.

■ Upload a KMS leaf certificate or the KMS CA certificate to vCenter Server by using the **Upload KMS certificate** option.

NOTE If you upload the root CA certificate or the intermediate CA certificate, vCenter Server trusts all certificates that are signed by that CA. For strong security, upload a leaf certificate or an intermediate CA certificate that the KMS vendor controls.

Procedure

1 Log in to the vSphere Web Client, and select a vCenter Server system.

2   Click **Configure** and select **Key Management Servers**.

3   Select the KMS instance with which you want to establish a trusted connection.

4   To establish the trust relationship, refresh or upload the KMS certificate.

Option	Action
Refresh KMS certificate	a Click **All Actions**, and select **Refresh KMS certificate**.
	b In the dialog box that appears, click **Trust**.
Upload KMS certificate	a Click **All Actions**, and select **Upload KMS Certificate**.
	b In the dialog box that appears, click **Upload file**, upload a certificate file, and click **OK**.

Enable Encryption on a New Virtual SAN Cluster

You can enable encryption when you configure a new Virtual SAN cluster.

Prerequisites

■   Required privileges:

- ■   **Host.Inventory.EditCluster**

- ■   **Cryptographer.ManageEncryptionPolicy**

- ■   **Cryptographer.ManageKMS**

- ■   **Cryptographer.ManageKeys**

■   You must have set up a KMS cluster and established a trusted connection between vCenter Server and the KMS.

Procedure

1   Navigate to an existing cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General** and click the **Configure vSAN** button.

4   On the **vSAN capabilites** page, select the **Encryption** check box, and select a KMS cluster.

> **NOTE**   Make sure the **Erase disks before use** check box is deselected, unless you want to wipe existing data from the storage devices as they are encrypted.

5   On the **Claim disks** page, specify which disks to claim for the Virtual SAN cluster.

a   Select a flash device to be used for capacity and click the **Claim for capacity tier** icon (⛁).

b   Select a flash device to be used as cache and click the **Claim for cache tier** icon (⚡).

6   Complete your cluster configuration.

Encryption of data at rest is enabled on the Virtual SAN cluster. Virtual SAN encrypts all data added to the Virtual SAN datastore.

Generate New Encryption Keys

You can generate new encryption keys, in case a key expires or becomes compromised.

The following options are available when you generate new encryption keys for your Virtual SAN cluster.

■   If you generate a new KEK, all hosts in the Virtual SAN cluster receive the new KEK from the KMS. Each host's DEK is re-encrypted with the new KEK.

- If you choose to re-encrypt all data using new keys, a new KEK and new DEKs are generated. A rolling disk re-format is required to re-encrypt data.

Prerequisites

- Required privileges:
  - **Host.Inventory.EditCluster**
  - **Cryptographer.ManageKeys**

- You must have set up a KMS cluster and established a trusted connection between vCenter Server and the KMS.

Procedure

1   Navigate to the Virtual SAN host cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General**.

4   In the vSAN is turned ON pane, click the **Generate new encryption keys** button.

5   To generate a new KEK, click **OK**. The DEKs will be re-encrypted with the new KEK.

   - To generate a new KEK and new DEKs, and re-encrypt all data in the Virtual SAN cluster, select the following check box: **Also re-encrypt all data on the storage using new keys**.

   - If your Virtual SAN cluster has limited resources, select the **Allow Reduced Redundancy** check box. If you allow reduced redundancy, your data might be at risk during the disk reformat operation.

Enable Virtual SAN Encryption on Existing Virtual SAN Cluster

You can enable encryption by editing the configuration parameters of an existing Virtual SAN cluster.

Prerequisites

- Required privileges:
  - **Host.Inventory.EditCluster**
  - **Cryptographer.ManageEncryptionPolicy**
  - **Cryptographer.ManageKMS**
  - **Cryptographer.ManageKeys**

- You must have set up a KMS cluster and established a trusted connection between vCenter Server and the KMS.

- The cluster's disk-claiming mode must be set to manual.

Procedure

1   Navigate to the Virtual SAN host cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, select **General**.

4   In the vSAN is turned ON pane, click the **Edit** button.

5   On the Edit vSAN settings dialog, check the **Encryption** check box, and select a KMS cluster.

6   (Optional) If the storage devices in your cluster contain sensitive data, select the **Erase disks before use** check box.

   This setting directs Virtual SAN to wipe existing data from the storage devices as they are encrypted.

7   Click **OK**.

A rolling reformat of all disk groups takes places as Virtual SAN encrypts all data in the Virtual SAN datastore.

Virtual SAN Encryption and Core Dumps

If your Virtual SAN cluster uses encryption, and if an error occurs on the ESXi host, the resulting core dump is encrypted to protect customer data. Core dumps that are included in the vm-support package are also encrypted.

NOTE   Core dumps can contain sensitive information. Follow your organization's data security and privacy policy when handling core dumps.

Core Dumps on ESXi Hosts

When an ESXi host crashes, an encrypted core dump is generated and the host reboots. The core dump is encrypted with the host key that is in the ESXi key cache. What you can do next depends on several factors.

■   In most cases, vCenter Server retrieves the key for the host from the KMS and attempts to push the key to the ESXi host after reboot. If the operation is successful, you can generate the vm-support package and you can decrypt or re-encrypt the core dump.

■   If vCenter Server cannot connect to the ESXi host, you might be able to retrieve the key from the KMS.

■   If the host used a custom key, and that key differs from the key that vCenter Server pushes to the host, you cannot manipulate the core dump. Avoid using custom keys.

Core Dumps and vm-support Packages

When you contact VMware Technical Support because of a serious error, your support representative usually asks you to generate a vm-support package. The package includes log files and other information, including core dumps. If support representatives cannot resolve the issues by looking at log files and other information, you can decrypt the core dumps to make relevant information available. Follow your organization's security and privacy policy to protect sensitive information, such as host keys.

Core Dumps on vCenter Server Systems

A core dump on a vCenter Server system is not encrypted. vCenter Server already contains potentially sensitive information. At the minimum, ensure that the Windows system on which vCenter Server runs or the vCenter Server Appliance is protected. You also might consider turning off core dumps for the vCenter Server system. Other information in log files can help determine the problem.

Collect a vm-support Package for an ESXi Host in an Encrypted Virtual SAN Cluster

If encryption is enabled on a Virtual SAN cluster, any core dumps in the vm-support package are encrypted. You can collect the package from the vSphere Web Client, and you can specify a password if you expect to decrypt the core dump later.

The vm-support package includes log files, core dump files, and more.

Prerequisites

Inform your support representative that encryption is enabled for the Virtual SAN cluster. Your support representative might ask you to decrypt core dumps to extract relevant information.

NOTE   Core dumps can contain sensitive information. Follow your organization's security and privacy policy to protect sensitive information such as host keys.

Procedure

1   Log in to vCenter Server with the vSphere Web Client.

2   Click **Hosts and Clusters**, and right-click the ESXi host.

3   Select **Export System Logs**.

4   In the dialog box, select **Password for encrypted core dumps**, and specify and confirm a password.

5   Leave the defaults for other options or make changes if requested by VMware Technical Support, and click **Finish**.

6   Specify a location for the file.

7   If your support representative asked you to decrypt the core dump in the `vm-support` package, log in to any ESXi host and follow these steps.

   a   Log in to the ESXi and connect to the directory where the `vm-support` package is located.

       The filename follows the pattern **esx.date_and_time.tgz**.

   b   Make sure that the directory has enough space for the package, the uncompressed package, and the recompressed package, or move the package.

   c   Extract the package to the local directory.

       ```
       vm-support -x *.tgz .
       ```

       The resulting file hierarchy might contain core dump files for the ESXi host, usually in `/var/core`, and might contain multiple core dump files for virtual machines.

   d   Decrypt each encrypted core dump file separately.

       ```
       crypto-util envelope extract --offset 4096 --keyfile vm-support-incident-key-file
       --password encryptedZdump decryptedZdump
       ```

       vm-support-incident-key-file is the incident key file that you find at the top level in the directory.

       encryptedZdump is the name of the encrypted core dump file.

       decryptedZdump is the name for the file that the command generates. Make the name similar to the *encryptedZdump* name.

   e   Provide the password that you specified when you created the `vm-support` package.

   f   Remove the encrypted core dumps, and compress the package again.

       ```
       vm-support --reconstruct
       ```

8   Remove any files that contain confidential information.

Decrypt or Re-encrypt an Encrypted Core Dump

You can decrypt or re-encrypt an encrypted core dump on your ESXi host by using the `crypto-util` CLI.

You can decrypt and examine the core dumps in the `vm-support` package yourself. Core dumps might contain sensitive information. Follow your organization's security and privacy policy to protect sensitive information, such as host keys.

For details about re-encrypting a core dump and other features of `crypto-util`, see the command-line help.

NOTE `crypto-util` is for advanced users.

Prerequisites

The ESXi host key that was used to encrypt the core dump must be available on the ESXi host that generated the core dump.

Procedure

1   Log directly in to the ESXi host on which the core dump occurred.

   If the ESXi host is in lockdown mode, or if SSH access is disabled, you might have to enable access first.

2   Determine whether the core dump is encrypted.

Option	Description
Monitor core dump	`crypto-util envelope describe vmmcores.ve`
zdump file	`crypto-util envelope describe --offset 4096 zdumpFile`

3   Decrypt the core dump, depending on its type.

Option	Description
Monitor core dump	`crypto-util envelope extract vmmcores.ve vmmcores`
zdump file	`crypto-util envelope extract --offset 4096 zdumpEncrypted zdumpUnencrypted`

Upgrading the Virtual SAN Cluster

10

Upgrading Virtual SAN is a multistage process, in which you must perform the upgrade procedures in the order described here.

Before you attempt to upgrade, make sure you understand the complete upgrade process clearly to ensure a smooth and uninterrupted upgrade. If you are not familiar with the general vSphere upgrade procedure, you should first review the *vSphere Upgrade* documentation.

NOTE Failure to follow the sequence of upgrade tasks described here will lead to data loss and cluster failure.

The Virtual SAN cluster upgrade proceeds in the following sequence of tasks.

1 Upgrade the vCenter Server. See the *vSphere Upgrade* documentation.

2 Upgrade the ESXi hosts. See "Upgrade the ESXi Hosts," on page 91. For information about migrating and preparing your ESXi hosts for upgrade, see the *vSphere Upgrade* documentation.

3 Upgrade the Virtual SAN disk format. Upgrading the disk format is optional, but for best results, upgrade the objects to use the latest version. The on-disk format exposes your environment to the complete feature set of Virtual SAN. See "Upgrade Virtual SAN Disk Format Using RVC," on page 96.

This chapter includes the following topics:

- "Before You Upgrade Virtual SAN," on page 89
- "Upgrade the vCenter Server," on page 91
- "Upgrade the ESXi Hosts," on page 91
- "About the Virtual SAN Disk Format," on page 93
- "Verify the Virtual SAN Cluster Upgrade," on page 97
- "Using the RVC Upgrade Command Options," on page 98

Before You Upgrade Virtual SAN

Plan and design your upgrade to be fail-safe. Before you attempt to upgrade Virtual SAN, verify that your environment meets the vSphere hardware and software requirements.

Upgrade Prerequisite

Consider the aspects that could delay the overall upgrade process. For guidelines and best practices, see the *vSphere Upgrade* documentation.

Review the key requirements before you upgrade your cluster to Virtual SAN 6.6.

Table 10-1. Upgrade Prerequisite

Upgrade Prerequisites	Description
Software, hardware, drivers, firmware, and storage I/O controllers	Verify that the software and hardware components, drivers, firmware, and storage I/O controllers that you plan on using are supported by Virtual SAN for 6.6 and later, and are listed on the VMware Compatibility Guide Web site at http://www.vmware.com/resources/compatibility/search.php.
Virtual SAN version	Verify that you are using the latest version of Virtual SAN. If you are currently running a beta version and plan on upgrading Virtual SAN to 6.6, your upgrade will fail. When you upgrade from a beta version, you must perform a fresh deployment of Virtual SAN.
Disk space	Verify that you have enough space available to complete the software version upgrade. The amount of disk storage needed for the vCenter Server installation depends on your vCenter Server configuration. For guidelines about the disk space required for upgrading vSphere, see the *vSphere Upgrade* documentation.
Virtual SAN disk format	Verify that you have enough capacity available to upgrade the disk format. If you do not have free space equal to the consumed capacity of the largest disk group, with the space available on disk groups other than the disk groups that are being converted, you must choose **Allow reduced redundancy** as the data migration option.
	For example, the largest disk group in a cluster has 10 TB of physical capacity, but only 5 TB is being consumed. An additional 5 TB of spare capacity will be needed elsewhere in the cluster, excluding the disk groups that are being migrated. When upgrading the Virtual SAN disk format, verify that the hosts are not in maintenance mode. When any member host of a Virtual SAN cluster enters maintenance mode, the cluster capacity is automatically reduced, because the member host no longer contributes storage to the cluster and the capacity on the host is unavailable for data. For information about various evacuation modes, see the "Place a Member of Virtual SAN Cluster in Maintenance Mode," on page 112.
Virtual SAN hosts	Verify that you have placed the Virtual SAN hosts in maintenance mode and selected the **Ensure data accessibility** or **Evacuate all data** option.
	You can use the vSphere Update Manager for automating and testing the upgrade process. However, when you use vSphere Update Manager to upgrade Virtual SAN, the default evacuation mode is **Ensure data accessibility**. When you use the **Ensure data accessibility** mode, your data is not completely protected, and if you encounter a failure while upgrading Virtual SAN, you might experience unexpected data loss. However, the **Ensure data accessibility** mode is faster than the **Evacuate all data** mode, because you do not need to move all data to another host in the cluster. For information about various evacuation modes, see the "Place a Member of Virtual SAN Cluster in Maintenance Mode," on page 112.
Virtual Machines	Verify that you have backed up your virtual machines.

Recommendations

Consider the following recommendations when deploying ESXi hosts for use with Virtual SAN:

■ If ESXi hosts are configured with memory capacity of 512 GB or less, use SATADOM, SD, USB, or hard disk devices as the installation media.

■ If ESXi hosts are configured with memory capacity greater than 512 GB, use a separate magnetic disk or flash device as the installation device. If you are using a separate device, verify that Virtual SAN is not claiming the device.

■ When you boot a Virtual SAN host from a SATADOM device, you must use a single-level cell (SLC) device and the size of the boot device must be at least 16 GB.

Virtual SAN 6.5 and later enables you to adjust the boot size requirements for an ESXI host in a Virtual SAN cluster. For more information, see the VMware knowledge base article at http://kb.vmware.com/kb/2147881.

Upgrading the Witness Host in a Two Host or Stretched Cluster

The witness host for a two host cluster or stretched cluster is located outside of the Virtual SAN cluster, but it is managed by the same vCenter Server. You can use the same process to upgrade the witness host as you use for a Virtual SAN data host.

Do not upgrade the witness host until all data hosts have been upgraded and have exited maintenance mode.

Using vSphere Update Manager to upgrade hosts in parallel can result in the witness host being upgraded in parallel with one of the data hosts. To avoid upgrade problems, configure vSphere Update Manager so it does not upgrade the witness host in parallel with the data hosts.

Upgrade the vCenter Server

This first task to perform during the Virtual SAN upgrade is a general vSphere upgrade, which includes upgrading vCenter Server and ESXi hosts.

VMware supports in-place upgrades on 64-bit systems from vCenter Server 4.x, vCenter Server 5.0.x, vCenter Server 5.1.x, and vCenter Server 5.5 to vCenter Server 6.0 and later. The vCenter Server upgrade includes a database schema upgrade and an upgrade of the vCenter Server. Instead of performing an in-place upgrade to vCenter Server, you can use a different machine for the upgrade. For detailed instructions and various upgrade options, see the *vSphere Upgrade* documentation.

Upgrade the ESXi Hosts

After you upgrade the vCenter Server, the next task for the Virtual SAN cluster upgrade is upgrading the ESXi hosts to use the current version.

If you have multiple hosts in the Virtual SAN cluster, and you use vSphere Update Manager to upgrade the hosts, the default evacuation mode is **Ensure data accessibility**. If you use this mode, and while upgrading Virtual SAN you encounter a failure, your data will be at risk. For information about working with evacuation modes, see "Place a Member of Virtual SAN Cluster in Maintenance Mode," on page 112

For information about using vSphere Update Manager, see the documentation Web site at https://www.vmware.com/support/pubs/vum_pubs.html .

Before you attempt to upgrade the ESXi hosts, review the best practices discussed in the *vSphere Upgrade* documentation. VMware provides several ESXi upgrade options. Choose the upgrade option that works best with the type of host that you are upgrading. For more information about various upgrade options, see the *vSphere Upgrade* documentation.

Prerequisites

■ Verify that you have sufficient disk space for upgrading the ESXi hosts. For guidelines about the disk space requirement, see the *vSphere Upgrade* documentation.

■ Verify that you are using the latest version of ESXi. You can download the latest ESXi installer from the VMware product download Web site at https://my.vmware.com/web/vmware/downloads.

■ Verify that you are using the latest version of vCenter Server.

■ Verify the compatibility of the network configuration, storage I/O controller, storage device, and backup software.

■ Verify that you have backed up the virtual machines.

■ Use Distributed Resource Scheduler (DRS) to prevent virtual machine downtime during the upgrade. Verify that the automation level for each virtual machine is set to **Fully Automated** mode to help DRS migrate virtual machines when hosts are entering maintenance mode. Alternatively, you can also power off all virtual machines or perform manual migration.

Procedure

1   Place the host that you intend to upgrade in maintenance mode.

You must begin your upgrade path with ESXi 5.5 or later hosts in the Virtual SAN cluster.

a   Right-click the host in the vSphere Web Client navigator and select **Maintenance Mode > Enter Maintenance Mode**.

b   Select the **Ensure data accessibility** or **Evacuate all data** evacuation mode, depending on your requirement, and wait for the host to enter maintenance mode.

If you are using vSphere Update Manager to upgrade the host, or if you are working with a three-host cluster, the default evacuation mode available is **Ensure data accessibility**. This mode is faster than the **Evacuate all data** mode. However, the **Ensure data accessibility** mode does not fully protect your data. During failure your data might be at risk and you might experience downtime, and unexpected data loss.

2   Upload the software to the datastore of your ESXi host and verify that the file is available in the directory inside the datastore. For example, you can upload the software to `/vmfs/volumes/<datastore>/VMware-ESXi-6.0.0-1921158-depot.zip`.

3   Run the `esxcli` command `install -d /vmfs/volumes/53b536fd-34123144-8531-00505682e44d/depot/VMware-ESXi-6.0.0-1921158-depot.zip --no-sig-check`. Use the `esxcli` software VIB to run this command.

After the ESXi host has installed successfully, you see the following message:

`The update completed successfully, but the system needs to be rebooted for the changes to be effective.`

4   You must manually restart your ESXi host from the vSphere Web Client.

a   Navigate to the ESXi host in the vSphere Web Client inventory.

b   Right-click the host, select **Power > Reboot**, click **Yes** to confirm, and then wait for the host to restart.

c   Right-click the host, select **Connection > Disconnect**, and then select **Connection > Connect** to reconnect to the host.

To upgrade the remaining hosts in the cluster, repeat this procedure for each host.

If you have multiple hosts in your Virtual SAN cluster, you can use vSphere Update Manager to upgrade the remaining hosts.

5   Exit maintenance mode.

What to do next

1   (Optional) Upgrade the Virtual SAN disk format. See "Upgrade Virtual SAN Disk Format Using RVC," on page 96.

2   Verify the host license. In most cases, you must reapply your host license. You can use vSphere Web Client and vCenter Server for applying host licenses. For more information about applying host licenses, see the *vCenter Server and Host Management* documentation.

3   (Optional) Upgrade the virtual machines on the hosts by using the vSphere Web Client or vSphere Update Manager.

About the Virtual SAN Disk Format

The disk format upgrade is optional and a Virtual SAN cluster continues to run smoothly if you use a previous disk format version.

For best results, upgrade the objects to use the latest on-disk format. The latest on-disk format provides the complete feature set of Virtual SAN.

Depending on the size of disk groups, the disk format upgrade can be time-consuming because the disk groups are upgraded one at a time. For each disk group upgrade, all data from each device is evacuated and the disk group is removed from the Virtual SAN cluster. The disk group is then added back to Virtual SAN with the new on-disk format.

NOTE   Once you upgrade the on-disk format, you cannot roll back software on the hosts or add certain older hosts to the cluster.

When you initiate an upgrade of the on-disk format, Virtual SAN performs several operations that you can monitor from the Resyncing Components page. The table summarizes each process that takes place during the disk format upgrade.

Table 10-2. Upgrade Progress

Percentage of Completion	Description
0%-5%	Cluster check. Cluster components are checked and prepared for the upgrade. This process takes a few minutes. Virtual SAN verifies that no outstanding issues exist that can prevent completion of the upgrade. ■ All hosts are connected. ■ All hosts have the correct software version. ■ All disks are healthy. ■ All objects are accessible.
5%-10%	Disk group upgrade. Virtual SAN performs the initial disk upgrade with no data migration. This process takes a few minutes.
10%-15%	Object realignment. Virtual SAN modifies the layout of all objects to ensure that they are properly aligned. This process can take a few minutes for a small system with few snapshots. It can take many hours or even days for large a system with many snapshots, many fragmented writes, and many unaligned objects.
15%-95%	Disk group removal and reformat. Each disk group is removed from the cluster, reformatted, and added back to the cluster. The time required for this process varies, depending on the megabytes allocated and the system utilization. A system at or near its I/O capacity transfers slowly.
95%-100%	Final object version upgrade. Object conversion to the new on-disk format and resynchronization is completed. The time required for this process varies, depending on the amount of space used and whether the **Allow reduced redundancy** option is selected.

During the upgrade, you can monitor the upgrade process from the vSphere Web Client when you navigate to the Resyncing Components page. See "Monitor the Resynchronization Tasks in the Virtual SAN Cluster," on page 134. You also can use the RVC vsan.upgrade_status <cluster> command to monitor the upgrade. Use the optional –r <seconds> flag to refresh the upgrade status periodically until you press Ctrl+C. The minimum number of seconds allowed between each refresh is 60.

You can monitor other upgrade tasks, such as device removal and upgrade, from the vSphere Web Client in the Recent Tasks pane of the status bar.

The following considerations apply when upgrading the disk format:

■ If you upgrade a cluster with three hosts, and you want to perform a full evacuation, the evacuation fails for objects with a **Primary level of failures to tolerate** greater than zero. A three-host cluster cannot reprotect a disk group that is being fully evacuated using the resources of only two hosts. For example, when the **Primary level of failures to tolerate** is set to 1, Virtual SAN requires three protection components (two mirrors and a witness), where each protection component is placed on a separate host.

   For a three-host cluster, you must choose the **Ensure data accessibility** evacuation mode. When in this mode, any hardware failure might result in data loss.

   You also must ensure that enough free space is available. The space must be equal to the logical consumed capacity of the largest disk group. This capacity must be available on a disk group separate from the one that is being migrated.

■ When upgrading a three-host cluster or when upgrading a cluster with limited resources, allow the virtual machines to operate in a reduced redundancy mode. Run the RVC command with the option, `vsan.ondisk_upgrade --allow-reduced-redundancy`.

■ Using the `--allow-reduced-redundancy` command option means that certain virtual machines might be unable to tolerate failures during the migration. This lowered tolerance for failure also can cause data loss. Virtual SAN restores full compliance and redundancy after the upgrade is completed. During the upgrade, the compliance status of virtual machines and their redundancies is temporarily noncompliant. After you complete the upgrade and finish all rebuild tasks, the virtual machines will become compliant.

■ While the upgrade is in progress, do not remove or disconnect any host, and do not place a host in maintenance mode. These actions might cause the upgrade to fail.

For information about the RVC commands and command options, see the *RVC Command Reference Guide*.

Upgrade Virtual SAN Disk Format Using vSphere Web Client

After you have finished upgrading the Virtual SAN hosts, you can perform the disk format upgrade.
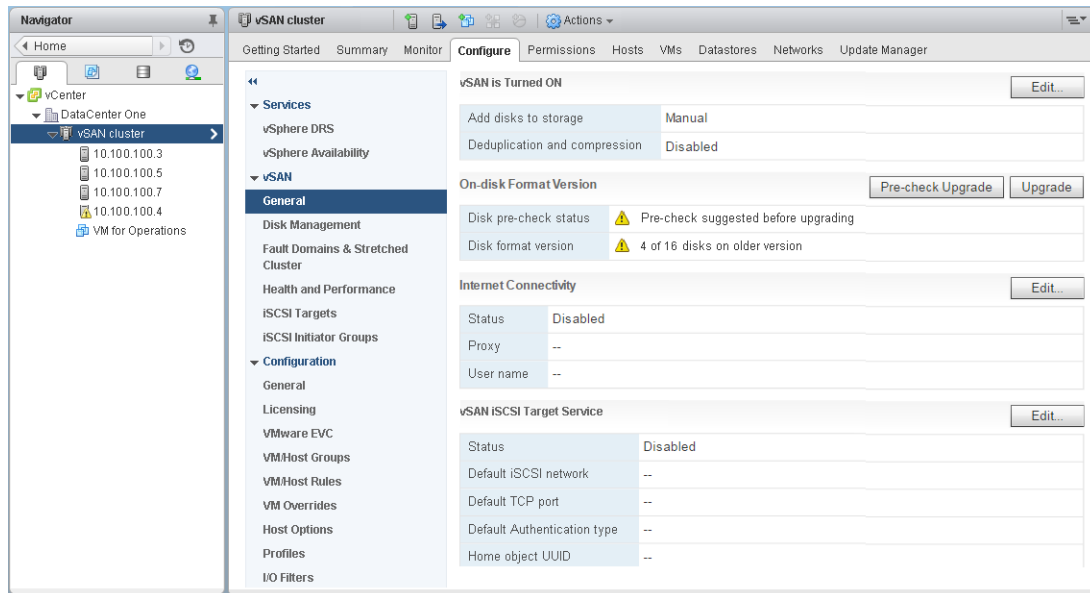
NOTE If you enable encryption or deduplication and compression on an existing Virtual SAN cluster, the on-disk format is automatically upgraded to the latest version. This procedure is not required. You can avoid reformatting the disk groups twice. See "Edit Virtual SAN Settings," on page 50.

Prerequisites

- Verify that you are using the updated version of vCenter Server.

- Verify that you are using the latest version of ESXi hosts.

- Verify that the disks are in a healthy state. Navigate to the Disk Management page in the vSphere Web Client to verify the object status.

- Verify that the hardware and software that you plan on using are certified and listed in the VMware Compatibility Guide Web site at http://www.vmware.com/resources/compatibility/search.php.

- Verify that you have enough free space to perform the disk format upgrade. Run the RVC command, `vsan.whatif_host_failures`, to determine whether you have enough capacity to successfully finish the upgrade or perform a component rebuild, in case you encounter any failure during the upgrade.

- Verify that your hosts are not in maintenance mode. When upgrading the disk format, you should not place the hosts in maintenance mode. When any member host of a Virtual SAN cluster enters maintenance mode, the available resource capacity in the cluster is reduced because the member host no longer contributes capacity to the cluster and the cluster upgrade might fail.

- Verify that there are no component rebuilding tasks currently in progress in the Virtual SAN cluster. See "Monitor the Resynchronization Tasks in the Virtual SAN Cluster," on page 134.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3 Under vSAN, select **General**.

4   (Optional) Under On-disk Format Version, click **Pre-check Upgrade**.

The upgrade pre-check analyzes the cluster to uncover any issues that might prevent a successful upgrade. Some of the items checked are host status, disk status, network status, and object status. Upgrade issues are displayed in the **Disk pre-check status** field.

5   Under On-disk Format Version, click **Upgrade**.

6   Click **Yes** on the Upgrade dialog to perform the upgrade of the on-disk format.

Virtual SAN performs a rolling reboot of each disk group in the cluster. The On-disk Format Version column displays the disk format version of storage devices in the cluster. The Disks with outdated version column indicates the number of devices using the new format. When the upgrade is successful, the Disks with outdated version will be 0.

If a failure occurs during the upgrade, you can check the Resyncing Components page in the vSphere Web Client. Wait for all resynchronizations to complete, and run the upgrade again. You also can check the cluster health using the health service. After you have resolved any issues raised by the health checks, you can run the upgrade again.

Upgrade Virtual SAN Disk Format Using RVC

After you have finished upgrading the Virtual SAN hosts, you can use the Ruby vSphere Console (RVC) to continue with the disk format upgrade.

Prerequisites

■   Verify that you are using the updated version of vCenter Server.

■   Verify that the version of the ESXi hosts running in the Virtual SAN cluster is 6.5 or later.

■   Verify that the disks are in a healthy state from the Disk Management page in the vSphere Web Client. You can also run the `vsan.disk_stats` RVC command to verify disk status.

■   Verify that the hardware and software that you plan on using are certified and listed in the VMware Compatibility Guide Web site at http://www.vmware.com/resources/compatibility/search.php.

■   Verify that you have enough free space to perform the disk format upgrade. Run the RVC `vsan.whatif_host_failures` command to determine that you have enough capacity to successfully finish the upgrade or perform a component rebuild in case you encounter failure during the upgrade.

■   Verify that you have PuTTY or similar SSH client installed for accessing RVC.

For detailed information about downloading the RVC tool and using the RVC commands, see the *RVC Command Reference Guide*.

■   Verify that your hosts are not in maintenance mode. When upgrading the on-disk format, do not place your hosts in maintenance mode. When any member host of a Virtual SAN cluster enters maintenance mode, the available resource capacity in the cluster is reduced because the member host no longer contributes capacity to the cluster and the cluster upgrade might fail.

■   Verify that there are no component rebuilding tasks currently in progress in the Virtual SAN cluster by running the RVC `vsan.resync_dashboard` command.

Procedure

1   Log in to your vCenter Server using RVC.

2   Run the `vsan.disks_stats /< vCenter IP address or hostname>/<data center name>/computers/<cluster name>` command to view the disk status.

For example: `vsan.disks_stats /192.168.0.1/BetaDC/computers/VSANCluster`

The command lists the names of all devices and hosts in the Virtual SAN cluster. The command also displays the current disk format and its health status. You can also check the current health of the devices in the **Health Status** column from the Disk Management page. For example, the device status appears as Unhealthy in the **Health Status** column for the hosts or disk groups that have failed devices.

3   Run the `vsan.ondisk_upgrade <path to vsan cluster>` command .

For example: `vsan.ondisk_upgrade /192.168.0.1/BetaDC/computers/VSANCluster`

4   Monitor the progress in RVC.

RVC upgrades one disk group at a time.

After the disk format upgrade has completed successfully, the following message appears.

`Done with disk format upgrade phase`

`There are n v1 objects that require upgrade Object upgrade progress: n upgraded, 0 left`

`Object upgrade completed: n upgraded`

`Done VSAN upgrade`

5   Run the `vsan.obj_status_report` command to verify that the object versions are upgraded to the new on-disk format.

Verify the Virtual SAN Disk Format Upgrade

After you finish upgrading the disk format, you must verify whether the Virtual SAN cluster is using the new on-disk format.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

The current disk format version appears in the Disk Format Version column. For example, if you are using disk format 2.0, it appears as version 2 in the Disk Format Version column. For on-disk format 3.0, the disk format version appears as version 3.

Verify the Virtual SAN Cluster Upgrade

The Virtual SAN cluster upgrade is not complete until you have verified that you are using the latest version of vSphere and Virtual SAN is available for use.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab, and verify that vSAN is listed.

◆   You also can navigate to your ESXi host and select **Summary** > **Configuration**, and verify that you are using the latest version of the ESXi host.

Using the RVC Upgrade Command Options

The `vsan.ondisk_upgrade` command provides various command options that you can use to control and manage the Virtual SAN cluster upgrade. For example, you can allow reduced redundancy to perform the upgrade when you have little free space available.

Run the `vsan.ondisk_upgrade --help` command to display the list of RVC command options.

Use these command options with the `vsan.ondisk_upgrade` command.

Table 10-3. Upgrade Command Options

Options	Description
`--hosts_and_clusters`	Use to specify paths to all host systems in the cluster or cluster's compute resources.
`--ignore-objects, -i`	Use to skip Virtual SAN object upgrade. You can also use this command option to eliminate the object version upgrade. When you use this command option, objects continue to use the current on-disk format version.
`--allow-reduced-redundancy, -a`	Use to remove the requirement of having a free space equal to one disk group during disk upgrade. With this option, virtual machines operate in a reduced redundancy mode during upgrade, which means certain virtual machines might be unable to tolerate failures temporarily and that inability might cause data loss. Virtual SAN restores full compliance and redundancy after the upgrade is completed.
`--force, -f`	Use to enable force-proceed and automatically answer all confirmation questions.
`--help, -h`	Use to display the help options.

For information about using the RVC commands, see the *RVC Command Reference Guide*.

Device Management in a Virtual SAN Cluster 11

You can perform various device management tasks in a Virtual SAN cluster. You can create hybrid or all-flash disk groups, enable Virtual SAN to claim devices for capacity and cache, enable or disable LED indicators on devices, mark devices as flash, mark remote devices as local, and so on.

This chapter includes the following topics:

■ "Managing Disk Groups And Devices," on page 99

■ "Working with Individual Devices," on page 101

Managing Disk Groups And Devices

When you enable Virtual SAN on a cluster, choose a disk-claiming mode to organize devices into groups.

Virtual SAN 6.6 and later releases have a uniform workflow for claiming disks across all scenarios. It groups all available disks by model and size, or by host. You must select which devices to use for cache and which to use for capacity.

Create a Disk Group on a Host

When you create disk groups, you must manually specify each host and each device to be used for the Virtual SAN datastore. You organize cache and capacity devices into disk groups.

To create a disk group, you define the disk group and individually select devices to include in the disk group. Each disk group contains one flash cache device and one or more capacity devices.

When you create a disk group, consider the ratio of flash cache to consumed capacity. Although the ratio depends on the requirements and workload of the cluster, consider using at least 10 percent of flash cache to consumed capacity ratio (not including replicas such as mirrors).

The Virtual SAN cluster initially contains a single Virtual SAN datastore with zero bytes consumed.

As you create disk groups on each host and add cache and capacity devices, the size of the datastore grows according to the amount of physical capacity added by those devices. Virtual SAN creates a single distributed Virtual SAN datastore using the local empty capacity available from the hosts added to the cluster.

If the cluster requires multiple flash cache devices, you must create multiple disk groups manually, because a maximum of one flash cache device is allowed per disk group.

NOTE If a new ESXi host is added to the Virtual SAN cluster, the local storage from that host is not added to the Virtual SAN datastore automatically. You have to manually create a disk group and add the devices to the disk group in order to use the new storage from the new ESXi host.

Claim Disks for the Virtual SAN Cluster

You can select multiple devices from your hosts, and Virtual SAN creates default disk groups for you.

When you add more capacity to the hosts or add new hosts with capacity to the Virtual SAN cluster, you can select the new devices to increase the capacity of the Virtual SAN datastore. In an all-flash cluster, you can mark flash devices for use as capacity.

After Virtual SAN has claimed devices, it creates the Virtual SAN shared datastore. The total size of the datastore reflects the capacity of all capacity devices in disk groups across all hosts in the cluster. Some capacity overhead is used for metadata.

Create a Disk Group on a Virtual SAN Host

You can manually combine specific cache devices with specific capacity devices to define disk groups on a particular host.

In this method, you manually select devices to create a disk group for a host. You add one cache device and at least one capacity device to the disk group.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3 Under vSAN, click **Disk Management**.

4 Select the host and click the **Create a new disk group** icon ( ).

   ■ Select the flash device to be used for cache.

   ■ From the **Capacity type** drop-down menu, select the type of capacity disks to use, depending on the type of disk group you want to create (HDD for hybrid or Flash for all-flash).

   ◆ Select the devices you want to use for capacity.

5 Click **OK**.

The new disk group appears in the list.

Claim Storage Devices for a Virtual SAN Cluster

You can select a group of cache and capacity devices, and Virtual SAN organizes them into default disk groups.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Click the **Claim Disks** icon ( ).

5   Select devices to add to the disk group.

- Each host that contributes storage to a hybrid disk group must contribute one flash cache device and one or more capacity devices. You can add only one flash cache device per disk group.

  - Select a flash device to be used as cache and click the **Claim for cache tier** icon ( ).

  - Select an HDD device to be used as capacity and click the **Claim for capacity tier** icon ( ).

  - Click **OK**.

- For all-flash disk groups, select flash devices for both capacity and cache.

  - Select a flash device to be used as cache and click the **Claim for cache tier** icon ( ).

  - Select a flash device to be used for capacity and click the **Claim for capacity tier** icon ( ).

  - Click **OK**.

To verify the role of each device added to the all-flash disk group, navigate to the Disk Role column at the bottom of the Disk Management page. The column shows the list of devices and their purpose in a disk group.

Virtual SAN claims the devices that you selected and organizes them into default disk groups that support the Virtual SAN datastore.

Working with Individual Devices

You can perform various device management tasks in the Virtual SAN cluster, such as adding devices to a disk group, removing devices from a disk group, enabling or disabling locator LEDs, and marking devices.
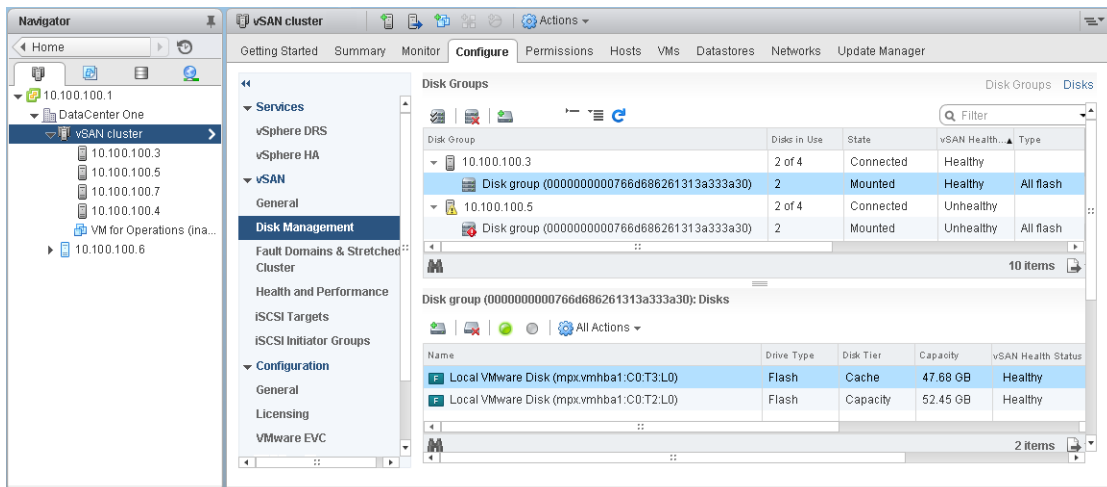
Add Devices to the Disk Group

When you configure Virtual SAN to claim disks in manual mode, you can add additional local devices to existing disk groups.

The devices must be the same type as the existing devices in the disk groups, such as SSD or magnetic disks.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Select the disk group, and click the **Add a disk to the selected disk group** icon (⊞).

5   Select the device that you want to add and click **OK**.

    If you add a used device that contains residual data or partition information, you must first clean the device. For information about removing partition information from devices, see "Remove Partition From Devices," on page 106. You can also run the `host_wipe_vsan_disks` RVC command to format the device. For more information about RVC commands, see the *RVC Command Reference Guide*.

Remove Disk Groups or Devices from Virtual SAN

You can remove selected devices from the disk group or an entire disk group.

Because removing unprotected devices might be disruptive for the Virtual SAN datastore and virtual machines in the datastore, avoid removing devices or disk groups.

Typically, you delete devices or disk groups from Virtual SAN when you are upgrading a device or replacing a failed device, or when you must remove a cache device. Other vSphere storage features can use any flash-based device that you remove from the Virtual SAN cluster.

Deleting a disk group permanently deletes the disk membership as well as the data stored on the devices.

NOTE   Removing one flash cache device or all capacity devices from a disk group removes the entire disk group.

Evacuating data from devices or disk groups might result in the temporary noncompliance of virtual machine storage policies.

Prerequisites

■   You can either place the Virtual SAN host in maintenance mode by selecting the **Evacuate all data** option or by selecting **Ensure data aaccessibility** when deleting a device or a disk group. If you select **No data evacuation** from the drop-down menu, your data might be at risk if a failure occurs during evacuation.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Remove a disk group or selected devices.

Option	Description
Remove the Disk Group	a Under Disk Groups, select the disk group to remove, and click the **Remove the disk group** icon (). b Select a data evacuation mode.
Remove the Selected Device	a Under Disk Groups, select the disk group that contains the device that you are removing. b Under Disks, select the device to remove, and click the **Remove the selected disk(s) from the disk group** icon (). c Select a data evacuation mode.

You can move the evacuated data to another disk or disk group on the same host.

5   Click **Yes** to confirm.

The data is evacuated from the selected devices or a disk group and is no longer available to Virtual SAN.

Using Locator LEDs

You can use locator LEDs to identify the location of storage devices.

Virtual SAN can light the locator LED on a failed device so that you can easily identify the device. This is particularly useful when you are working with multiple hot plug and host swap scenarios.

Consider using I/O storage controllers with pass-through mode, because controllers with RAID 0 mode require additional steps to enable the controllers to recognize locator LEDs.

For information about configuring storage controllers with RAID 0 mode, see your vendor documentation.

Enable and Disable Locator LEDs

You can turn locator LEDs on Virtual SAN storage devices on or off. When you turn on the locator LED, you can identify the location of a specific storage device.

When you no longer need a visual alert on your Virtual SAN devices, you can turn off locator LEDs on the selected devices.

Prerequisites

■   Verify that you have installed the supported drivers for storage I/O controllers that enable this feature. For information about the drivers that are certified by VMware, see the *VMware Compatibility Guide* at http://www.vmware.com/resources/compatibility/search.php.

■   In some cases, you might need to use third-party utilities to configure the Locator LED feature on your storage I/O controllers. For example, when you are using HP you should verify that the HP SSA CLI is installed.

For information about installing third-party VIBs, see the *vSphere Upgrade* documentation.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Select a host to view the list of devices.

5   At the bottom of the page, select one or more storage devices from the list, and enable or disable the locator LEDs on the selected devices.

Option	Action
Turns on the locator LED of the selected disk(s) icon	Enables locator LED on the selected storage device. You can enable locator LEDs from the **Manage** tab and click **Storage > Storage Devices**.
Turns off the locator LED of the selected disk(s) icon	Disables locator LED on the selected storage device. You can disable locator LEDs from the **Manage** tab and click **Storage > Storage Devices**.

Mark Devices as Flash

When flash devices are not automatically identified as flash by ESXi hosts, you can manually mark them as local flash devices.

This can also happen when they are enabled for RAID 0 mode rather than passthrough mode. When devices are not recognized as local flash, they are excluded from the list of devices offered for Virtual SAN and you cannot use them in the Virtual SAN cluster. Marking these devices as local flash makes them available to Virtual SAN.

Prerequisites

■   Verify that the device is local to your host.

■   Verify that the device is not in use.

■   Make sure that the virtual machines accessing the device are powered off and the datastore is unmounted.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Select the host to view the list of available devices.

5   From the **Show** drop-down menu at the bottom of the page, select **Not in Use**.

6   Select one or more flash devices from the list and click the **Mark the selected disks as flash disks** icon ( ).

7   Click **Yes** to save your changes.

The Drive type for the selected devices appear as Flash.

Mark Devices as HDD

When local magnetic disks are not automatically identified as HDD devices by ESXi hosts, you can manually mark them as local HDD devices.

If you marked a magnetic disk as a flash device, you can change the disk type of the device by marking it as a magnetic disk.

Prerequisites

■   Verify that the magnetic disk is local to your host.

■   Verify that the magnetic disk is not in use and is empty.

■   Verify that the virtual machines accessing the device are powered off.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3 Under vSAN, click **Disk Management**.

4 Select the host to view the list of available magnetic disks.

5 From the **Show** drop-down menu at the bottom of the page, select **Not in Use**.

6 Select one or more magnetic disks from the list and click **Mark the selected disks as HDD disks** icon
(HDD).

7 Click **Yes** to save.

The Drive Type for the selected magnetic disks appears as HDD.

Mark Devices as Local

When hosts are using external SAS enclosures, Virtual SAN might recognize certain devices as remote, and might be unable to automatically claim them as local.

In such cases, you can mark the devices as local.

Prerequisites

Make sure that the storage device is not shared.

Procedure

1 Browse to the Virtual SAN cluster in the vSphere Web Client navigator.

2 Click the **Configure** tab.

3 Under vSAN, click **Disk Management**.

4 Select a host to view the list of devices.

5 From the **Show** drop-down menu at the bottom of the page, select **Not in Use**.

6 From the list of devices, select one or more remote devices that you want to mark as local and click the **Mark the selected disks as local for the host** icon.

7 Click **Yes** to save your changes.

Mark Devices as Remote

Hosts that use external SAS controllers can share devices. You can manually mark those shared devices as remote, so that Virtual SAN does not claim the devices when it creates disk groups.

In Virtual SAN, you cannot add shared devices to a disk group.

Procedure

1 Browse to the Virtual SAN cluster in the vSphere Web Client navigator.

2 Click the **Configure** tab.

3 Under vSAN, click **Disk Management**.

4 Select a host to view the list of devices.

5 From the **Show** drop-down menu at the bottom of the page, select **Not in Use**.

6 Select one or more devices that you want to mark as remote and click the **Marks the selected disk(s) as remote for the host** icon.

7   Click **Yes** to confirm.

Add a Capacity Device

You can add a capacity device to an existing Virtual SAN disk group.

You cannot add a shared device to a disk group.

Prerequisites

Verify that the device is formatted and is not in use.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Select a disk group.

5   Click the **Add a disk to the selected disk group** icon (🖫) at the bottom of the page.

6   Select the capacity device that you want to add to the disk group.

7   Click **OK**.

The device is added to the disk group.

Remove Partition From Devices

You can remove partition information from a device so Virtual SAN can claim the device for use.

If you have added a device that contains residual data or partition information, you must remove all preexisting partition information from the device before you can claim it for Virtual SAN use. VMware recommends adding clean devices to disk groups.

When you remove partition information from a device, Virtual SAN deletes the primary partition that includes disk format information and logical partitions from the device.

Prerequisites

Verify that the device is not in use by ESXi as boot disk, VMFS datastore, or Virtual SAN.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Disk Management**.

4   Select a host to view the list of available devices.

5   From the **Show** drop-down menu at the bottom of the page, select **Ineligible**.

6   Select a device from the list, and click the **Erase partitions on the selected disks** icon (🖉).

7   Click **OK** to confirm.

The device is clean and does not include any partition information.

Expanding and Managing a Virtual SAN Cluster 12

After you have set up your Virtual SAN cluster, you can use the vSphere Web Client to add hosts and capacity devices, remove hosts and devices, and manage failure scenarios.

This chapter includes the following topics:

Expanding a Virtual SAN Cluster

You can expand an existing Virtual SAN cluster by adding hosts or adding devices to existing hosts, without disrupting any ongoing operations.

Use one of the following methods to expand your Virtual SAN cluster.

■ Add new ESXi hosts to the cluster that are configured using the supported cache and capacity devices. See "Add a Host to the Virtual SAN Cluster," on page 108. When you add a new device or a add a host with capacity, Virtual SAN does not automatically distribute data to the newly added device. To enable Virtual SAN to distribute data to the devices that are recently added, you must perform manual rebalance operation in the cluster by using the Ruby vSphere Console (RVC). See "Manual Rebalance," on page 142.

■ Move existing ESXi hosts to the Virtual SAN cluster by using host profile. See "Configuring Hosts Using Host Profile," on page 109. New cluster members add storage and compute capacity. You must manually create a subset of disk groups from the local capacity devices on the newly added host. See "Create a Disk Group on a Virtual SAN Host," on page 100.

   Verify that the hardware components, drivers, firmware, and storage I/O controllers that you plan on using are certified and listed in the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php. When adding capacity devices, make sure that the devices are unformatted and not partitioned, so that Virtual SAN can recognize and claim the devices.

■ Add new capacity devices to ESXi hosts that are cluster members. You must manually add the device to the disk group on the host. See "Add Devices to the Disk Group," on page 101.

Expanding Virtual SAN Cluster Capacity and Performance

If your Virtual SAN cluster is running out of storage capacity or when you notice reduced performance of the cluster, you can expand the cluster for capacity and performance.

■ Expand the storage capacity of your cluster either by adding storage devices to existing disk groups or by creating a new disk group. New disk groups require flash devices for the cache. For information about adding devices to disk groups, see "Add Devices to the Disk Group," on page 101. Adding capacity devices without increasing the cache might reduce your cache-to-capacity ratio to an unsupported level. See "Design Considerations for Flash Caching Devices in Virtual SAN," on page 26.

■ Improve the cluster performance by adding at least one cache device (flash) and one capacity device (flash or magnetic disk) to an existing storage I/O controller or to a new server host. You can add one or more servers with additional disk groups, which has the same performance impact after Virtual SAN completes a proactive rebalance in the Virtual SAN cluster.

Although compute-only hosts can exist in a Virtual SAN environment and consume capacity from other hosts in the cluster, add uniformly configured hosts to provide smooth operation.

For best results, add hosts configured with both cache and capacity devices. For information about adding devices to disk groups, see"Add Devices to the Disk Group," on page 101.

Add a Host to the Virtual SAN Cluster

You can add an ESXi host to a running Virtual SAN cluster without disrupting any ongoing operations. The host's resources become associated with the cluster.

Prerequisites

■ Verify that the resources, including drivers, firmware, and storage I/O controllers, are listed in the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php.

■ VMware recommends creating uniformly configured hosts in the Virtual SAN cluster, so you have an even distribution of components and objects across devices in the cluster. However, there might be situations where the cluster becomes unevenly balanced, particularly during maintenance or if you overcommit the capacity of the Virtual SAN datastore with excessive virtual machine deployments.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Right-click the cluster and select **Add Host**.

3 Enter the host name or IP address, and click **Next**.

4 Enter the user name and password associated with the host, and click **Next**.

5 View the summary information and click **Next**.

6 Assign a license key and click **Next**.

7 (Optional) Enable lockdown mode to prevent remote users from logging directly into the host.

   You can configure this option later by editing the Security Profile in host settings.

8 Select what to do with the host's virtual machines and resource pools.

   ■ **Put this host's virtual machines in the cluster's root resource pool**

vCenter Server removes all existing resource pools of the host. The virtual machines in the host's hierarchy are all attached to the root. Because share allocations are relative to a resource pool, you might have to manually change a virtual machine's shares, which destroys the resource pool hierarchy.

■ **Create a resource pool for this host's virtual machines and resource pools**

vCenter Server creates a top-level resource pool that becomes a direct child of the cluster and adds all children of the host to that new resource pool. You can type a name for that new top-level resource pool. The default is **Grafted from <host_name>**.

9   Review the settings and click **Finish**.

The host is added to the cluster.

Configuring Hosts Using Host Profile

When you have multiple hosts in the Virtual SAN cluster, you can reuse the profile of an existing Virtual SAN host and apply its profile settings to the rest of the hosts in the Virtual SAN cluster.

The host profile includes information about storage configuration, network configuration, and other characteristics of the host. Typically, if you are planning to create a cluster with a large number of hosts, such as 8, 16, 32, or 64 hosts, use the host profile feature to add more than one host at a time to the Virtual SAN cluster.

Prerequisites

■ Verify that the host is in maintenance mode.

■ Verify that the hardware components, drivers, firmware, and storage I/O controllers are listed in the VMware Compatibility Guide at http://www.vmware.com/resources/compatibility/search.php.

Procedure

1   Create a host profile.

    a   Navigate to the Host Profiles view.

    b   Click the **Extract Profile from a Host** icon (+).

    c   Select the host that you intend to use as the reference host and click **Next**.

        The selected host must be an active host.

    d   Type a name and description for the new profile and click **Next**.

    e   Review the summary information for the new host profile and click **Finish**.

        The new profile appears in the Host Profiles list.

2   Attach the host to the intended host profile.

    a   From the Profile list in the Host Profiles view, select the host profile to be applied to the Virtual SAN host.

    b   Click the **Attach/Detach Hosts and clusters to a host profile** icon (🖼).

    c   Select the host from the expanded list and click **Attach** to attach the host to the profile.

        The host is added to the Attached Entities list.

    d   Click **Next**.

    e   Click **Finish** to complete the attachment of the host to the profile.

3   Detach the referenced Virtual SAN host from the host profile.

When a host profile is attached to a cluster, the host or hosts within that cluster are also attached to the host profile. However, when the host profile is detached from the cluster, the association between the host or hosts in the cluster and that of the host profile remains intact.

a   From the Profile List in the Host Profiles view, select the host profile to be detached from a host or cluster.

b   Click the **Attach/Detach Hosts and clusters to a host profile** icon (🖳).

c   Select the host or cluster from the expanded list and click **Detach**.

d   Click **Detach All** to detach all the listed hosts and clusters from the profile.

e   Click **Next**.

f   Click **Finish** to complete the detachment of the host from the host profile.

4   Verify the compliance of the Virtual SAN host to its attached host profile and determine if any configuration parameters on the host are different from those specified in the host profile.

a   Navigate to a host profile.

The **Objects** tab lists all host profiles, the number of hosts attached to that host profile, and the summarized results of the last compliance check.

b   Click the **Check Host Profile Compliance** icon (🟢).

To view specific details about which parameters differ between the host that failed compliance and the host profile, click the **Monitor** tab and select the Compliance view. Expand the object hierarchy and select the non-compliant host. The parameters that differ are displayed in the Compliance window, below the hierarchy.

If compliance fails, use the Remediate action to apply the host profile settings to the host. This action changes all host profile-managed parameters to the values that are contained in the host profile attached to the host.

c   To view specific details about which parameters differ between the host that failed compliance and the host profile, click the **Monitor** tab and select the Compliance view.

d   Expand the object hierarchy and select the failing host.

The parameters that differ are displayed in the Compliance window, below the hierarchy.

5   Remediate the host to fix compliance errors on the host.

a   Select the **Monitor** tab and click **Compliance**.

b   Right-click the host or hosts to remediate and select **All vCenter Actions > Host Profiles > Remediate**.

You can update or change the user input parameters for the host profiles policies by customizing the host.

c   Click **Next**.

d   Review the tasks that are necessary to remediate the host profile and click **Finish**.

The host is part of the Virtual SAN cluster and its resources are accessible to the Virtual SAN cluster. The host can also access all existing Virtual SAN storage I/O policies in the Virtual SAN cluster.

Working with Maintenance Mode

Before you shut down, reboot, or disconnect a host that is a member of a Virtual SAN cluster, you must put the host in maintenance mode.

When working with maintenance mode, consider the following guidelines:

- When you place an ESXi host in maintenance mode, you must select a data evacuation mode, such as **Ensure data accessibility from other hosts** or **Evacuate all data to other hosts**.

- When any member host of a Virtual SAN cluster enters maintenance mode, the cluster capacity automatically reduces as the member host no longer contributes storage to the cluster.

- A virtual machine's compute resources might not reside on the host that is being placed in maintenance mode, and the storage resources for virtual machines might be located anywhere in the cluster.

- The **Ensure data accessibility** mode is faster than the **Evacuate all data** mode because the **Ensure data accessibility** migrates only the components from the hosts that are essential for running the virtual machines. When in this mode, if you encounter a failure, the availability of your virtual machine is affected. Selecting the **Ensure data accessibility** mode does not reprotect your data during failure and you might experience unexpected data loss.

- When you select the **Evacuate all data** mode, your data is automatically reprotected against a failure, if the resources are available and the **Primary level of failures to tolerate** set to 1 or more. When in this mode, all components from the host are migrated and, depending on the amount of data you have on the host, the migration could take longer. With **Evacuate all data** mode, your virtual machines can tolerate failures, even during planned maintenance.

- When working with a three-host cluster, you cannot place a server in maintenance mode with **Evacuate all data**. Consider designing a cluster with four or more hosts for maximum availability.

Before you place a host in maintenance mode, you must verify the following:

- If you are using **Evacuate all data** mode, verify that you have enough hosts and capacity available in the cluster to meet the **Primary level of failures to tolerate** policy requirements.

- Verify that you have enough flash capacity on the remaining hosts to handle any flash read cache reservations. You can run the `vsan.whatif_host_failures` RVC command to analyze the current capacity utilization per host and whether a single host failure could make the cluster run out of space, and impact the cluster capacity, cache reservation, and cluster components. For information about the RVC commands, see the *RVC Command Reference Guide*.

- Verify that you have enough capacity devices in the remaining hosts to handle stripe width policy requirements, if selected.

- Make sure you have enough free capacity on the remaining hosts to handle the amount of data that must be migrated from the host entering maintenance mode.
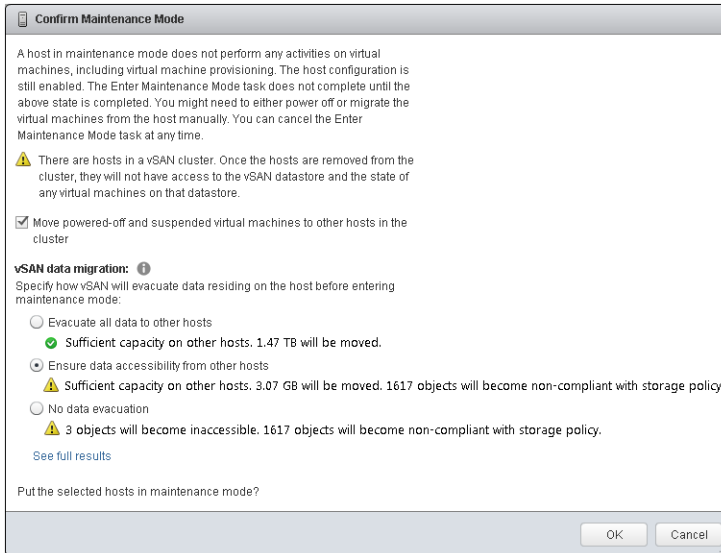
The Confirm Maintenance Mode dialog box provides information to guide your maintenance activities. You can view the impact of each data evacuation option.

■   Whether or not sufficient capacity is available to perform the operation.

■   How much data will be moved.

■   How many objects will become non-compliant.

■   How many objects will become inaccessible.

Place a Member of Virtual SAN Cluster in Maintenance Mode

Before you shut down, reboot, or disconnect a host that is a member of a Virtual SAN cluster, you must place the host in maintenance mode. When you place a host in maintenance mode, you must select a data evacuation mode, such as **Ensure data accessibility from other hosts** or **Evacuate all data to other hosts**.

When any member host of a Virtual SAN cluster enters maintenance mode, the cluster capacity is automatically reduced, because the member host no longer contributes capacity to the cluster.

Prerequisites

Verify that your environment has the capabilities required for the option you select.

Procedure

1   Right-click the host and select **Maintenance Mode > Enter Maintenance Mode**.

2   Select a data evacuation mode and click **OK**.

Option	Description
Ensure data accessibility from other hosts	This is the default option. When you power off or remove the host from the cluster, Virtual SAN ensures that all accessible virtual machines on this host remain accessible. Select this option if you want to take the host out of the cluster temporarily, for example, to install upgrades, and plan to have the host back in the cluster. This option is not appropriate if you want to permanently remove the host from the cluster.
	Typically, only partial data evacuation is required. However, the virtual machine might no longer be fully compliant to a VM storage policy during evacuation. That means, it might not have access to all its replicas. In case of a failure while the host is in maintenance mode and the **Primary level of failures to tolerate** is set to 1, you might experience data loss in the cluster.
	NOTE This is the only evacuation mode available if you are working with a three-host cluster or a Virtual SAN cluster configured with three fault domains.
Evacuate all data to other hosts	Virtual SAN evacuates all data to other hosts in the cluster, maintains or fixes availability compliance for the affected components, and protects data when sufficient resources exists in the cluster. Select this option if you plan to migrate the host permanently. When evacuating data from the last host in the cluster, make sure you migrate the virtual machines to another datastore and then place the host in maintenance mode.
	This evacuation mode results in the largest amount of data transfer and consumes the most time and resources. All of the components on the local storage of the selected host will be migrated elsewhere in the cluster so that when the host enters maintenance mode, all virtual machines will have access to their storage components and will still be compliant to their assigned storage policies.
	NOTE If a virtual machine object that has data on the host is not accessible and is not fully evacuated, the host will be unable to enter the maintenance mode.
No data evacuation	Virtual SAN does not evacuate any data from this host. If you power off or remove the host from the cluster, some virtual machines might become unaccessible.

A cluster with three fault domains has the same restrictions that a three-host cluster has, such as the inability to use **Evacuate all data** mode or to reprotect data after a failure.

What to do next

You can track the progress of data migration in the cluster. See

Managing Fault Domains in Virtual SAN Clusters

If your Virtual SAN cluster spans across multiple racks or blade server chassis in a data center and you want to make sure that your hosts are protected against rack or chassis failure, you can create fault domains and add one or more hosts to each fault domain.

A fault domain consists of one or more Virtual SAN hosts grouped together according to their physical location in the data center. When configured, fault domains enable Virtual SAN to tolerate failures of entire physical racks as well as failures of a single host, capacity device, network link or a network switch dedicated to a fault domain.

The **Primary level of failures to tolerate** policy for the cluster depends on the number of failures a virtual machine is provisioned to tolerate. For example, when a virtual machine is configured with the **Primary level of failures to tolerate** set to 1 (PFTT = 1) and using multiple fault domains, Virtual SAN can tolerate a single failure of any kind and of any component in a fault domain, including the failure of an entire rack.

When you configure fault domains on a rack and provision a new virtual machine, Virtual SAN ensures that protection objects, such as replicas and witnesses, are placed in different fault domains. For example, if a virtual machine's storage policy has the **Primary level of failures to tolerate** set to N (PFTT = n), Virtual SAN requires a minimum of 2*n+1 fault domains in the cluster. When virtual machines are provisioned in a cluster with fault domains using this policy, the copies of the associated virtual machine objects are stored across separate racks.

A minimum of three fault domains are required. For best results, configure four or more fault domains in the cluster. A cluster with three fault domains has the same restrictions that a three host cluster has, such as the inability to reprotect data after a failure and the inability to use the **Full data migration** mode. For information about designing and sizing fault domains, see "Designing and Sizing Virtual SAN Fault Domains," on page 34.

Consider a scenario where you have a Virtual SAN cluster with 16 hosts. The hosts are spread across 4 racks, that is, 4 hosts per rack. In order to tolerate an entire rack failure, you should create a fault domain for each rack. A cluster of such capacity can be configured to tolerate the **Primary level of failures to tolerate** set to 1. If you want to configure the cluster to allow for virtual machines with **Primary level of failures to tolerate** set to 2, you need to configure five fault domains in a cluster.

When a rack fails, all resources including the CPU, memory in the rack become unavailable to the cluster. To reduce the impact of a potential rack failure, you should configure fault domains of smaller sizes. This increases the total amount of resource availability in the cluster after a rack failure.

When working with fault domains, follow these best practices.

■ Configure a minimum of three fault domains in the Virtual SAN cluster. For best results, configure four or more fault domains.

■ A host not included in any fault domain is considered to reside in its own single-host fault domain.

■ You do not need to assign every Virtual SAN host to a fault domain. If you decide to use fault domains to protect the Virtual SAN environment, consider creating equal sized fault domains.

■ When moved to another cluster, Virtual SAN hosts retain their fault domain assignments.

■ When designing a fault domain, it is recommended that you configure fault domains with uniform number of hosts.

For guidelines about designing fault domains, see "Designing and Sizing Virtual SAN Fault Domains," on page 34.

■ You can add any number of hosts to a fault domain. Each fault domain must contain at least one host.

Create a New Fault Domain in Virtual SAN Cluster

To ensure that the virtual machine objects continue to run smoothly during a rack failure, you can group hosts in different fault domains.

When you provision a virtual machine on the cluster with fault domains, Virtual SAN distributes protection components, such as witnesses and replicas of the virtual machine objects across different fault domains. As a result, the Virtual SAN environment becomes capable of tolerating entire rack failures in addition to a single host, storage disk, or network failure.

Prerequisites

■ Choose a unique fault domain name. Virtual SAN does not support duplicate fault domain names in a cluster.

■ Verify the version of your ESXi hosts. You can only include hosts that are 6.0 or later in fault domains.

■ Verify that your Virtual SAN hosts are online. You cannot assign hosts to a fault domain that is offline or unavailable due to hardware configuration issue.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Click the **Create a new fault domain** icon (➕).

5   Type the fault domain name.

6   From the **Show** drop-down menu, select **Hosts not in fault domain** to view the list of hosts that are not assigned to a fault domain or select **Show All Hosts** to view all hosts in the cluster.

7   Select one or more hosts to add to the fault domain.

    A fault domain cannot be empty. You must select at least one host to include in the fault domain.

8   Click **OK.**

    The selected hosts appear in the fault domain.

Move Host into Selected Fault Domain

You can move a host into a selected fault domain in the Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select the fault domain and click the **Move hosts into selected fault domain** icon (🗂️).

5   From the **Show** drop-down menu at the bottom of the page, select **Hosts not in fault domain** to view the hosts that are available to be added to fault domains, or select **Show All Hosts** to view all hosts in the cluster.

6   Select the host that you want to add to the fault domain.

7   Click **OK**.

    The selected host appears in the fault domain.

Move Hosts into an Existing Fault Domain

You can move a host into an existing fault domain in the Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select one or more hosts and click the **Move hosts into fault domain** icon (🗂️).

5   Select a fault domain and click **OK**.

    Each fault domain must contain at least one host. If the host that you move is the only host in the source fault domain, Virtual SAN deletes the empty fault domain from the cluster.

Move Hosts out of a Fault Domain

Depending on your requirement, you can move hosts out of a fault domain.

Prerequisites

Verify that the host is online. You cannot move hosts that are offline or unavailable from a fault domain.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select the host that you want to move and click the **Move hosts out of fault domain** icon ( ).

5   Click **Yes**.

The selected host is no longer part of the fault domain. Any host that is not part of a fault domain is considered to reside in its own single-host fault domain.

What to do next

You can add hosts to fault domains. See "Move Hosts into an Existing Fault Domain," on page 115.

Rename a Fault Domain

You can change the name of an existing fault domain in your Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select the fault domain that you want to rename and click the **Rename selected fault domain** icon ( ).

5   Enter a new fault domain name.

6   Click **OK.**

The new name appears in the list of fault domains.

Remove Selected Fault Domains

When you no longer need a fault domain, you can remove it from the Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under vSAN, click **Fault Domains and Stretched Cluster**.

4   Select the fault domain that you want to delete and click the **Remove selected fault domains** icon ( ).

5   Click **Yes**.

All hosts in the fault domain are removed and the selected fault domain is deleted from the Virtual
SAN cluster. Each host that is not part of a fault domain is considered to reside in its own single-host
fault domain.

Using the Virtual SAN iSCSI Target Service

Use the iSCSI target service to enable hosts and physical workloads that reside outside the Virtual SAN
cluster to access the Virtual SAN datastore.

This feature enables an iSCSI initiator on a remote host to transport block-level data to an iSCSI target on a
storage device in the Virtual SAN cluster.

After you configure the Virtual SAN iSCSI target service, you can discover the Virtual SAN iSCSI targets
from a remote host. To discover Virtual SAN iSCSI targets, use the IP address of any host in the Virtual SAN
cluster, and the TCP port of the iSCSI target. To ensure high availability of the Virtual SAN iSCSI target,
configure multipath support for your iSCSI application. You can use the IP addresses of two or more hosts
to configure the multipath.

NOTE Virtual SAN iSCSI target service does not support other vSphere or ESXi clients or initiators, third
party hypervisors, or migrations using raw device mapping (RDMs).

Virtual SAN iSCSI target service supports the following CHAP authentication methods:

CHAP	In CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target.
Mutual CHAP	In mutual CHAP authentication, an additional level of security enables the initiator to authenticate the target.

iSCSI Targets

You can add one or more iSCSI targets that provide storage blocks as logical unit numbers (LUNs). Virtual
SAN identifies each iSCSI target by a unique iSCSI qualified Name (IQN). You can use the IQN to present
the iSCSI target to a remote iSCSI initiator so that the initiator can access the LUN of the target.

Each iSCSI target contains one or more LUNs. You define the size of each LUN, assign a Virtual SAN storage
policy to each LUN, and enable the iSCSI target service on a Virtual SAN cluster. You can configure a
storage policy to use as the default policy for the home object of the Virtual SAN iSCSI target service.

iSCSI Initiator Groups

You can define a group of iSCSI initiators that have access to a specified iSCSI target. The iSCSI initiator
group restricts access to only those initiators that are members of the group. If you do not define an iSCSI
initiator or initiator group, then each target is accessible to all iSCSI initiators.

A unique name identifies each iSCSI initiator group. You can add one or more iSCSI initiators as members of
the group. Use the IQN of the initiator as the member initiator name.

Enable the iSCSI Target Service

Before you can create iSCSI targets and LUNs and define iSCSI initiator groups, you must enable the iSCSI
target service on the Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab. Under Virtual SAN, click **General**.

3   Click the **Edit** button for Virtual SAN iSCSI Target Service.

4   Select the **Enable Virtual SAN iSCSI target service** check box. You can select the default network, TCP port, and Authentication method at this time. You also can select a Virtual SAN Storage Policy.

5   Click **OK**.

What to do next

After the iSCSI target service is enabled, you can create iSCSI targets and LUNs, and define iSCSI initiator groups.

Create an iSCSI Target

You can create or edit an iSCSI target and its associated LUN.

Prerequisites

Verify that the iSCSI target service is enabled.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab. Under Virtual SAN, click **iSCSI Targets**.

3   In the Virtual SAN iSCSI Targets section, click the **Add a new iSCSI target** (✚) icon.

    The New iSCSI Target dialog box is displayed. The target IQN is generated automatically.

4   Enter a target alias. You also can edit the network, TCP port, and authentication method for this target.

5   (Optional) To define the LUN for the target, click the **Add your first LUN to the iSCSI target** check box, and enter the size of the LUN.

6   Click **OK**.

What to do next

Define a list of iSCSI initiators that can access this target.

Add a LUN to an iSCSI Target

You can add one or more LUNs to an iSCSI target, or edit an existing LUN.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab. Under Virtual SAN, click **iSCSI Targets**.

3   Select the **LUNs** tab in the Target Details section of the page.

4   Click the **Add a new iSCSI LUN to the target** (✚) icon.

    The Add LUN to Target dialog box is displayed.

5   Enter the size of the LUN.

    The Virtual SAN Storage Policy configured for the iSCSI target service is assigned automatically. You can assign a different policy to each LUN.

6   Click **OK**.

Create an iSCSI Initiator Group

You can create an iSCSI initiator group to provide access control for iSCSI targets. Only iSCSI initiators that are members of the initiator group can access the iSCSI targets.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab. Under Virtual SAN, click **iSCSI Initiator Groups**.

3   In the Virtual SAN iSCSI Initiator Groups section, click the **Add a new iSCSI initiator group** (✚) icon.

    The New Virtual SAN iSCSI Initiator Group dialog box is displayed.

4   Enter a name for the iSCSI initiator group.

5   (Optional) To add members to the initiator group, enter the IQN of each member.

    Use the following format to enter the member IQN:

    iqn.YYYY-MM.domain:name

    Where:

    ■   YYYY = year, such as 2016

    ■   MM = month, such as 09

    ■   domain = domain where the initiator resides

    ■   name = member name (optional)

6   Click **OK**.

What to do next

Add members to the iSCSI initiator group.

Assign a Target to an iSCSI Initiator Group

You can assign an iSCSI target to an iSCSI initiator group. Only those initiators that are members of the initiator group can access the assigned targets.

Prerequisites

Verify that you have an existing iSCSI initiator group.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Configure** tab. Under Virtual SAN, click **iSCSI Initiator Groups**.

3   In the Group Details section, select the **Accessible Targets** tab.

4   Click the **Add a new accessible target for iSCSI Initiator group** (✚) icon.

    The Allow Access to Target For Initiator Group dialog box is displayed.

5   On the **Filter** tab, select a target from the list of available targets.

    The Selected Objects tab displays the currently selected targets.

6   Click **OK**.

Monitor Virtual SAN iSCSI Target Service

You can monitor the iSCSI target service to view the physical placement of iSCSI target components and to check for failed components. You also can monitor the health status of the iSCSI target service.

Prerequisites

Verify that you have enabled the Virtual SAN iSCSI target service and created targets and LUNs.

Procedure

1 Browse to the Virtual SAN cluster in the vSphere Web Client navigator.

2 Click **Monitor** and select **Virtual SAN**.

3 Click **iSCSI Targets**.

 iSCSI targets and LUNs are listed at the top of the page.

4 Click a target alias and view its status.

 The Physical Disk Placement tab at the bottom of the page shows where the data components of the target are located. The Compliance Failures tab shows failed components.

5 Click a LUN and view its status.

 The Physical Disk Placement tab at the bottom of the page shows where the data components of the target are located. The Compliance Failures tab shows failed components.

Migrate a Hybrid Virtual SAN Cluster to an All-Flash Cluster

You can migrate the disk groups in a hybrid Virtual SAN cluster to all-flash disk groups.

The Virtual SAN hybrid cluster uses magnetic disks for the capacity layer and flash devices for the cache layer. You can change the configuration of the disk groups in the cluster so that it uses flash devices on the cache layer and the capacity layer.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Remove the hybrid disk groups for each host in the cluster.

 a Click the **Configure** tab.

 b Under Virtual SAN, click **Disk Management**.

 c Under Disk Groups, select the disk group to remove, and click the **Remove the disk group** icon
 ().

 d Select **Full data migration** as a migration mode and click **Yes**.

3 Remove the physical HDD disks from the host.

4 Add the flash devices to the host.

 Verify that no partitions exist on the flash devices.

5 Create the all-flash disk groups on each host.

Power off a Virtual SAN Cluster

You can power off a Virtual SAN cluster.

Prerequisites

If the vCenter Server VM is running on the Virtual SAN cluster, migrate the VM to the first host, or record the host where it is currently running.

Procedure

1   Power off all virtual machines that are running on the Virtual SAN cluster.

The vCenter Server virtual machine must be powered off last.

2   Place all ESXi hosts that compose the cluster in maintenance mode.

Run the `esxcli` command to set the Virtual SAN mode for entering the maintenance state.

`esxcli system maintenanceMode set –e true –m noAction`

3   Power off the ESXi hosts.

Using Virtual SAN Policies

<div style="text-align: right; font-size: 3em; font-weight: bold;">13</div>

When you use Virtual SAN, you can define virtual machine storage requirements, such as performance and availability, in a policy. Virtual SAN ensures that each virtual machine deployed to Virtual SAN datastores is assigned at least one storage policy.

After they are assigned, the storage policy requirements are pushed to the Virtual SAN layer when a virtual machine is created. The virtual device is distributed across the Virtual SAN datastore to meet the performance and availability requirements.

Virtual SAN uses storage providers to supply information about underlying storage to the vCenter Server. This information helps you to make appropriate decisions about virtual machine placement, and to monitor your storage environment.

This chapter includes the following topics:

- "About Virtual SAN Policies," on page 123
- "View Virtual SAN Storage Providers," on page 126
- "About the Virtual SAN Default Storage Policy," on page 127
- "Assign a Default Storage Policy to Virtual SAN Datastores," on page 128
- "Define a Virtual Machine Storage Policy for Virtual SAN," on page 129

About Virtual SAN Policies

Virtual SAN storage polices define storage requirements for your virtual machines. These policies determine how the virtual machine storage objects are provisioned and allocated within the datastore to guarantee the required level of service.

When you enable Virtual SAN on a host cluster, a single Virtual SAN datastore is created and a default storage policy is assigned to the datastore.

When you know the storage requirements of your virtual machines, you can create a storage policy referencing capabilities that the datastore advertises. You can create several policies to capture different types or classes of requirements.

Each virtual machine deployed to Virtual SAN datastores is assigned at least one virtual machine storage policy. You can assign storage policies when you create or edit virtual machines.

NOTE   If you do not assign a storage policy to a virtual machine, Virtual SAN assigns a default policy. The default policy has **Primary level of failures to tolerate** set to 1, a single disk stripe per object, and a thin-provisioned virtual disk.

The VM swap object and the VM snapshot memory object do not adhere to the storage policies assigned to a VM. These objects are configured with **Primary level of failures to tolerate** set to 1. These objects might not have the same availability as other objects that have been assigned a policy with a different value for **Primary level of failures to tolerate**.

Table 13-1. Storage Policy Attributes

Capability	Description
Number of disk stripes per object	The minimum number of capacity devices across which each replica of a virtual machine object is striped. A value higher than 1 might 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 Virtual SAN environment has sufficient capacity devices present to accommodate the request.
Flash read cache reservation	Flash capacity reserved as read cache for the virtual machine object. 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 is shared fairly among all objects. Use this option only to address specific performance issues.
	You do not have to set a reservation to get cache. Setting read cache reservations might 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 supported only for hybrid configurations. You must not use this attribute when defining a VM storage policy for an all-flash cluster.
	Default value is 0%. Maximum value is 100%.
	NOTE By default, Virtual SAN dynamically allocates read cache to storage objects based on demand. This feature represents the most flexible and the most optimal use of resources. As a result, typically, you do not need to change the default 0 value for this parameter.
	To increase the value when solving a performance problem, exercise caution. Over-provisioned cache reservations across several virtual machines can cause flash device space to be wasted on over-reservations. These cache reservations are not available to service the workloads that need the required space at a given time. This space wasting and unavailability might lead to performance degradation.

Table 13-1. Storage Policy Attributes (Continued)

Capability	Description
Primary level of failures to tolerate	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 using RAID 5 or RAID 6. When provisioning a virtual machine, if you do not choose a storage policy, Virtual SAN assigns this policy as the default virtual machine storage policy. 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. **NOTE** If you do not want Virtual SAN to protect a single mirror copy of virtual machine objects, you can specify the **Primary level of failures to tolerate** = 0. However, the host might experience unusual delays when entering maintenance mode. The delays occur because Virtual SAN must evacuate the object from the host for the maintenance operation to complete successfully. Setting the **Primary level of failures to tolerate** = 0 means that your data is unprotected, and you might lose data when the Virtual SAN cluster encounters a device failure. **NOTE** If you create a storage policy and you do not specify a value for **Primary level of failures to tolerate**, Virtual SAN creates a single mirror copy of the VM objects. IT can tolerate a single failure. However, if multiple component failures occur, your data might be at risk. In a stretched cluster, defines the number of host and device failures that a virtual machine object can tolerate. You can use the **Primary level of failures to tolerate** with the **Secondary level of failures to tolerate** to provide local fault protection for objects within a single site.
Secondary level of failures to tolerate	In a stretched cluster, this rule defines the number of host and object failures that a virtual machine object can tolerate within a single site. Default value is 1. Maximum value is 3.
Affinity	In a stretched cluster, this rule is available only if the **Primary level of failures to tolerate** is set to 0. You can set the Affinity rule to **None**, **Preferred**, or **Secondary**. This rule enables you to limit virtual machine objects to a selected site in the stretched cluster. Default value is None.
Force provisioning	If the option is set to **Yes**, the object is provisioned even if the **Primary level 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 datastore. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is no longer possible. The default **No** is acceptable for most production environments. Virtual SAN 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	Percentage of the logical size of the virtual machine disk (vmdk) object that must be reserved, or thick provisioned when deploying virtual machines. Default value is 0%. Maximum value is 100%.

Table 13-1. Storage Policy Attributes (Continued)

Capability	Description
Disable object checksum	If the option is set to **No**, the object calculates checksum information to ensure the integrity of its data. If this option is set to **Yes**, the object does not calculate checksum information.
	Virtual SAN uses end-to-end checksum to 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, Virtual SAN repairs the data or reports the error.
	If a checksum mismatch is detected, Virtual SAN 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	Specifies whether the data replication method optimizes for Performance or Capacity. If you select **RAID-1 (Mirroring) - Performance**, Virtual SAN 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**, Virtual SAN 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 **Primary level 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 **Primary level of failures to tolerate** to 2.
	In stretched clusters with **Secondary level of failures to tolerate** configured, this rule is applies only to the **Secondary level of failures to tolerate**.
	For more information about RAID 5 or RAID 6, see "Using RAID 5 or RAID 6 Erasure Coding," on page 76.
IOPS limit for object	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 represents two I/O operations.
	When calculating IOPS, read and write are considered equivalent, but cache hit ratio and sequentiality are not considered. If a disk's IOPS exceeds the limit, I/O operations are throttled. If the **IOPS limit for object** is set to 0, IOPS limits are not enforced.
	Virtual SAN allows the object to double the rate of the IOPS limit during the first second of operation or after a period of inactivity.

When working with virtual machine storage policies, you must understand how the storage capabilities affect the consumption of storage capacity in the Virtual SAN cluster. For more information about designing and sizing considerations of storage policies, see Chapter 4, "Designing and Sizing a Virtual SAN Cluster," on page 23.

View Virtual SAN Storage Providers

Enabling Virtual SAN automatically configures and registers a storage provider for each host in the Virtual SAN cluster.

Virtual SAN storage providers are built-in software components that communicate datastore capabilities to vCenter Server. A storage capability is typically represented by a key-value pair, where the key is a specific property offered by the datastore. The value is a number or range that the datastore can provide for a provisioned object, such as a virtual machine home namespace object or a virtual disk. You can also use tags to create user-defined storage capabilities and reference them when defining a storage policy for a virtual machine. For information about how to apply and use tags with datastores, see the *vSphere Storage* documentation.

The Virtual SAN storage providers report a set of underlying storage capabilities to vCenter Server. They also communicate with the Virtual SAN layer to report the storage requirements of the virtual machines. For more information about storage providers, see the *vSphere Storage* documentation.

Virtual SAN registers a separate storage provider for each host in the Virtual SAN cluster, using the following URL:

http://*host_ip*:8080/version.xml

where *host_ip* is the actual IP of the host.

Verify that the storage providers are registered.

Procedure

1   Browse to vCenter Server in the vSphere Web Client navigator.

2   Click the **Configure** tab, and click **Storage Providers**.

The storage providers for Virtual SAN appear on the list. Each host has a storage provider, but only one storage provider is active. Storage providers that belong to other hosts are in standby. If the host that currently has the active storage provider fails, the storage provider for another host becomes active.

NOTE   You cannot manually unregister storage providers used by Virtual SAN. If you need to remove or unregister the Virtual SAN storage providers, remove corresponding hosts from the Virtual SAN cluster and then add the hosts back. Make sure at least one storage provider is active.

About the Virtual SAN Default Storage Policy

Virtual SAN requires that the virtual machines deployed on the Virtual SAN datastores are assigned at least one storage policy. When provisioning a virtual machine, if you do not explicitly assign a storage policy to the virtual machine the Virtual SAN Default Storage Policy is assigned to the virtual machine.

The default policy contains Virtual SAN rulesets and a set of basic storage capabilities, typically used for the placement of virtual machines deployed on Virtual SAN datastores.

Table 13-2. Virtual SAN Default Storage Policy Specifications

Specification	Setting
Primary level of failures to tolerate	1
Number of disk stripes per object	1
Flash read cache reservation, or flash capacity used for read cache	0
Object space reservation	0 **NOTE** Setting the Object space reservation to zero means that the virtual disk will be thin provisioned, by default.
Force provisioning	No

You can review the configuration settings for the default virtual machine storage policy from the vSphere Web Client when you navigate to the **VM Storage Policies** > **Virtual SAN Default Storage Policy** > **Manage** > **Rule-Set 1: VSAN**.

For best results, consider creating and using your own VM storage policies, even if the requirements of the policy are same as those defined in the default storage policy. For information about creating a user-defined VM storage policy, see "Define a Virtual Machine Storage Policy for Virtual SAN," on page 129.

When you assign a user-defined storage policy as the default policy to a datastore, Virtual SAN automatically removes the association to the default storage policy and applies the settings for the user-defined policy on the specified datastore. At any point, you can assign only one virtual machine storage policy as the default policy to the Virtual SAN datastore.

Characteristics

The following characteristics apply to the Virtual SAN Default Storage Policy.

■ The Virtual SAN default storage policy is assigned to all virtual machine objects if you do not select any other Virtual SAN policy when you provision a virtual machine, that is when the **VM Storage Policy** field is set to **Datastore default** on the Select Storage page. For more information about using storage policies, see the *vSphere Storage* documentation.

> **Note** VM swap and VM memory objects receive the Virtual SAN Default Storage Policy with **Force provisioning** set to **Yes**.

■ The Virtual SAN default policy only applies to Virtual SAN datastores. You cannot apply the default storage policy to non-Virtual SAN datastores, such as NFS or a VMFS datastore.

■ Because the default virtual machine storage policy is compatible with any Virtual SAN datastore in the vCenter Server, you can move your virtual machine objects provisioned with the default policy to any Virtual SAN datastore in the vCenter Server.

■ You can clone the default policy and use it as a template to create a user-defined storage policy.

■ You can edit the default policy, if you have the StorageProfile.View privilege. You must have at least one Virtual SAN enabled cluster that contains at least one host. VMware highly recommends that you do not edit the settings of the default storage policy.

■ You cannot edit the name and description of the default policy, or the Virtual SAN storage provider specification. All other parameters including the policy rules are editable.

■ You cannot delete the default policy.

■ The default storage policy is assigned when the policy that you assign during virtual machine provisioning does not include rules specific to Virtual SAN.

Assign a Default Storage Policy to Virtual SAN Datastores

You can assign a user-defined storage policy as the default policy to a datastore, to reuse a storage policy that matches your requirements.

Prerequisites

Verify that the VM storage policy you want to assign as the default policy to the Virtual SAN datastore meets the requirements of virtual machines in the Virtual SAN cluster.

Procedure

1   Navigate to the Virtual SAN datastore in the vSphere Web Client.

2   Click the **Configure** tab.

3   Under General, click the Default Storage Policy **Edit** button, and select the storage policy that you want to assign as the default policy to the Virtual SAN datastore.

The vSphere Web Client displays a list of storage policies that are compatible with the Virtual SAN datastore, such as the Virtual SAN Default Storage Policy and user-defined storage policies that have Virtual SAN rule sets defined.

4   Select a policy and click **OK**.

The storage policy is applied as the default policy when you provision new virtual machines without explicitly specifying a storage policy for a datastore.

What to do next

You can define a new storage policy for virtual machines. See "Define a Virtual Machine Storage Policy for Virtual SAN," on page 129.

Define a Virtual Machine Storage Policy for Virtual SAN

You can create a storage policy that defines storage requirements for a VM and its virtual disks. In this policy, you reference storage capabilities supported by the Virtual SAN datastore.

Prerequisites

■   Verify that the Virtual SAN storage provider is available. See "View Virtual SAN Storage Providers," on page 126.

■   Ensure that the virtual machine storage policies are enabled. For information about storage policies, see the *vSphere Storage* documentation.

■   Required privileges: **Profile-driven storage.Profile-driven storage view** and **Profile-driven storage.Profile-driven storage update**

Procedure

1   From the vSphere Web Client home, click **Policies and Profiles**, then click **VM Storage Policies**.

2   Click the **Create a new VM storage policy** icon ( ).

3   On the Name and description page, select a vCenter Server.

4   Type a name and a description for the storage policy, and click **Next**.

5   On the Policy structure page, click **Next**.

6   On the **Common rules for data services provided by hosts** page, click **Next**.

7   On the Rule-Set 1 page, define the first rule set.

    a   Select the **Use rule-sets in the storage policy** check box.

    b   Select **VSAN** from the **Storage type** drop-down list.

         The page expands as you add rules for the Virtual SAN datastore.

    c   Select a rule from the **Add rule** drop-down list.

         Make sure that the values you provide are within the range of values advertised by storage capabilities of the Virtual SAN datastore.

         From the Storage Consumption model, you can review the virtual disk size available for use and the corresponding flash cache and storage capacity, including the reserved storage space your virtual machines would potentially consume when you apply the specified storage policy.

    d   (Optional) Add tag-based capabilities.

8   (Optional) Click the **Add another rule set** button to add another rule set.

9   Click **Next**.

10   On the Storage compatibility page, review the list of datastores that match this policy and click **Next**.

    To be eligible, a datastore does not need to satisfy all rule sets within the policy. The datastore must satisfy at least one rule set and all rules within this set. Verify that the Virtual SAN datastore meets the requirements set in the storage policy and that it appears on the list of compatible datastores.

11   On the Ready to complete page, review the policy settings, and click **Finish**.

The new policy is added to the list.

What to do next

Assign this policy to a virtual machine and its virtual disks. Virtual SAN will place the virtual machine objects in accordance with the requirements specified in the policy. For information about applying the storage policies to virtual machine objects, see the *vSphere Storage* documentation.

Monitoring Virtual SAN

14

You can monitor your Virtual SAN environment from the vSphere Web Client.

You can monitor all of the objects in a Virtual SAN environment, including hosts that participate in a Virtual SAN cluster and the Virtual SAN datastore. For more information about monitoring objects and storage resources in a Virtual SAN cluster, see the *vSphere Monitoring and Performance* documentation.

This chapter includes the following topics:

Monitor the Virtual SAN Cluster

You can monitor the Virtual SAN cluster and all the objects related to it.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Monitor** tab and click **vSAN**.

3   Select **Physical Disks** to review all hosts, cache devices, and capacity devices in the cluster.

    Virtual SAN displays information about capacity devices, such as total capacity, used capacity, reserved capacity, functional status, physical location, and so on. The physical location is based on the hardware location of cache and capacity and devices on Virtual SAN hosts.

4   Select a capacity device and click **Virtual Disks** to review the virtual machines that use the device.

    You can monitor many aspects of virtual machine objects, including their current state and whether they are compliant with the storage policies assigned to them.

5   Select **Capacity** to review information about the amount of capacity provisioned and used in the cluster, and also to review a breakdown of the used capacity by object type or by data type.

6   Select the **Configure** tab and select **General** to check the status of the Virtual SAN cluster, verify Internet connectivity, and review the on-disk format used in the cluster.

Monitor Virtual SAN Capacity

You can monitor the capacity of the Virtual SAN datastore, deduplication and compression efficiency, and a breakdown of capacity usage.

The vSphere Web Client cluster Summary tab includes a summary of Virtual SAN capacity. You also can view more detailed information in the Capacity monitor.
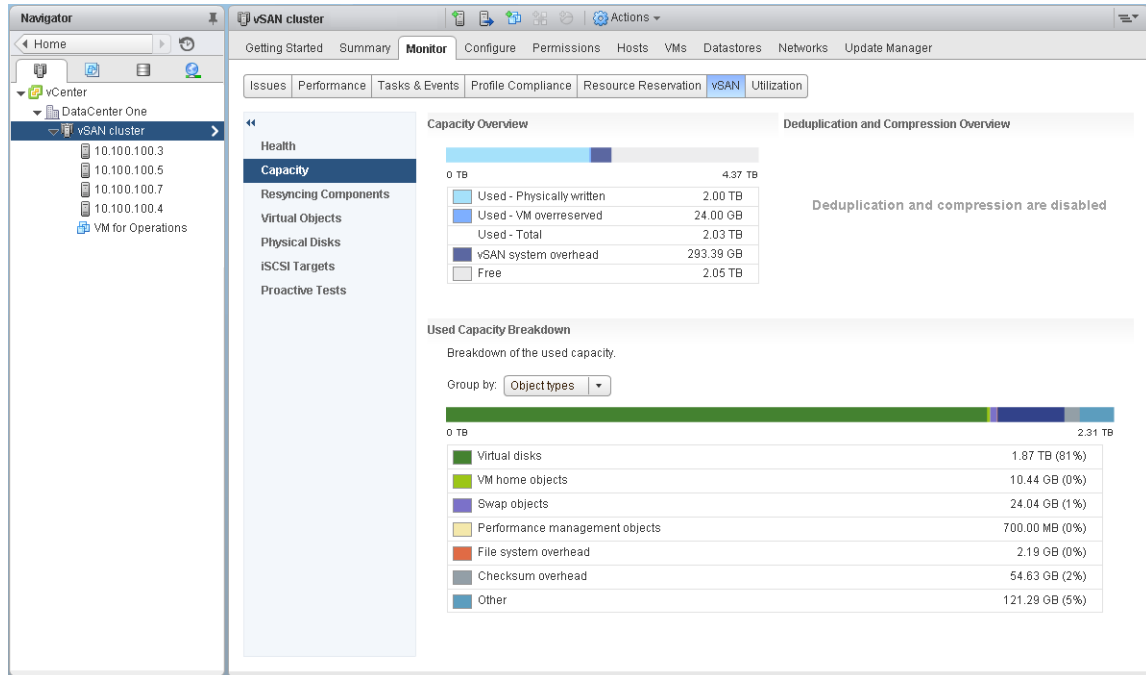
Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Monitor** tab and click **vSAN**.

3   Select **Capacity** to view Virtual SAN capacity information.

The Capacity Overview displays the storage capacity of the Virtual SAN datastore, including used space and free space. The Used Capacity Breakdown displays the percentage of capacity used by different object types or data types. If you select Data types, Virtual SAN displays the percentage of capacity used by primary VM data, Virtual SAN overhead, and temporary overhead. If you select Object types, Virtual SAN displays the percentage of capacity used by the following object types:

■   Virtual disks

■   VM home objects

■   Swap objects

■   Performance management objects

■   .vmem files

■   File system overhead

- Checksum overhead

- Snapshot memory

- Deduplication and compression overhead

- Space under dedup engine consideration

- iSCSI home and target objects, and iSCSI LUNs

- Other, such as user-created files, VM templates, and so on

If you enable deduplication and compression on the cluster, the Deduplication and Compression Overview displays capacity information related to that feature. When deduplication and compression are enabled, it might take several minutes for capacity updates to be reflected in the Capacity monitor as disk space is reclaimed and reallocated. For more information about deduplication and compression, see "Using Deduplication and Compression," on page 71.

Monitor Virtual Devices in the Virtual SAN Cluster

You can view the status of virtual disks in the Virtual SAN cluster.

When one or more hosts are unable to communicate with the Virtual SAN datastore, the information about virtual devices is not displayed.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Monitor** tab and click **vSAN**.

3   Select **Virtual Disks** to view all hosts and the corresponding virtual disks in the Virtual SAN cluster, including which hosts, cache and capacity devices their components are currently consuming.

4   Select the **VM home** folder on one of the virtual machines and click the **Physical Disk Placement** tab to view device information, such as name, identifier or UUID, and so on.

    Click the **Compliance Failures** tab to check the compliance status of your virtual machine.

5   Select **hard disk** on one of the virtual machines and click the **Physical Disk Placement** tab to view the device information, such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts.

    Click the **Compliance Failures** tab to check the compliance status of your virtual device.

6   Click the **Compliance Failures** tab to check the compliance status of your virtual machines.

About Virtual SAN Cluster Resynchronization

You can monitor the status of virtual machine objects that are being resynchronized in the Virtual SAN cluster.

When a hardware device, host, or network fails, or if a host is placed into maintenance mode, Virtual SAN initiates resynchronization in the Virtual SAN cluster. However, Virtual SAN might briefly wait for the failed components to come back online before initiating resynchronization tasks.

The following events trigger resynchronization in the cluster:

- Editing a virtual machine (VM) storage policy. When you change VM storage policy settings, Virtual SAN might initiate object recreation and subsequent resynchronization of the objects.

Certain policy changes might cause Virtual SAN to create another version of an object and synchronize it with the previous version. When the synchronization is complete, the original object is discarded.

Virtual SAN ensures that VMs continue to run and are not interrupted by this process. This process might require additional temporary capacity.

■ Restarting a host after a failure.

■ Recovering hosts from a permanent or long-term failure. If a host is unavailable for more than 60 minutes (by default), Virtual SAN creates copies of data to recover the full policy compliance.

■ Evacuating data by using the Full data migration mode before you place a host in maintenance mode.

■ Exceeding the utilization threshold of a capacity device. Resynchronization is triggered when capacity device utilization in the Virtual SAN cluster approaches or exceeds the threshold level of 80 percent.

If a VM is not responding due to latency caused by resynchronization, you can throttle the IOPS used for resynchronization.

Monitor the Resynchronization Tasks in the Virtual SAN Cluster

To evaluate the status of objects that are being resynchronized, you can monitor the resynchronization tasks that are currently in progress.

Prerequisites

Verify that hosts in your Virtual SAN cluster are running ESXi 6.5 or later.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Select the **Monitor** tab and click **vSAN**.

3 Select **Resyncing Components** to track the progress of resynchronization of virtual machine objects and the number of bytes that are remaining before the resynchronization is complete.

You can also view information about the number of objects that are currently being synchronized in the cluster, the estimated time to finish the resynchronization, the time remaining for the storage objects to fully comply with the assigned storage policy, and so on.

If your cluster has connectivity issues, the data on the Resyncing Components page might not get refreshed as expected and the fields might reflect inaccurate information.

Throttle Resynchronization Activity in the Virtual SAN Cluster

You can reduce the number of IOPS used to perform resynchronization on disk groups in the Virtual SAN cluster. Resynchronization throttling is a cluster-wide setting, and it is applied on a per disk group basis.

If VMs are not responding due to latency caused by resynchronization, you can throttle the number of IOPS used for resynchronization. Consider resynchronization throttling only if latencies are rising in the cluster due to resynchronization, or if resynchronization traffic is too high on a host.

Resynchronization throttling can increase the time required to complete resynchronization. Reprotection of non-compliant VMs might be delayed.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Select the **Monitor** tab and click **vSAN**.

3 Select **Resyncing Components** and click **Resync Throttling**.

4 (Optional) Click **Show current resync traffic per host** to view resynchronization activity.

5   Select the **Enable throttling for resyncing components traffic** check box.

6   Move the slider to set the throttle, as follows:

- Move the slider to the right to increase the amount of IOPS allowed for resynchronization.

- Move the slider to the left to decrease the amount of IOPS allowed for resynchronization.

A general rule is to throttle the IOPS by half and allow some time for the cluster to adapt. If further action is needed, throttle the IOPS by half again until the cluster stabilizes.

7   Click **OK**.

Monitor Devices that Participate in Virtual SAN Datastores

Verify the status of the devices that back up the Virtual SAN datastore. You can check whether the devices experience any problems.

Procedure

1   Navigate to Storage in the vSphere Web Client.

2   Select the Virtual SAN datastore.

3   Click the **Configure** tab.

You can view general information about the Virtual SAN datastore, including capacity, capabilities, and the default storage policy.

4   Click **Device Backing** and select the disk group to display local devices in the Disks table at the bottom of the page.

5   To display columns that are not visible, right-click the column heading and select **Show/Hide Columns**.

6   Select the columns that you want to display and click **OK**.

The selected columns are displayed in the Disks table.

Monitoring Virtual SAN Health

You can check the health of the Virtual SAN cluster.

You can use the Virtual SAN health checks to monitor the status of cluster components, diagnose issues, and troubleshoot problems. The health checks cover hardware compatibility, network configuration and operation, advanced Virtual SAN configuration options, storage device health, and virtual machine objects.
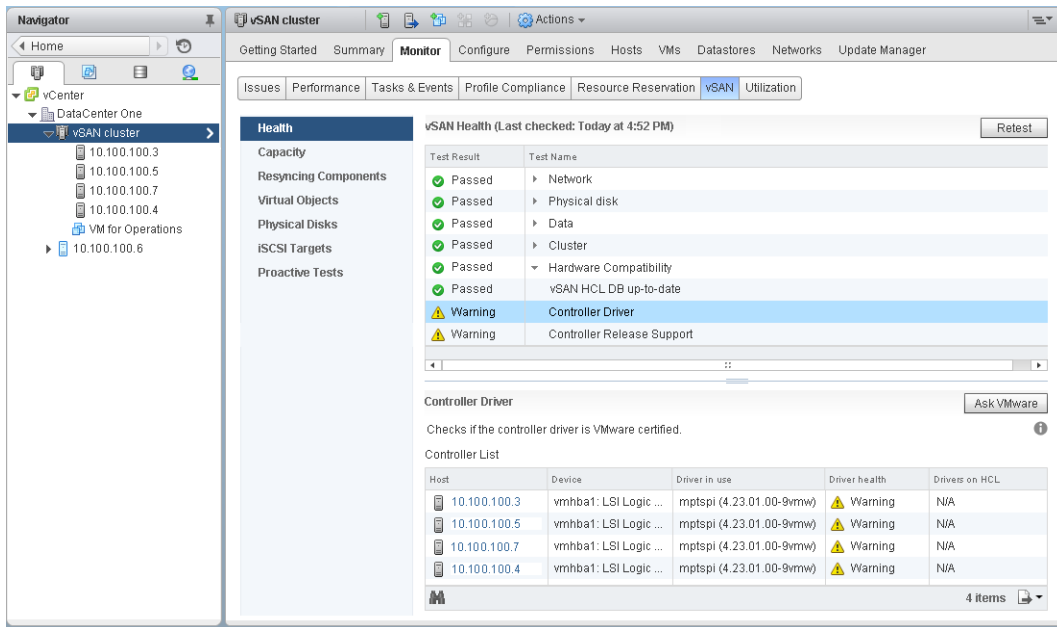
The Virtual SAN health checks are divided into categories. Each category contains individual health checks.

Table 14-1. Virtual SAN Health Check Categories

Health Check Category	Description
Hardware Compatibility	Monitor the cluster components to ensure that they are using supported hardware, software, and drivers.
Performance Service	Monitor the health of a Virtual SAN performance service.
Network	Monitor Virtual SAN network health.
Physical disk	Monitor the health of physical devices in the cluster.
Data	Monitor Virtual SAN data health.
Cluster	Monitor Virtual SAN cluster health.
Limits	Monitor Virtual SAN cluster limits.
Online health	Monitor Virtual SAN cluster health and send to VMware's analytics backend system for advanced analysis. You must participate in the Customer Experience Improvement Program to use online health checks.
Virtual SAN iSCSI target service	Monitor the iSCSI target service, including the network configuration and runtime status.
Encryption	Monitor Virtual SAN encryption health.
Stretched cluster	Monitor the health of a stretched cluster, if applicable.

Virtual SAN periodically retests each health check and updates the results. To run the health checks and update the results immediately, click the **Retest** button.

If you participate in the Customer Experience Improvement Program, you can run health checks and send the data to VMware for advanced analysis. Click the **Retest with Online health** button.

For more information about Virtual SAN health checks, see *VMware Virtual SAN Health Check Plugin Guide*.

Monitoring Virtual SAN Health on a Host

The ESXi host client is a browser-based interface for managing a single ESXi host. It enables you to manage the host when vCenter Server is not available. The host client provides tabs for managing and monitoring Virtual SAN at the host level.

- The **vSAN** tab displays basic Virtual SAN configuration.

- The **Hosts** tab displays the hosts participating in the Virtual SAN cluster.

- The **Health** tab displays host-level health checks.

Configure Virtual SAN Health Service

You can configure the health check interval for the Virtual SAN health service.

The Virtual SAN health service is turned on by default. You can turn periodical health checks off or on, and set the health-check interval.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Configure** tab.

3 Under vSAN, select **Health and Performance**.

4 Click the Health Services **Edit settings** button.

　a To turn off periodical health checks, deselect **Turn ON periodical health check**.

　　You also can set the time interval between health checks.

　b To turn on periodical health checks, select **Turn ON periodical health check**.

Check Virtual SAN Health

You can view the status of Virtual SAN health checks to verify the configuration and operation of your Virtual SAN cluster.

Prerequisites

Virtual SAN health service must be turned on before you can view the health checks.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click the **Monitor** tab and click **vSAN**.

3 Select **Health** to review the Virtual SAN health check categories.

　If the Test Result column displays Warning or Failed, expand the category to review the results of individual health checks.

4 Select an individual health check and check the detailed information at the bottom of the page.

　You can click the **Ask VMware** button to open a knowledge base article that describes the health check and provides information about how to resolve the issue.

Monitor Virtual SAN from ESXi Host Client

You can monitor Virtual SAN health and basic configuration through the ESXi host client.

Prerequisites

Virtual SAN health service must be turned on before you can view the health checks.

Procedure

1   Open a browser and enter the IP address of the host.

    The browser redirects to the login page for the host client.

2   Enter the username and password for the host, and click **Login**.

3   In the host client navigator, click **Storage**.

4   In the main page, click the vSAN datastore to display the Monitor link in the navigator.

5   Click the tabs to view Virtual SAN information for the host.

    a   Click the **vSAN** tab to display basic Virtual SAN configuration.

    b   Click the **Hosts** tab to display the hosts participating in the Virtual SAN cluster.

    c   Click the **Health** tab to display host-level health checks.

6   (Optional) On the **vSAN** tab, click **Edit Settings** to correct configuration issues at the host level. Select the values that match the configuration of your Virtual SAN cluster.

    Select the values that match the configuration of your Virtual SAN cluster, and click **Save**.

Monitoring Virtual SAN Performance

You can use Virtual SAN performance service to monitor the performance of your Virtual SAN environment, and investigate potential problems.

The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.
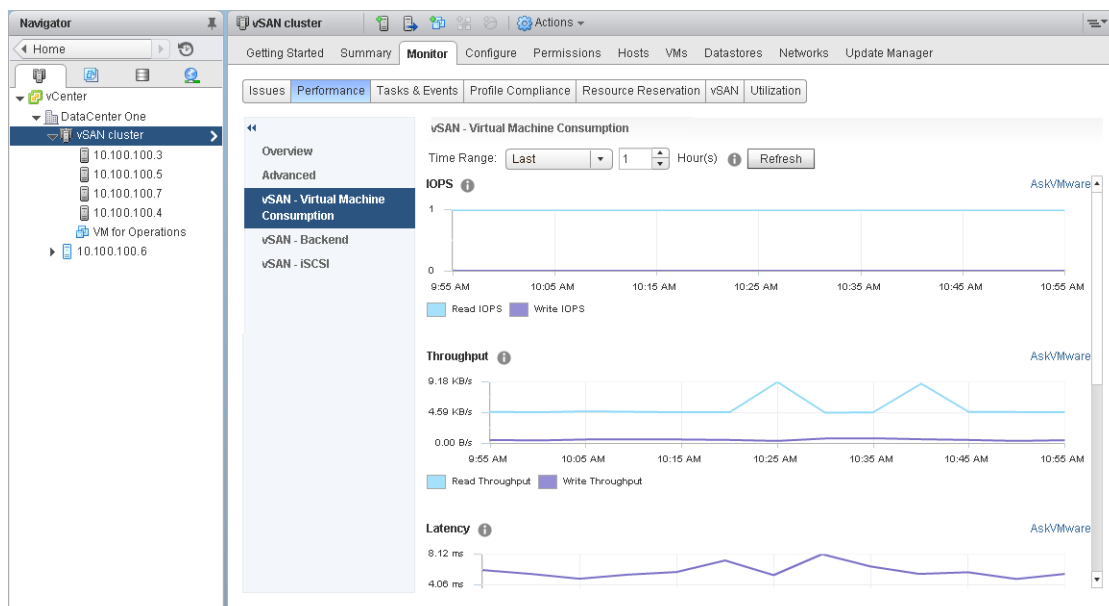
When the Virtual SAN performance service is turned on, the cluster summary displays an overview of Virtual SAN performance statistics, including IOPS, throughput, and latency. You can view detailed performance statistics for the cluster, and for each host, disk group, and disk in the Virtual SAN cluster. You also can view performance charts for virtual machines and virtual disks.

Turn on Virtual SAN Performance Service

When you create a Virtual SAN cluster, the performance service is disabled. Turn on Virtual SAN performance service to monitor the performance of Virtual SAN clusters, hosts, disks, and VMs.

When you turn on the performance service, Virtual SAN places a Stats database object in the datastore to collect statistical data. The Stats database is a namespace object in the cluster's Virtual SAN datastore.

Prerequisites

- All hosts in the Virtual SAN cluster must be running ESXi 6.5 or later.

- Before you enable the Virtual SAN performance service, make sure that the cluster is properly configured and has no unresolved health problems.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client navigator.

2   Click the **Configure** tab.

3   Under vSAN, select **Health and Performance**.

4   Click **Edit** to edit the performance service settings.

5   Select the **Turn On vSAN performance service** check box.

    You can turn off the Virtual SAN performance service by deselecting the check box.

6   Select a storage policy for the Stats database object and click **OK**.

Use Saved Time Range

You can select saved time ranges from the time range picker in performance views.

You can manually save a time range with customized name. When you run a storage performance test, the selected time range is saved automatically. You can save a time range for any of the performance views.

Prerequisites

- The Virtual SAN performance service must be turned on.

- All hosts in the Virtual SAN cluster must be running ESXi 6.6 or later.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client Navigator.

2   Click the **Monitor** tab and click **Performance**.

3   Select any tab, such as **vSAN - Backend**. In the time range drop-down, select **Save time range...**

4   Enter a name for the selected time range.

5   Click **OK**.

View Virtual SAN Cluster Performance

You can use the Virtual SAN cluster performance charts to monitor the workload in your cluster and determine the root cause of problems.

When the performance service is turned on, the cluster summary displays an overview of Virtual SAN performance statistics, including Virtual SAN IOPS, throughput, and latency. At the cluster level, you can view detailed statistical charts for virtual machine consumption and the Virtual SAN back end.

Prerequisites

The Virtual SAN performance service must be turned on before you can view performance charts.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client Navigator.

2   Click the **Monitor** tab and click **Performance**.

3   Select **vSAN - Virtual Machine Consumption**. Select a time range for your query.

Virtual SAN displays performance charts for clients running on the cluster, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster.

4   Select **vSAN - Backend**. Select a time range for your query.

Virtual SAN displays performance charts for the cluster back-end operations, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster.

5   Select **vSAN - iSCSI** and select an iSCSI target or LUN. Select a time range for your query.

NOTE   To view iSCSI performance charts, all hosts in the Virtual SAN cluster must be running ESXi 6.6 or later.

Virtual SAN displays performance charts for iSCSI targets or LUNs, including IOPS, bandwidth, latency, and outstanding I/O.

View Virtual SAN Host Performance

You can use the Virtual SAN host performance charts to monitor the workload on your hosts and determine the root cause of problems. You can view Virtual SAN performance charts for hosts, disk groups, and individual storage devices.

When the performance service is turned on, the host summary displays performance statistics for each host and its attached disks. At the host level, you can view detailed statistical charts for virtual machine consumption and the Virtual SAN back end, including IOPS, throughput, latency, and congestion. Additional charts are available to view the local client cache read IOPS and hit rate. At the disk group level, you can view statistics for the disk group. At the disk level, you can view statistics for an individual storage device.

Prerequisites

The Virtual SAN performance service must be turned on before you can view performance charts.

To view the following performance charts, hosts in the Virtual SAN cluster must be running ESXi 6.6 or later: Physical Adapters, VMkernal Adapters, VMkernal Adapters Aggregation, iSCSI, vSAN - Backend resync I/Os, resync IOPS, resync throughput, Disk Group resync latency.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client Navigator, and select a host.

2   Click the **Monitor** tab and click **Performance**.

3   Select **vSAN - Virtual Machine Consumption**. Select a time range for your query.

    Virtual SAN displays performance charts for clients running on the host, including IOPS, throughput, latency, congestions, and outstanding I/Os.

4   Select **vSAN - Backend**. Select a time range for your query.

    Virtual SAN displays performance charts for the host back-end operations, including IOPS, throughput, latency, congestions, outstanding I/Os, and resync I/Os.

5   Select **vSAN - Disk Group**, and select a disk group. Select a time range for your query.

    Virtual SAN displays performance charts for the disk group, including front end (Guest) IOPS, throughput, and latency, as well as overhead IOPS and latency. It also displays the read-cached hit rate, evictions, write-buffer free percentage, capacity and usage, cache disk destage rate, congestions, outstanding I/O, outstanding I/O size, delayed I/O percentage, delayed I/O average latency, internal queue IOPS, internal queue throughput, resync IOPS, resync throughput, and resync latency.

6   Select **vSAN - Disk**, and select a disk. Select a time range for your query.

    Virtual SAN displays performance charts for the disk, including a physical/firmware layer IOPS, throughput, and latency.

7   Select **vSAN - Physical Adapters**, and select a NIC. Select a time range for your query.

    Virtual SAN displays performance charts for the physical NIC (pNIC), including throughput, packets per second, and packets loss rate.

8   Select **vSAN - VMkernel Adapters**, and select a VMkernel adapter, such as vmk1. Select a time range for your query.

    Virtual SAN displays performance charts for the VMkernel adapter, including throughput, packets per second, and packets loss rate.

9   Select **vSAN - VMkernel Adapters Aggregation**. Select a time range for your query.

    Virtual SAN displays performance charts for all network I/Os processed in the network adapters used by Virtual SAN, including throughput, packets per second, and packets loss rate.

10 Select **vSAN - iSCSI**. Select a time range for your query.

    Virtual SAN displays performance charts for all the iSCSI services on the host, including IOPS, bandwidth, latency, and outstanding I/Os.

View Virtual SAN VM Performance

You can use the Virtual SAN VM performance charts to monitor the workload on your virtual machines and virtual disks.

When the performance service is turned on, you can view detailed statistical charts for virtual machine performance and virtual disk performance. VM performance statistics cannot be collected during migration between hosts, so you might notice a gap of several minutes in the VM performance chart.

NOTE   The performance service supports only virtual SCSI controllers for virtual disks. Virtual disks using other controllers, such as IDE, are not supported.

Prerequisites

The Virtual SAN performance service must be turned on before you can view performance charts.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client Navigator, and select a VM.

2   Click the **Monitor** tab and click **Performance**.

3   Select **vSAN - Virtual Machine Consumption**. Select a time range for your query.

   Virtual SAN displays performance charts for the VM, including IOPS, throughput, and latency.

4   Select **vSAN - Virtual Disk**. Select a time range for your query.

   Virtual SAN displays performance charts for the virtual disks, including IOPS, delayed normalized IOPS, virtual SCSI IOPS, virtual SCSI throughput, and virtual SCSI latency.

About Virtual SAN Cluster Rebalancing

When any capacity device in your cluster reaches 80 percent utilization, Virtual SAN automatically rebalances the cluster, until the utilization of all capacity devices is below the threshold.

Cluster rebalancing evenly distributes resources across the cluster to maintain consistent performance and availability.

Other operations can initiate cluster rebalancing:

■   If Virtual SAN detects hardware failures on the cluster

■   If Virtual SAN hosts are placed in maintenance mode with the **Evacuate all data** option

■   If Virtual SAN hosts are placed in maintenance mode with **Ensure data accessibility** when objects assigned PFTT=0 reside on the host.

NOTE   To provide enough space for maintenance and reprotection, and to minimize automatic rebalancing events in the Virtual SAN cluster, consider keeping 30-percent capacity available at all times.

You can manually rebalance the Virtual SAN cluster by using the Ruby vSphere Console (RVC). See

Automatic Rebalance

By default, Virtual SAN automatically rebalances the Virtual SAN cluster when a capacity device reaches 80 percent utilization. Rebalancing also occurs when you place a Virtual SAN host in maintenance mode.

Run the following RVC commands to monitor the rebalance operation in the cluster:

■   `vsan.check_limits`. Verifies whether the disk space utilization is balanced in the cluster.

■   `vsan.whatif_host_failures`. Analyzes the current capacity utilization per host, interprets whether a single host failure can force the cluster to run out of space for reprotection, and analyzes how a host failure might impact cluster capacity, cache reservation, and cluster components.

   The physical capacity usage shown as the command output is the average usage of all devices in the Virtual SAN cluster.

■   `vsan.resync_dashboard`. Monitors any rebuild tasks in the cluster.

For information about the RVC command options, see the *RVC Command Reference Guide.*

Manual Rebalance

You can manually rebalance through the cluster health check, or by using RVC commands.

If the Virtual SAN disk balance health check fails, you can initiate a manual rebalance in the vSphere Web Client. Under Cluster health, access the vSAN Disk Balance health check, and click the **Rebalance Disks** button.

Use the following RVC commands to manually rebalance the cluster:

- `vsan.check_limits`. Verifies whether any capacity device in the Virtual SAN cluster is approaching the 80 percent threshold limit.

- `vsan.proactive_rebalance [opts]<Path to ClusterComputeResource> --start`. Manually starts the rebalance operation. When you run the command, Virtual SAN scans the cluster for the current distribution of components, and begins to balance the distribution of components in the cluster. Use the command options to specify how long to run the rebalance operation in the cluster, and how much data to move each hour for each Virtual SAN host. For more information about the command options for managing the rebalance operation in the Virtual SAN cluster, see the *RVC Command Reference Guide*.

  Because cluster rebalancing generates substantial I/O operations, it can be time-consuming and can affect the performance of virtual machines.

NOTE   When you manually rebalance the disks, the operation runs for the selected time period, until no more data needs to be moved. The default time period is 24 hours. If no data is being moved, Virtual SAN ends the rebalancing task.

You can configure an alarm that notifies you when the provisioned space reaches a certain threshold. See "Creating a vCenter Server Alarm for a Virtual SAN Event," on page 146.

Balance the Disk Use in the Virtual SAN Cluster

If your Virtual SAN cluster becomes unbalanced, you can rebalance the disk use.

If you remove capacity devices from the Virtual SAN cluster and add new capacity devices, the Virtual SAN cluster might become unbalanced from a capacity utilization perspective. After Virtual SAN health monitoring warns you about any unbalances, you can rebalance your cluster.

Prerequisites

Perform the rebalance operation during non-production hours to avoid excessive impact on the cluster.

Procedure

1   Navigate to the Virtual SAN cluster in the vSphere Web Client.

2   Click the **Monitor** tab and click **vSAN**.

3   Click **Health**.

4   In the vSAN health service table, select **Warning: Virtual SAN Disk Balance**.

    You can review the disk balance of the hosts.

5   Click the **Rebalance Disks** button to rebalance your cluster.

This operation moves components from the over-utilized disks to the under-utilized disks.

Using the Virtual SAN Default Alarms

You can use the default Virtual SAN alarms to monitor the cluster, hosts, and existing Virtual SAN licenses.

The default alarms are automatically triggered when the events corresponding to the alarms are activated or if one or all of the conditions specified in the alarms are met. You cannot edit the conditions or delete the default alarms. To configure alarms that are specific to your requirements, create custom alarms for Virtual SAN. See "Creating a vCenter Server Alarm for a Virtual SAN Event," on page 146.

The table lists the default Virtual SAN alarms.

Table 14-2. Default Virtual SAN Alarms

Virtual SAN Alarms	Description
Expired Virtual SAN time-limited license	Monitor Virtual SAN trial licenses.
Registration/unregistration of a VASA Vendor provider on a Virtual SAN hosts failed	Register or unregister VASA providers on the failed Virtual SAN hosts.
Expired Virtual SAN license	Monitor expired Virtual SAN licenses.
Errors occurred on the disk(s) of a Virtual SAN host	Monitor errors on Virtual SAN devices.
Virtual SAN Health Service Alarm for Group Test 'Cluster health'	Monitor Virtual SAN cluster health.
Virtual SAN Health Service Alarm for Group Test 'Data health'	Monitor Virtual SAN cluster data health.
Virtual SAN Health Service Alarm for Group Test 'Limits health'	Monitor Virtual SAN cluster limits.
Virtual SAN Health Service Alarm for Group Test 'Network health'	Monitor Virtual SAN network health.
Virtual SAN Health Service Alarm for Group Test 'Physical disk health'	Monitor the health of physical devices in the cluster.
Virtual SAN Health Service Alarm for Group Test 'Virtual SAN HCL health'	Monitor the cluster components to ensure they are using supported hardware, software and drivers.
Virtual SAN Health Service Alarm for Group Test 'software state health'	Monitor the health of the software currently used in the cluster.
Virtual SAN Health Service Alarm for Group Test 'Unexpected Virtual SAN health'	Monitor any unexpected cluster health issues.
Virtual SAN Health Service Alarm for Group Test 'Virtual SAN CLOMD liveness'	Monitor that the CLOMD (Cluster Level Object Manager Daemon), which runs on ESXi hosts and is responsible for data moves and evacuations, is alive or not.
Virtual SAN Health Service Alarm for Group Test 'Virtual SAN cluster partition'	Monitor Virtual SAN cluster partition.

For information about monitoring alarms, events, and editing existing alarm settings, see the *vSphere Monitoring and Performance* documentation.

View Virtual SAN Default Alarms

Use the default Virtual SAN alarms to monitor your cluster, hosts, analyze any new events, and assess the overall cluster health.

Procedure

1 Navigate to the Virtual SAN cluster in the vSphere Web Client.

2 Click **Configure** and then click **Alarm Definitions**.

3 In the search box, type **Virtual SAN** as the search term to display the alarms that are specific to Virtual SAN.

Type Virtual SAN Health Service Alarm to search for Virtual SAN health service alarms.

The default Virtual SAN alarms are displayed.

4 From the list of alarms, click on each alarm to view the alarm definition.

Using the VMkernel Observations for Creating Alarms

VMkernel Observations (VOBs) are system events that you can use to set up Virtual SAN alarms for monitoring and troubleshooting performance and networking issues in the Virtual SAN cluster. In Virtual SAN, these events are known as observations.

VMware ESXi Observation IDs for Virtual SAN

Each VOB event is associated with an identifier (ID). Before you create a Virtual SAN alarm in the vCenter Server, you must identify an appropriate VOB ID for the Virtual SAN event for which you want to create an alert. You can create alerts in the VMware ESXi Observation Log file (`vobd.log`). For example, you should use the following VOB IDs to create alerts for any device failures in the cluster.

- `esx.problem.vob.vsan.lsom.diskerror`

- `esx.problem.vob.vsan.pdl.offline`

To review the list of VOB IDs for Virtual SAN, open the `vobd.log` file located on your ESXi host in the `/var/log` directory. The log file contains the following VOB IDs that you can use for creating Virtual SAN alarms.

Table 14-3. VOB IDs for Virtual SAN

VOB ID	Description
esx.audit.vsan.clustering.enabled	The Virtual SAN clustering service is enabled.
esx.clear.vob.vsan.pdl.online	The Virtual SAN device has come online.
esx.clear.vsan.clustering.enabled	The Virtual SAN clustering services is enabled.
esx.clear.vsan.vsan.network.available	Virtual SAN has one active network configuration.
esx.clear.vsan.vsan.vmknic.ready	A previously reported vmknic has acquired a valid IP.
esx.problem.vob.vsan.lsom.componentthreshold	Virtual SAN reaches the near node component count limit.
esx.problem.vob.vsan.lsom.diskerror	A Virtual SAN device is in a permanent error state.
esx.problem.vob.vsan.lsom.diskgrouplimit	Virtual SAN fails to create a new disk group.
esx.problem.vob.vsan.lsom.disklimit	Virtual SAN fails to add devices to a disk group.
esx.problem.vob.vsan.lsom.diskunhealthy	Virtual SAN disk is unhealthy.
esx.problem.vob.vsan.pdl.offline	A Virtual SAN device is offline.
esx.problem.vsan.clustering.disabled	Virtual SAN clustering services are disabled.
esx.problem.vsan.lsom.congestionthreshold	Virtual SAN device memory or SSD congestion has been updated.
esx.problem.vsan.net.not.ready	A vmknic is added to Virtual SAN network configuration without a valid IP address. This happens when the Virtual SAN network is not ready.
esx.problem.vsan.net.redundancy.lost	The Virtual SAN network configuration does not have the required redundancy.
esx.problem.vsan.no.network.connectivity	Virtual SAN does not have existing networking configuration, which is in use.
esx.problem.vsan.vmknic.not.ready	A vmknic is added to the Virtual SAN network configuration without a valid IP address.

Creating a vCenter Server Alarm for a Virtual SAN Event

You can create alarms to monitor events on the selected Virtual SAN object, including the cluster, hosts, datastores, networks, and virtual machines.

Prerequisites

You must have the required privilege level of `Alarms.Create Alarm` or `Alarm.Modify Alarm`

Procedure

1 Select the vCenter Server object in the inventory that you want to monitor.

2 Click the **Configure** tab > **Alarm Definitions** > click the ➕ icon.

3 Type a name and description for the new alarm.

4 From the **Monitor** drop-down menu, select the object on which you want to configure an alarm.

5 Click the **specific event occurring on this object for example VM Power On** and click **Next**.

6 Click **Triggers** to add a Virtual SAN event that will trigger the alarm. The options on the Triggers page change depending on the type of activity you plan to monitor.

7 Click the **Add** icon (➕).

8 Click in the **Event** column, and select an option from the drop-down menu.

9 Click in the **Status** column, and select an option from the drop-down menu.

10 (Optional) Configure additional conditions to be met before the alarm triggers.

    a Click the **Add** icon to add an argument.

    b Click in the **Argument** column, and select an option from the drop-down menu.

    c Click in the **Operator** column, and select an option from the drop-down menu.

    d Click in the **Value** column, and enter a value in the text field.

       You can add more than one argument.

11 Click **Next**.

You selected and configured alarm triggers.

Handling Failures and Troubleshooting Virtual SAN

<div style="text-align: right; font-size: large;">15</div>

If you encounter problems when using Virtual SAN, you can use troubleshooting topics. The topics help you understand the problem and offer you a workaround, when it is available.

This chapter includes the following topics:

- "Using esxcli Commands with Virtual SAN," on page 147
- "Virtual SAN Configuration on an ESXi Host Might Fail," on page 150
- "Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly," on page 150
- "Virtual SAN Cluster Configuration Issues," on page 151
- "Handling Failures in Virtual SAN," on page 152
- "Shutting Down the Virtual SAN Cluster," on page 164

Using esxcli Commands with Virtual SAN

Use esxcli commands to obtain information about Virtual SAN and to troubleshoot your Virtual SAN environment.

The following commands are available:

Command	Description
`esxcli vsan network list`	Verify which VMkernel adapters are used for Virtual SAN communication.
`esxcli vsan storage list`	List storage disks that were claimed by Virtual SAN.
`esxcli vsan cluster get`	Get Virtual SAN cluster information.
`esxcli vsan health`	Get Virtual SAN cluster health status.
`esxcli vsan debug`	Get Virtual SAN cluster debug information.

The `esxcli vsan debug` commands can help you debug and troubleshoot the Virtual SAN cluster, especially when vCenter Server is not available.

Usage: `esxcli vsan debug {cmd} [cmd options]`

Debug commands:

Command	Description
`esxcli vsan debug disk`	Debug Virtual SAN physical disks.
`esxcli vsan debug object`	Debug Virtual SAN objects.
`esxcli vsan debug resync`	Debug Virtual SAN resyncing objects.

Command	Description
`esxcli vsan debug controller`	Debug Virtual SAN disk controllers.
`esxcli vsan debug limit`	Debug Virtual SAN limits.
`esxcli vsan debug vmdk`	Debug Virtual SAN VMDKs.

Example `esxcli vsan debug` commands:

```
esxcli vsan debug disk summary get
    Overall Health: green
    Component Metadata Health: green
    Memory Pools (heaps): green
    Memory Pools (slabs): green

esxcli vsan debug disk list
UUID: 52e1d1fa-af0e-0c6c-f219-e5e1d224b469
    Name: mpx.vmhba1:C0:T1:L0
    SSD: False
    Overall Health: green
    Congestion Health:
          State: green
          Congestion Value: 0
          Congestion Area: none
    In Cmmds: true
    In Vsi: true
    Metadata Health: green
    Operational Health: green
    Space Health:
          State: green
          Capacity: 107365793792 bytes
          Used: 1434451968 bytes
          Reserved: 150994944 bytes

esxcli vsan debug object health summary get
    Health Status                                       Number Of Objects
    -------------------------------------------------   -----------------
    reduced-availability-with-no-rebuild-delay-timer                    0
    reduced-availability-with-active-rebuild                            0
    inaccessible                                                        0
    data-move                                                          0
    healthy                                                            1
    nonavailability-related-incompliance                               0
    nonavailability-related-reconfig                                   0
    reduced-availability-with-no-rebuild                               0

esxcli vsan debug object list
    Object UUID: 47cbdc58-e01c-9e33-dada-020010d5dfa3
       Version: 5
       Health: healthy
       Owner:
       Policy:
          stripeWidth: 1
          CSN: 1
          spbmProfileName: Virtual SAN Default Storage Policy
          spbmProfileId: aa6d5a82-1c88-45da-85d3-3d74b91a5bad
          forceProvisioning: 0
          cacheReservation: 0
```

```
            proportionalCapacity: [0, 100]
            spbmProfileGenerationNumber: 0
            hostFailuresToTolerate: 1

      Configuration:
        RAID_1
          Component: 47cbdc58-6928-333f-0c51-020010d5dfa3
            Component State: ACTIVE,  Address Space(B): 273804165120 (255.00GB),
            Disk UUID: 52e95956-42cf-4d30-9cbe-763c616614d5, Disk Name: mpx.vmhba1..
            Votes: 1,  Capacity Used(B): 373293056 (0.35GB),
            Physical Capacity Used(B): 369098752 (0.34GB),  Host Name: sc-rdops...
          Component: 47cbdc58-eebf-363f-cf2b-020010d5dfa3
            Component State: ACTIVE,  Address Space(B): 273804165120 (255.00GB),
            Disk UUID: 52d11301-1720-9901-eb0a-157d68b3e4fc,  Disk Name: mpx.vmh...
            Votes: 1,  Capacity Used(B): 373293056 (0.35GB),
            Physical Capacity Used(B): 369098752 (0.34GB),  Host Name: sc-rdops-vm..
        Witness: 47cbdc58-21d2-383f-e45a-020010d5dfa3
          Component State: ACTIVE,  Address Space(B): 0 (0.00GB),
          Disk UUID: 52bfd405-160b-96ba-cf42-09da8c2d7023,  Disk Name: mpx.vmh...
          Votes: 1,  Capacity Used(B): 12582912 (0.01GB),
          Physical Capacity Used(B): 4194304 (0.00GB),  Host Name: sc-rdops-vm...

      Type: vmnamespace
      Path: /vmfs/volumes/vsan:52134fafd48ad6d6-bf03cb6af0f21b8d/New Virtual Machine
      Group UUID: 00000000-0000-0000-0000-000000000000
      Directory Name: New Virtual Machine
esxcli vsan debug controller list
    Device Name: vmhba1
      Device Display Name: LSI Logic/Symbios Logic 53c1030 PCI-X Fusion-MPT Dual Ult..
      Used By VSAN: true
      PCI ID: 1000/0030/15ad/1976
      Driver Name: mptspi
      Driver Version: 4.23.01.00-10vmw
      Max Supported Queue Depth: 127

esxcli vsan debug limit get
      Component Limit Health: green
      Max Components: 750
      Free Components: 748
      Disk Free Space Health: green
      Lowest Free Disk Space: 99 %
      Used Disk Space: 1807745024 bytes
      Used Disk Space (GB): 1.68 GB
      Total Disk Space: 107365793792 bytes
      Total Disk Space (GB): 99.99 GB
      Read Cache Free Reservation Health: green
```

```
        Reserved Read Cache Size: 0 bytes
        Reserved Read Cache Size (GB): 0.00 GB
        Total Read Cache Size: 0 bytes
        Total Read Cache Size (GB): 0.00 GB

esxcli vsan debug vmdk list
    Object: 50cbdc58-506f-c4c2-0bde-020010d5dfa3
        Health: healthy
        Type: vdisk
        Path: /vmfs/volumes/vsan:52134fafd48ad6d6-bf03cb6af0f21b8d/47cbdc58-e01c-9e33-
            dada-020010d5dfa3/New Virtual Machine.vmdk
        Directory Name: N/A

esxcli vsan debug resync list
    Object           Component              Bytes Left To Resync  GB Left To Resync
    ---------------- ---------------------- --------------------- -----------------
    31cfdc58-e68d... Component:23d1dc58...             536870912  0.50
    31cfdc58-e68d... Component:23d1dc58...            1073741824  1.00
    31cfdc58-e68d... Component:23d1dc58...            1073741824  1.00
```

Virtual SAN Configuration on an ESXi Host Might Fail

In certain circumstances, the task of configuring Virtual SAN on a particular host might fail.

Problem

An ESXi host that joins a Virtual SAN cluster fails to have Virtual SAN configured.

Cause

If a host does not meet hardware requirements or experiences other problems, Virtual SAN might fail to configure the host. For example, insufficient memory on the host might prevent Virtual SAN from being configured.

Solution

1   Place the host that causes the failure in Maintenance Mode.

2   Move the host out of the Virtual SAN cluster.

3   Resolve the problem that prevents the host to have Virtual SAN configured.

4   Exit Maintenance Mode.

5   Move the host back into the Virtual SAN cluster.

Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly

When you use the **Check Compliance** button, a virtual machine object does not change its status from Not Compliant to Compliant even though Virtual SAN resources have become available and satisfy the virtual machine profile.

Problem

When you use a force provisioning option, you can provision a virtual machine object even when the policy specified in the virtual machine profile is not satisfiable with the resources currently available in the Virtual SAN cluster. The object is created, but remains in the non-compliant status.

Virtual SAN is expected to bring the object into compliance when storage resources in the cluster become available, for example, when you add a host. However, the object's status does not change to compliant immediately after you add resources.

Cause

This occurs because Virtual SAN regulates the pace of the reconfiguration to avoid overloading the system. The amount of time it takes for compliance to be achieved depends on the number of objects in the cluster, the I/O load on the cluster and the size of the object in question. In most cases, compliance will be achieved within a reasonable time.

Virtual SAN Cluster Configuration Issues

After you make any changes to Virtual SAN configuration, vCenter Server performs validation checks for Virtual SAN configuration. Validation checks are also performed as a part of a host synchronization process. If vCenter Server detects any configuration problems, it displays error messages.

Problem

A number of error messages indicate that vCenter Server has detected a problem with Virtual SAN configuration.

Solution

Use the following methods to fix Virtual SAN configuration problems.

Table 15-1. Virtual SAN Configuration Errors and Solutions

Virtual SAN Configuration Error	Solution
Host with the VSAN service enabled is not in the vCenter cluster	Add the host to the Virtual SAN cluster. 1 Right-click the host, and select **Move To**. 2 Select the Virtual SAN cluster and click **OK**.
Host is in a VSAN enabled cluster but does not have VSAN service enabled	Verify whether Virtual SAN network is properly configured and enabled on the host. See "Configuring Virtual SAN Network," on page 43.
VSAN network is not configured	Configure Virtual SAN network. See "Configuring Virtual SAN Network," on page 43.
Host cannot communicate with all other nodes in the VSAN enabled cluster	Might be caused by network isolation. See "Networking Requirements for Virtual SAN," on page 21 documentation.
Found another host participating in the VSAN service which is not a member of this host's vCenter cluster.	Make sure that the Virtual SAN cluster configuration is correct and all Virtual SAN hosts are in the same subnet. See "Designing the Virtual SAN Network," on page 32.

Handling Failures in Virtual SAN

Virtual SAN handles failures of the storage devices, hosts and network in the cluster according to the severity of the failure. You can diagnose problems in Virtual SAN by observing the performance of the Virtual SAN datastore and network.

Failure Handling in Virtual SAN

Virtual SAN implements mechanisms for indicating failures and rebuilding unavailable data for data protection.

Failure States of Virtual SAN Components

In Virtual SAN, components that have failed can be in absent or degraded state. According to the component state, Virtual SAN uses different approaches for recovering virtual machine data.

Virtual SAN also provides alerts about the type of component failure. See "Using the VMkernel Observations for Creating Alarms," on page 145 and "Using the Virtual SAN Default Alarms," on page 143.

Virtual SAN supports two types of failure states for components:

Table 15-2. Failure States of Components in Virtual SAN

Component Failure State	Description	Recovery	Cause
Degraded	A component is in degraded state if Virtual SAN detects a permanent component failure and assumes that the component is not going to recover to working state.	Virtual SAN starts rebuilding the affected components immediately.	■ Failure of a flash caching device ■ Magnetic or flash capacity device failure ■ Storage controller failure
Absent	A component is in absent state if Virtual SAN detects a temporary component failure where the component might recover and restore its working state.	Virtual SAN starts rebuilding absent components if they are not available within certain timeout. By default, Virtual SAN starts rebuilding absent components after 60 minutes.	■ Lost network connectivity ■ Failure of a physical network adapter ■ ESXi host failure ■ Unplugged flash caching device ■ Unplugged magnetic disk or flash capacity device

Examine the Failure State of a Component

Use the vSphere Web Client to examine whether a component is in the absent or degraded failure state.

If a failure occurs in the cluster, Virtual SAN marks the components for an object as absent or degraded based on the failure severity.

Procedure

1 In the vSphere Web Client, navigate to the Virtual SAN cluster.

2 On the **Monitor** tab, click **vSAN** and select **Virtual Disks**.

The home directories and virtual disks of the virtual machines in the cluster appear.

3 Select a virtual machine object.

4   On the **Physical Disk Placement** tab, examine the Component State property of the components for the selected object.

If a failure has occurred in the Virtual SAN cluster, the Component State property is equal to Absent or Degraded.

Object States That Indicate Problems in Virtual SAN

Examine the compliance status and the operational state of a virtual machine object to determine how a failure in the cluster affects the virtual machine.

Table 15-3. Object State

Object State Type	Description
Compliance Status	The compliance status of a virtual machine object indicates whether it meets the requirements of the assigned VM storage policy.
Operational State	The operational state of an object can be healthy or unhealthy. It indicates the type and number of failures in the cluster. An object is healthy if an intact replica is available and more than 50 percent of the object's votes are still available. An object is unhealthy if an entire replica is not available or less than 50 percent of the object's votes are unavailable. For example, an object might become unhealthy if a network failure occurs in the cluster and a host becomes isolated.

To determine the overall influence of a failure on a virtual machine, examine the compliance status and the operational state. If the operational state remains healthy although the object is noncompliant, the virtual machine can continue using the Virtual SAN datastore. If the operational state is unhealthy, the virtual machine cannot use the datastore.

Examine the Health of an Object in Virtual SAN

Use the vSphere Web Client to examine whether a virtual machine is healthy. A virtual machine is considered as healthy when a replica of the VM object and more than 50 percent of the votes for an object are available.

Procedure

1   In the vSphere Web Client, navigate to the Virtual SAN cluster.

2   On the **Monitor** tab, click **vSAN** and select **Virtual Disks**.

The home directories and virtual disks of the virtual machines in the cluster appear.

3   For a virtual machine object, examine the value of the Operational State property.

If the Operational State is Unhealthy, the vSphere Web Client indicates the reason for the unhealthy state in brackets.

Examine the Compliance of a Virtual Machine in Virtual SAN

Use the vSphere Web Client to examine whether a virtual machine object is compliant with the assigned VM storage policy.

Procedure

1   Examine the compliance status of a virtual machine.

a   Browse to the virtual machine in the vSphere Web Client navigator.

b   On the **Summary** tab, examine the value of the VM Storage Policy Compliance property under VM Storage Policies.

2   Examine the compliance status of the objects of the virtual machine.

   a   In the vSphere Web Client, navigate to the Virtual SAN cluster.

   b   On the **Monitor** tab, click **vSAN** and select **Virtual Disks**.

   c   Select a virtual machine object.

   d   Examine the value of the Compliance Status property for the object. If the Compliance Status is different from Compliant, determine the cause for the noncompliance.

   ■   Examine the Operational State of the object to verify whether the object is healthy.

   ■   On the **Compliance Failure** tab, examine which requirements from the VM storage policy that the object cannot satisfy.

   ■   On the **Physical Disk Placement** tab, examine the state of the object components.

Accessibility of Virtual Machines Upon a Failure in Virtual SAN

If a virtual machine uses Virtual SAN storage, its storage accessibility might change according to the type of failure in the Virtual SAN cluster.

Changes in the accessibility occur when the cluster experiences more failures than the policy for a virtual machine object tolerates.

As a result from a failure in the Virtual SAN cluster, a virtual machine object might become inaccessible. An object is inaccessible if a full replica of the object is not available because the failure affects all replicas, or when less than 50 percent of the object's votes are available.

According to the type of object that is inaccessible, virtual machines behave in the following ways:

Table 15-4. Inaccessibility of Virtual Machine Objects

Object Type	Virtual Machine State	Virtual Machine Symptoms
VM Home Namespace	■ Inaccessible ■ Orphaned if vCenter Server or the ESXi host cannot access the `.vmx` file of the virtual machine.	The virtual machine process might crash and the virtual machine might be powered off.
VMDK	Inaccessible	The virtual machine remains powered on but the I/O operations on the VMDK are not being executed. After a certain timeout passes, the guest operating system ends the operations.

Virtual machine inaccessibility is not a permanent state. After the underlying issue is resolved, and a full replica and more than 50 percent of the object's votes are restored, the virtual machine automatically becomes accessible again.

Storage Device is Failing in a Virtual SAN Cluster

Virtual SAN monitors the performance of each storage device and proactively isolates unhealthy devices. It detects gradual failure of a storage device and isolates the device before congestion builds up within the affected host and the entire Virtual SAN cluster.

If a disk experiences sustained high latencies or congestion, Virtual SAN considers the device as a dying disk, and evacuates data from the disk. Virtual SAN handles the dying disk by evacuating or rebuilding data. No user action is required, unless the cluster lacks resources or has inaccessible objects.

Component Failure State and Accessibility

The Virtual SAN components that reside on the magnetic disk or flash capacity device are marked as absent.

Behavior of Virtual SAN

Virtual SAN responds to the storage device failure in the following ways.

Parameter	Behavior
Alarms	An alarm is generated from each host whenever an unhealthy device is diagnosed. A warning is issued whenever a disk is suspected of being unhealthy.
Health check	The **Overall disk health** check issues a warning for the dying disk.
Health status	On the Disk Management page, the health status of the dying disk is listed as **Unhealthy**. When Virtual SAN completes evacuation of data, the health status is listed as **DyingDiskEmpty**.
Rebuilding data	Virtual SAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, Virtual SAN starts the recovery process immediately because the components are marked as degraded.
	If resources are available, Virtual SAN automatically reprotects the data.

If Virtual SAN detects a disk with a permanent error, it makes a limited number of attempts to revive the disk by unmounting and mounting it.

Capacity Device Not Accessible in a Virtual SAN Cluster

When a magnetic disk or flash capacity device fails, Virtual SAN evaluates the accessibility of the objects on the device and rebuilds them on another host if space is available and the **Primary level of failures to tolerate** is set to 1 or more.

Component Failure State and Accessibility

The Virtual SAN components that reside on the magnetic disk or flash capacity device are marked as degraded.

Behavior of Virtual SAN

Virtual SAN responds to the capacity device failure in the following ways.

Parameter	Behavior
Primary level of failures to tolerate	If the **Primary level of failures to tolerate** in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, Virtual SAN starts an automatic reprotection.
	If the **Primary level of failures to tolerate** is set to 0, a virtual machine object is inaccessible if one of the object's components resides on the failed capacity device.
	Restore the virtual machine from a backup.
I/O operations on the capacity device	Virtual SAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.
	If Virtual SAN determines that the object is available, all running I/O operations are resumed.
Rebuilding data	Virtual SAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, Virtual SAN starts the recovery process immediately because the components are marked as degraded.
	If resources are available, an automatic reprotect will occur.

A Flash Caching Device Is Not Accessible in a Virtual SAN Cluster

When a flash caching device fails, Virtual SAN evaluates the accessibility of the objects on the disk group that contains the cache device, and rebuilds them on another host if possible and the **Primary level of failures to tolerate** is set to 1 or more.

Component Failure State and Accessibility

Both cache device and capacity devices that reside in the disk group, for example, magnetic disks, are marked as degraded. Virtual SAN interprets the failure of a single flash caching device as a failure of the entire disk group.

Behavior of Virtual SAN

Virtual SAN responds to the failure of a flash caching device in the following way:

Parameter	Behavior
Primary level of failures to tolerate	If the **Primary level of failures to tolerate** in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, Virtual SAN starts an automatic reprotection.
	If the **Primary level of failures to tolerate** is set to 0, a virtual machine object is inaccessible if one of the object's components is on the failed disk group.
I/O operations on the disk group	Virtual SAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.
	If Virtual SAN determines that the object is available, all running I/O operations are resumed.
Rebuilding data	Virtual SAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, Virtual SAN starts the recovery process immediately because the components are marked as degraded.

A Host Is Not Responding in a Virtual SAN Cluster

If a host stops responding because of a failure or reboot of the host, Virtual SAN waits for the host to recover before Virtual SAN rebuilds the components on the host elsewhere in the cluster.

Component Failure State and Accessibility

The Virtual SAN components that reside on the host are marked as absent.

Behavior of Virtual SAN

Virtual SAN responds to the host failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	If the **Primary level of failures to tolerate** in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, Virtual SAN starts an automatic reprotection.
	If the **Primary level of failures to tolerate** is set to 0, a virtual machine object is inaccessible if the object's components reside on the failed host.
I/O operations on the host	Virtual SAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.
	If Virtual SAN determines that the object is available, all running I/O operations are resumed.
Rebuilding data	If the host does not rejoin the cluster within 60 minutes, Virtual SAN examines whether some of the other hosts in the cluster can satisfy the requirements for cache, space and placement rules for the objects on the inaccessible host. If such a host is available, Virtual SAN starts the recovery process.
	If the host rejoins the cluster after 60 minutes and recovery has started, Virtual SAN evaluates whether to continue the recovery or stop it and resynchronize the original components.

Network Connectivity Is Lost in the Virtual SAN Cluster

When the connectivity between the hosts in the cluster is lost, Virtual SAN determines the active partition and rebuilds the components from the isolated partition on the active partition if the connectivity is not restored.

Component Failure State and Accessibility

Virtual SAN determines the partition where more than 50 percent of the votes of an object are available. The components on the isolated hosts are marked as absent.

Behavior of Virtual SAN

Virtual SAN responds to a network failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	If the **Primary level of failures to tolerate** in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, Virtual SAN starts an automatic reprotection.
	If the **Primary level of failures to tolerate** is set to 0, a virtual machine object is inaccessible if the object's components are on the isolated hosts.
I/O operations on the isolated hosts	Virtual SAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.
	If Virtual SAN determines that the object is available, all running I/O operations are resumed.
Rebuilding data	If the host rejoins the cluster within 60 minutes, Virtual SAN synchronizes the components on the host.
	If the host does not rejoin the cluster within 60 minutes, Virtual SAN examines whether some of the other hosts in the cluster can satisfy the requirements for cache, space and placement rules for the objects on the inaccessible host. If such a host is available, Virtual SAN starts the recovery process.
	If the host rejoins the cluster after 60 minutes and recovery has started, Virtual SAN evaluates whether to continue the recovery or stop it and resynchronize the original components.

A Storage Controller Fails in a Virtual SAN Cluster

When a storage controller fails, Virtual SAN evaluates the accessibility of the objects on the disk groups that are attached to the controller and rebuilds them on another host.

Symptoms

If a host contains a single storage controller and multiple disk groups, and all devices in all disk groups are failed, then you might assume that a failure in the common storage controller is the root cause. Examine the VMkernel log messages to determine the nature of the fault.

Component Failure State and Accessibility

When a storage controller fails, the components on the flash caching devices and capacity devices in all disk groups that are connected to the controller are marked as degraded.

If a host contains multiple controllers, and only the devices that are attached to an individual controller are inaccessible, then you might assume that this controller has failed.

Behavior of Virtual SAN

Virtual SAN responds to a storage controller failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	If the **Primary level of failures to tolerate** in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, Virtual SAN starts an automatic reprotection.
	If the **Primary level of failures to tolerate** is set to 0, a virtual machine object is inaccessible if the object's components reside on the disk groups that are connected to the storage controller.
Rebuilding data	Virtual SAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, Virtual SAN starts the recovery process immediately because the components are marked as degraded.

Stretched Cluster Site Fails or Loses Network Connection

A Virtual SAN stretched cluster manages failures that occur due to the loss of a network connection between sites or the temporary loss of one site.

Stretched Cluster Failure Handling

In most cases, the stretched cluster continues to operate during a failure and automatically recovers after the failure is resolved.

Table 15-5. How Stretched Cluster Handles Failures

Type of Failure	Behavior
Network Connection Lost Between Active Sites	If the network connection fails between the two active sites, the witness host and the preferred site continue to service storage operations, and keep data available. When the network connection returns, the two active sites are resynchronized.
Secondary Site Fails or Loses Network Connection	If the secondary site goes offline or becomes isolated from the preferred site and the witness host, the witness host and the preferred site continue to service storage operations, and keep data available. When the secondary site returns to the cluster, the two active sites are resynchronized.
Preferred Site Fails or Loses Network Connection	If the preferred site goes offline or becomes isolated from the secondary site and the witness host, the secondary site continues storage operations as long as it remains connected to the witness host. When the preferred site returns to the cluster, the two active sites are resynchronized.
Witness Host Fails or Loses Network Connection	If the witness host goes offline or becomes isolated from the preferred site or the secondary site, objects become noncompliant but data remains available. VMs that are currently running are not affected.

Troubleshooting Virtual SAN

Examine the performance and accessibility of virtual machines to diagnose problems in the Virtual SAN cluster.

Verify Drivers, Firmware, Storage I/O Controllers Against the *VMware Compatibility Guide*

Use the Virtual SAN Health Service to verify whether your hardware components, drivers, and firmware are compatible with Virtual SAN.

Using hardware components, drivers, and firmware that are not compatible with Virtual SAN might cause problems in the operation of the Virtual SAN cluster and the virtual machines running on it.

The hardware compatibility health checks verify the your hardware against the *VMware Compatibility Guide*. For more information about using the Virtual SAN health service, see "Monitoring Virtual SAN Health," on page 135.

Examining Performance in a Virtual SAN Cluster

Monitor the performance of virtual machines, hosts, and the Virtual SAN datastore to identify potential storage problems.

Monitor regularly the following performance indicators to identify faults in Virtual SAN storage, for example, by using the performance charts in the vSphere Web Client:

■   Datastore. Rate of I/O operations on the aggregated datastore.

■   Virtual Machine. I/O operations, memory and CPU usage, network throughput and bandwidth.

You can use the Virtual SAN performance service to access detailed performance charts. For information about using the performance service, see "Monitoring Virtual SAN Performance," on page 138. For more information about using performance data in a Virtual SAN cluster, see the *Virtual SAN Troubleshooting Reference Manual*.

Network Misconfiguration Status in a Virtual SAN Cluster

After you enable Virtual SAN on a cluster, the datastore is not assembled correctly because of a detected network misconfiguration.

Problem

After you enable Virtual SAN on a cluster, on the **Summary** tab for the cluster the Network Status for Virtual SAN appears as `Misconfiguration detected`.

Cause

One or more members of the cluster cannot communicate because of either of the following reasons:

■   A host in the cluster does not have a VMkernel adapter for Virtual SAN.

■   The hosts cannot connect each other in the network.

Solution

Join the members of the cluster to the same network. See "Configuring Virtual SAN Network," on page 43.

Virtual Machine Appears as Noncompliant, Inaccessible or Orphaned in Virtual SAN

The state of a virtual machine that stores data on a Virtual SAN datastore appears as noncompliant, inaccessible or orphaned because of failures in the Virtual SAN cluster.

Problem

A virtual machine on a Virtual SAN datastore is in one of the following states that indicate a fault in the Virtual SAN cluster.

■   The virtual machine is non-compliant and the compliance status of some of its object is noncompliant. See "Examine the Compliance of a Virtual Machine in Virtual SAN," on page 153.

■   The virtual machine object is inaccessible or orphaned. See "Examine the Failure State of a Component," on page 152.

If an object replica is still available on another host, Virtual SAN forwards the I/O operations of the virtual machine to the replica.

Cause

If the object of the virtual machine can no longer satisfy the requirement of the assigned VM storage policy, Virtual SAN considers it noncompliant. For example, a host might temporarily lose connectivity. See "Object States That Indicate Problems in Virtual SAN," on page 153.

If Virtual SAN cannot locate a full replica or more than 50 percent of the votes for the object, the virtual machine becomes inaccessible. If a Virtual SAN detects that the .vmx file is not accessible because the VM Home Namespace is corrupted, the virtual machine becomes orphaned. See "Accessibility of Virtual Machines Upon a Failure in Virtual SAN," on page 154.

Solution

If the cluster contains enough resources, Virtual SAN automatically recovers the corrupted objects if the failure is permanent.

If the cluster does not have enough resources to rebuild the corrupted objects, extend the space in the cluster. See "Expanding Virtual SAN Cluster Capacity and Performance," on page 108 and "Add a Host to the Virtual SAN Cluster," on page 108.

Attempt to Create a Virtual Machine on Virtual SAN Fails

When you try to deploy a virtual machine in a Virtual SAN cluster, the operation fails with an error that the virtual machine files cannot be created.

Problem

The operation for creating a virtual machine fails with an error status: `Cannot complete file creation operation`.

Cause

The deployment of a virtual machine on Virtual SAN might fail for several reasons.

■ Virtual SAN cannot allocate space for the virtual machine storage policies and virtual machine objects. Such a failure might occur if the datastore does not have enough usable capacity, for example, if a physical disk is temporarily disconnected from the host.

■ The virtual machine has very large virtual disks and the hosts in the cluster cannot provide storage for them based on the placement rules in the VM storage policy

For example, if the **Primary level of failures to tolerate** in the VM storage policy is set to 1, Virtual SAN must store two replicas of a virtual disk in the cluster, each replica on a different host. The datastore might have this space after aggregating the free space on all hosts in the cluster. However, no two hosts can be available in the cluster, each providing enough space to store a separate replica of the virtual disk.

Virtual SAN does not move components between hosts or disks groups to free space for a new replica, even though the cluster might contain enough space for provisioning the new virtual machine.

Solution

◆ Verify the state of the capacity devices in the cluster.

a In the vSphere Web Client, navigate to the Virtual SAN cluster.

b On the **Monitor** tab, click **vSAN** and select **Physical Disks**.

c Examine the capacity and health status of the devices on the hosts in the cluster.

Stretched Cluster Configuration Error When Adding a Host

Before adding new hosts to a stretched cluster, all current hosts must be connected. If a current host is disconnected, the configuration of the new host is incomplete.

Problem

After you add a new host to a stretched cluster in which some hosts are disconnected, on the Summary tab for the cluster the Configuration Status for Virtual SAN appears as `Unicast agent unset on host`.

Cause

When a new host joins a stretched cluster, Virtual SAN must update the configuration on all hosts in the cluster. If one or more hosts are disconnected from the vCenter Server, the update fails. The new host successfully joins the cluster, but its configuration is incomplete.

Solution

Verify all hosts are connected to vCenter Server, and click on the link provided in the Configuration Status message to update the configuration of the new host.

If you cannot rejoin the disconnected host, remove the disconnected host from the cluster, and click on the link provided in the Configuration Status message to update the configuration of the new host.

Stretched Cluster Configuration Error When Using RVC to Add a Host

If you use the RVC tool to add a new host to a stretched cluster, the configuration of the new host is incomplete.

Problem

After you use the RVC tool to add a new host to a stretched cluster, on the Summary tab for the cluster the Configuration Status for Virtual SAN appears as `Unicast agent unset on host`.

Cause

When a new host joins a stretched cluster, Virtual SAN must update the configuration on all hosts in the cluster. If you use the RVC tool to add the host, the update does not occur. The new host successfully joins the cluster, but its configuration is incomplete.

Solution

Verify all hosts are connected to vCenter Server, and click on the link provided in the Configuration Status message to update the configuration of the new host.

Cannot Add or Remove the Witness Host in a Stretched Cluster

Before adding or removing the witness host in a stretched cluster, all current hosts must be connected. If a current host is disconnected, you cannot add or remove the witness host.

Problem

When you add or remove a witness host in a stretched cluster in which some hosts are disconnected, the operation fails with an error status: `The operation is not allowed in the current state. Not all hosts in the cluster are connected to Virtual Center.`

Cause

When the witness host joins or leaves a stretched cluster, Virtual SAN must update the configuration on all hosts in the cluster. If one or more hosts are disconnected from the vCenter Server, the witness host cannot be added or removed.

Solution

Verify all hosts are connected to vCenter Server, and retry the operation. If you cannot rejoin the disconnected host, remove the disconnected host from the cluster, and then you can add or remove the witness host.

Disk Group Becomes Locked

In an encrypted Virtual SAN cluster, when communication between a host and the KMS is lost, the disk group can become locked if the host reboots.

Problem

Virtual SAN locks a host's disk groups when the host reboots and it cannot get the KEK from the KMS. The disks behave as if they are unmounted. Objects on the disks become inaccessible.

You can view a disk group's health status on the Disk Management page in the vSphere Web Client. An Encryption health check warning notifies you that a disk is locked.

Cause

Hosts in an encrypted Virtual SAN cluster do not store the KEK on disk. If a host reboots and cannot get the KEK from the KMS, Virtual SAN locks the host's disk groups.

Solution

To exit the locked state, you must restore communication with the KMS and reestablish the trust relationship.

Replacing Existing Hardware Components

Under certain conditions, you must replace hardware components, drivers, firmware, and storage I/O controllers in the Virtual SAN cluster.

In Virtual SAN, you should replace hardware devices when you encounter failures or if you must upgrade your cluster.

Replace a Flash Caching Device on a Host

You should replace a flash caching device if you detect a failure or when you must upgrade it. Before you physically unplug a flash device from the host, you must manually remove the device from Virtual SAN.

> **Caution**   If you decommission the flash caching device without removing it from Virtual SAN first, Virtual SAN uses smaller amount of cache than expected. As a result, the cluster performance becomes degraded.

When you replace a flash caching device, the virtual machines on the disk group become inaccessible and the components on the group are marked as degraded. See "A Flash Caching Device Is Not Accessible in a Virtual SAN Cluster," on page 156.

Prerequisites

■ Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

  If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.

■ If you upgrade the flash caching device, verify the following requirements:

  ■ If you upgrade the flash caching device, verify that the cluster contains enough space to migrate the data from the disk group that is associated with the flash device.

■ Place the host in maintenance mode. See "Place a Member of Virtual SAN Cluster in Maintenance Mode," on page 112.

Procedure

1 In the vSphere Web Client, navigate to the Virtual SAN cluster.

2 On the **Configure** tab, click **Disk Management** under vSAN.

3 Select the disk group that contains the device that you want to replace.

4 Select the flash caching device and click **Remove selected disk(s) from disk group**.

After the flash caching device is deleted from the Virtual SAN cluster, the cluster details reflect the current cluster capacity and configuration settings. Virtual SAN discards the disk group memberships, deletes partitions, and removes stale data from all devices.

What to do next

1 Add a new device to the host.

The host automatically detects the device.

2 If the host is unable to detect the device, perform a device rescan.

Replace a Capacity Device

You should replace a flash capacity device or a magnetic disk if you detect a failure or when you upgrade it. Before you physically remove the device from the host, you must manually delete the device from Virtual SAN.

When you unplug a capacity device without removing it from the Virtual SAN cluster, the virtual machines on the disk group become inaccessible and the components on the group are marked as absent.

If the capacity device fails, the virtual machines become inaccessible and the components on the group are marked as degraded. See "Capacity Device Not Accessible in a Virtual SAN Cluster," on page 155.

Prerequisites

■ Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.

■ If you upgrade the capacity device, verify the following requirements:

■ Verify that the cluster contains enough space to migrate the data from the capacity device.

■ Place the host in maintenance mode. See "Place a Member of Virtual SAN Cluster in Maintenance Mode," on page 112.

Procedure

1 In the vSphere Web Client, navigate to the Virtual SAN cluster.

2 On the **Configure** tab, click **Disk Management** under vSAN.

3 Select the disk group that contains the device that you want to replace.

4 Select the flash capacity device or magnetic disk, and click **Remove selected disk(s) from disk group**.

What to do next

1 Add a new device to the host.

The host automatically detects the device.

2   If the host is unable to detect the device, perform a device rescan.

Remove a Device from a Host by Using an ESXCLI Command

If you detect a failed storage device or if you upgrade a device, you can manually remove it from a host by using an ESXCLI command.

If you remove a flash caching device, Virtual SAN deletes the disk group that is associated with the flash device and all its member devices.

Prerequisites

Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.

Procedure

1   Open an SSH connection to the ESXi host.

2   To identify the device ID of the failed device, run this command and learn the device ID from the output.

```
esxcli vsan storage list
```

3   To remove the device from Virtual SAN, run this command.

```
esxcli vsan storage remove –d device_id
```

What to do next

1   Add a new device to the host.

    The host automatically detects the device.

2   If the host is unable to detect the device, perform a device rescan.

Shutting Down the Virtual SAN Cluster

When necessary, you can shut down the entire Virtual SAN cluster.

If you plan to shut down the Virtual SAN cluster, you do not need to manually disable Virtual SAN on the cluster.

Procedure

1   Power off all virtual machines (VMs) running in the Virtual SAN cluster.

2   Place the ESXi hosts in maintenance mode.

    a   Right-click the host and select **Enter Maintenance Mode**.

    b   Select the **No data migration** evacuation mode and click **OK**.

3   In the Confirm Maintenance Mode wizard, deselect the **Move powered-off and suspended virtual machines to other hosts in the cluster** check box.

    When you deselect this check box, Virtual SAN does not migrate the VMs to other hosts. If you plan to shut down the entire cluster and put all hosts in maintenance mode, you do not need to move or migrate the VM storage objects to other hosts or devices in the cluster.

4   Power off the hosts after they have successfully entered maintenance mode.

5   Power on the ESXi hosts.

    a   On the physical box where ESXi is installed, press the power button until the power-on sequence begins.

       The ESXi host starts, locates its VMs, and functions normally.

After you power on the hosts, the Virtual SAN cluster is automatically recreated.

If you navigate to the ESXi host and click **Summary**, you might see that the Network Status of the cluster appears as `Misconfiguration detected`.

You can ignore the status message if you did not make network configuration changes and the Virtual SAN cluster was working as expected before you shut down the cluster. The message disappears after at least three hosts join the cluster.

6   Take the hosts out of maintenance mode.

7   Restart the VMs.

Index