

# Administering VMware vSAN

Update 1

16 OCT 2018

VMware vSphere 6.7

VMware vSAN 6.7



vmware®

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

<https://docs.vmware.com/>

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

[docfeedback@vmware.com](mailto:docfeedback@vmware.com)

**VMware, Inc.**  
3401 Hillview Ave.  
Palo Alto, CA 94304  
[www.vmware.com](http://www.vmware.com)

Copyright © 2015–2018 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

# Contents

	About Administering VMware vSAN	5
<b>1</b>	<b>Introduction to vSAN</b>	<b>6</b>
<b>2</b>	<b>Configuring and Managing a vSAN Cluster</b>	<b>7</b>
	Configure a Cluster for vSAN Using the vSphere Client	7
	Configure a Cluster for vSAN Using the vSphere Web Client	9
	Enable vSAN on an Existing Cluster	11
	Disable vSAN	11
	Edit vSAN Settings	12
	View vSAN Datastore	13
<b>3</b>	<b>Using vSAN Policies</b>	<b>15</b>
	About vSAN Policies	15
	Host Pinning	18
	View vSAN Storage Providers	19
	About the vSAN Default Storage Policy	20
	Assign a Default Storage Policy to vSAN Datastores	21
	Define a Storage Policy for vSAN Using vSphere Client	22
	Define a Storage Policy for vSAN Using the vSphere Web Client	24
<b>4</b>	<b>Expanding and Managing a vSAN Cluster</b>	<b>26</b>
	Expanding a vSAN Cluster	26
	Working with Maintenance Mode	31
	Managing Fault Domains in vSAN Clusters	33
	Using the vSAN iSCSI Target Service	37
	Migrate a Hybrid vSAN Cluster to an All-Flash Cluster	43
	Power off a vSAN Cluster	43
<b>5</b>	<b>Device Management in a vSAN Cluster</b>	<b>44</b>
	Managing Disk Groups and Devices	44
	Working with Individual Devices	47
<b>6</b>	<b>Increasing Space Efficiency in a vSAN Cluster</b>	<b>55</b>
	Introduction to vSAN Space Efficiency	55
	Reclaiming Space with SCSI Unmap	55
	Using Deduplication and Compression	56
	Using RAID 5 or RAID 6 Erasure Coding	62

RAID 5 or RAID 6 Design Considerations 63

## **7 Using Encryption on a vSAN Cluster 64**

How vSAN Encryption Works 64

Design Considerations for vSAN Encryption 65

Set Up the KMS Cluster 65

Enable Encryption on a New vSAN Cluster 71

Generate New Encryption Keys 72

Enable vSAN Encryption on Existing vSAN Cluster 72

vSAN Encryption and Core Dumps 73

## **8 Upgrading the vSAN Cluster 77**

Before You Upgrade vSAN 77

Upgrade the vCenter Server 79

Upgrade the ESXi Hosts 79

About the vSAN Disk Format 81

Verify the vSAN Cluster Upgrade 88

Using the RVC Upgrade Command Options 88

vSAN Build Recommendations for vSphere Update Manager 89

# About Administering VMware vSAN

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

## Intended Audience

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

For more information about vSAN and how to create a vSAN cluster, see the *vSAN Planning and Deployment Guide*.

For more information about monitoring a vSAN cluster and fixing problems, see the *vSAN Monitoring and Troubleshooting Guide*.

## vSphere Client and vSphere Web Client

Instructions in this guide reflect the vSphere Client (an HTML5-based GUI). You can also use the instructions to perform the tasks by using the vSphere Web Client (a Flex-based GUI).

Tasks for which the workflow differs significantly between the vSphere Client and the vSphere Web Client have duplicate procedures that provide steps according to the respective client interface. The procedures that relate to the vSphere Web Client, contain vSphere Web Client in the title.

---

**Note** In vSphere 6.7 Update 1, almost all of the vSphere Web Client functionality is implemented in the vSphere Client. For an up-to-date list of any remaining unsupported functionality, see [Functionality Updates for the vSphere Client](#).

---

# Introduction to vSAN

1

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

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

# Configuring and Managing a vSAN Cluster

# 2

You can configure and manage a vSAN cluster by using the vSphere Client, esxcli commands, and other tools.

This chapter includes the following topics:

- [Configure a Cluster for vSAN Using the vSphere Client](#)
- [Configure a Cluster for vSAN Using the vSphere Web Client](#)
- [Enable vSAN on an Existing Cluster](#)
- [Disable vSAN](#)
- [Edit vSAN Settings](#)
- [View vSAN Datastore](#)

## Configure a Cluster for vSAN Using the vSphere Client

You can use the Configure vSAN wizard in the HTML5-based vSphere Client to complete the basic configuration of your vSAN cluster.

**Configure vSAN**

1 Configuration type

**2 Services**

3 Claim disks

4 Create fault domains

5 Ready to complete

**Services**

Select the services to enable.

These settings require all disks to be reformatted. Moving large amount of stored data might be slow and temporarily decrease the performance of the cluster.

**Deduplication and Compression Services**  ⓘ

**Encryption**  ⓘ

Erase disks before use ⓘ

KMS cluster: \_\_\_\_\_ v

Options:

Allow Reduced Redundancy ⓘ

CANCEL BACK NEXT

## Prerequisites

Create a cluster and add hosts to the cluster before using the Configure vSAN wizard to complete the basic configuration.

## Procedure

- 1 Navigate to an existing cluster in the vSphere Client.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services** and click the **Configure** button.
- 4 Select the configuration type, and click **Next**.
  - Single site cluster. All hosts at one site, with shared witness functions.
  - Two host vSAN cluster. One host at each site, and a witness host at another site.
  - Stretched cluster. Two active data sites, each with an even number of hosts and storage devices, and a witness host at a third site.
- 5 On the **Services** page, configure vSAN services, and click **Next**.
  - a (Optional) Enable **Deduplication and Compression** on the cluster.
  - b (Optional) Enable **Encryption**, and select a KMS.
  - c You can select the **Allow Reduced Redundancy** check box to enable encryption or deduplication and compression on a vSAN 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.
- 6 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.
- 7 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 vSAN Clusters](#).

For more information about stretched clusters, see "Introduction to Stretched Clusters" in *vSAN Planning and Deployment*.
- 8 On the **Ready to complete** page, review the configuration, and click **Finish**.

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

## What to do next

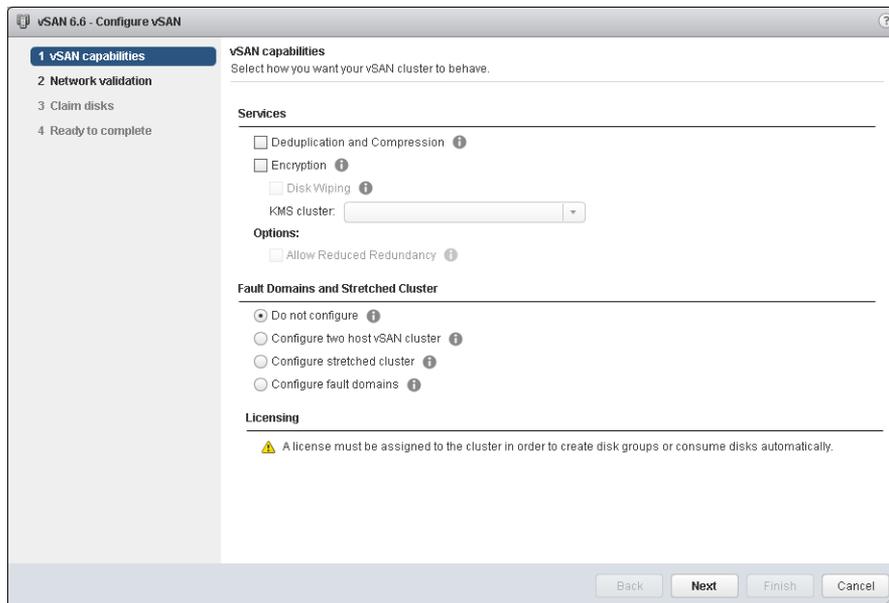
Verify that the vSAN datastore has been created. See [View vSAN Datastore](#).

Verify that the vSAN storage provider is registered. See [View vSAN Storage Providers](#).

Claim the storage devices or create disk groups. See *Administering VMware vSAN*.

# Configure a Cluster for vSAN Using the vSphere Web Client

You can use the Configure vSAN wizard to complete the basic configuration of your vSAN cluster.



## Prerequisites

You must create a cluster and add hosts to the cluster before using the Configure vSAN 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 vSAN 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
<b>Do not configure</b>	Default setting used for a single-site vSAN cluster.
<b>2 host vSAN cluster</b>	Provides fault tolerance for a cluster that has two hosts at a remote office, with a witness host in the main office. Set the <b>Primary level of failures to tolerate</b> policy to 1.
<b>Stretched cluster</b>	Supports two active sites, each with an even number of hosts and storage devices, and a witness host at a third site.
<b>Configure fault domains</b>	Supports fault domains that you can use to group vSAN 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 vSAN 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 vSAN 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 and stretched clusters, see *Administering VMware vSAN*.

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

## Enable vSAN on an Existing Cluster

You can edit cluster properties to enable vSAN for an existing cluster.

### Prerequisites

Verify that your environment meets all requirements. See "Requirements for Enabling vSAN" in *Administering VMware vSAN*.

### Procedure

- 1 Navigate to an existing host cluster.
- 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>Services</b>.</li> <li>b (Optional) Enable deduplication and compression on the cluster. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> <li>c (Optional) Enable encryption on the cluster, and select a KMS server. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> <li>d (Optional) Select Allow Reduced Redundancy. If needed, vSAN will reduce the protection level of your VMs, while enabling deduplication and compression or encryption.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>General</b>.</li> <li>b Click the vSAN is turned ON <b>Edit</b> button.</li> <li>c (Optional) If you want to enable deduplication and compression on the cluster, select the <b>Deduplication and compression</b> check box. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> <li>d (Optional) If you want to enable encryption on the cluster, select the <b>Encryption</b> check box, and select a KMS server. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> </ol>

- 3 Click **OK** or **Apply** to confirm your selections.

### What to do next

Claim the storage devices or create disk groups. See *Administering VMware vSAN*.

## Disable vSAN

You can turn off vSAN for a host cluster.

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

**Prerequisites**

Verify that the hosts are in maintenance mode.

**Procedure**

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>Services</b>.</li> <li>b Click <b>Turn Off vSAN</b>.</li> <li>c On the Turn Off vSAN dialog, confirm your selection.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>General</b>.</li> <li>b Click the vSAN is turned ON <b>Edit</b> button.</li> <li>c Deselect the vSAN <b>Turn On</b> check box.</li> </ol>

**Edit vSAN Settings**

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

Edit the settings of an existing vSAN 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 vSAN host cluster.

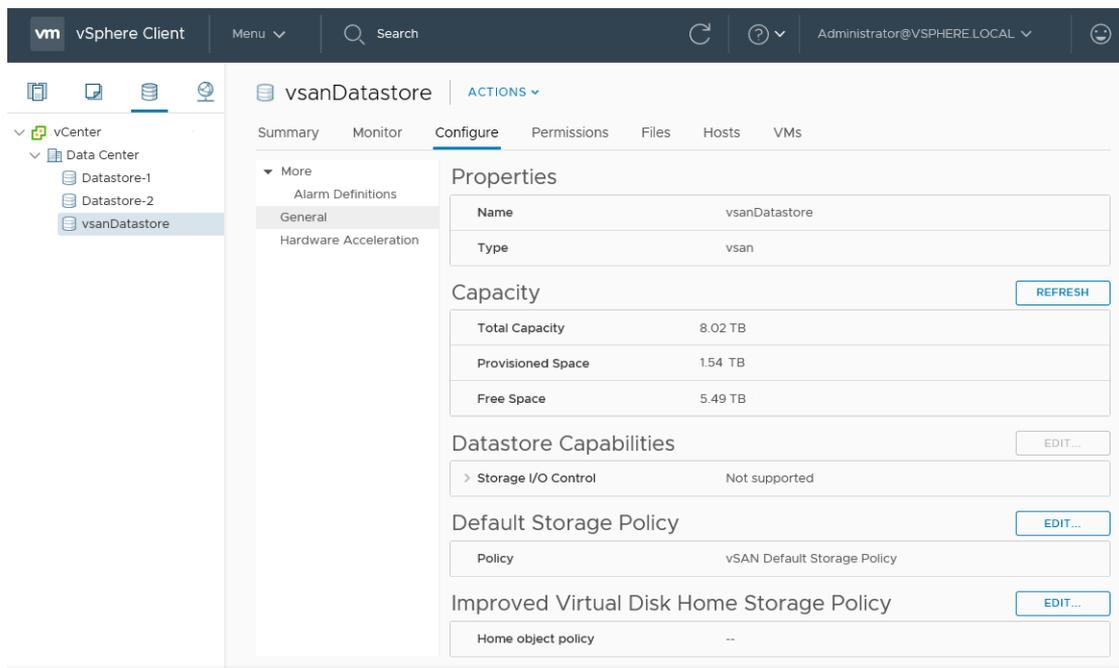
2 Click the **Configure** tab.

Option	Description
<b>vSphere Client</b>	<ol style="list-style-type: none"> <li>Under vSAN, select <b>Services</b>.</li> <li>Click the <b>Edit</b> button for the service you want to configure.                             <ul style="list-style-type: none"> <li>■ Enable or disable deduplication and compression.</li> <li>■ Configure vSAN encryption.</li> <li>■ Configure vSAN performance service.</li> <li>■ Configure iSCSI target service.</li> <li>■ Configure advanced options, such as object repair time, site read locality, thin swap provisioning, and large cluster support.</li> </ul> </li> <li>Modify the settings to match your requirements.</li> </ol>
<b>vSphere Web Client</b>	<ol style="list-style-type: none"> <li>Under vSAN, select <b>General</b>.</li> <li>Click the vSAN is turned ON <b>Edit</b> button.</li> <li>(Optional) If you want to enable deduplication and compression on the cluster, select the <b>Deduplication and compression</b> check box. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> <li>(Optional) If you want to enable encryption on the cluster, select the <b>Encryption</b> check box, and select a KMS server. vSAN will automatically upgrade the on-disk format, causing a rolling reformat of every disk group in the cluster.</li> </ol>

3 Click **OK** or **Apply** to confirm your selections.

## View vSAN Datastore

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



## Prerequisites

Activate vSAN and configure disk groups.

## Procedure

- 1 Navigate to Storage.
- 2 Select the vSAN datastore.
- 3 Click the **Configure** tab.
- 4 Review the vSAN datastore capacity.

The size of the vSAN 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 is  $7 \times 2 \text{ TB} \times 8 = 112 \text{ TB}$ . 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

Create a storage policy for virtual machines using the storage capabilities of the vSAN datastore. For information, see the *vSphere Storage* documentation.

# Using vSAN Policies

When you use vSAN, you can define virtual machine storage requirements, such as performance and availability, in a policy. vSAN ensures that each virtual machine deployed to vSAN datastores is assigned at least one storage policy.

After they are assigned, the storage policy requirements are pushed to the vSAN layer when a virtual machine is created. The virtual device is distributed across the vSAN datastore to meet the performance and availability requirements.

vSAN 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 vSAN Policies](#)
- [Host Pinning](#)
- [View vSAN Storage Providers](#)
- [About the vSAN Default Storage Policy](#)
- [Assign a Default Storage Policy to vSAN Datastores](#)
- [Define a Storage Policy for vSAN Using vSphere Client](#)
- [Define a Storage Policy for vSAN Using the vSphere Web Client](#)

## About vSAN Policies

vSAN storage policies 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 vSAN on a host cluster, a single vSAN 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 vSAN 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, vSAN 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 3-1. Storage Policy Rules**

Capability	Description
Primary level of failures to tolerate (PFTT)	<p>Defines the number of host and device failures that a virtual machine object can tolerate. For <math>n</math> failures tolerated, each piece of data written is stored in <math>n+1</math> places, including parity copies if using RAID 5 or RAID 6.</p> <p>When provisioning a virtual machine, if you do not choose a storage policy, vSAN assigns this policy as the default virtual machine storage policy.</p> <p>If fault domains are configured, <math>2n+1</math> fault domains with hosts contributing capacity are required. A host which does not belong to a fault domain is considered its own single-host fault domain.</p> <p>Default value is 1. Maximum value is 3.</p> <p><b>Note</b> If you do not want vSAN to protect a single mirror copy of virtual machine objects, you can specify <b>PFTT</b> = 0. However, the host might experience unusual delays when entering maintenance mode. The delays occur because vSAN must evacuate the object from the host for the maintenance operation to complete successfully. Setting <b>PFTT</b> = 0 means that your data is unprotected, and you might lose data when the vSAN cluster encounters a device failure.</p> <p><b>Note</b> If you create a storage policy and you do not specify a value for <b>PFTT</b>, vSAN 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.</p> <p>In a stretched cluster, this rule defines the number of site failures that a virtual machine object can tolerate. You can use <b>PFTT</b> with the <b>SFTT</b> to provide local fault protection for objects within your data sites.</p> <p>The maximum value for a stretched cluster is 1.</p>
Secondary level of failures to tolerate (SFTT)	<p>In a stretched cluster, this rule defines the number of additional host failures that the object can tolerate after the number of site failures defined by <b>PFTT</b> is reached. If <b>PFTT</b> = 1 and <b>SFTT</b> = 2, and one site is unavailable, then the cluster can tolerate two additional host failures.</p> <p>Default value is 1. Maximum value is 3.</p>
Data Locality	<p>In a stretched cluster, this rule is available only if the <b>Primary level of failures to tolerate</b> is set to 0. You can set the <b>Data Locality</b> rule to <b>None</b>, <b>Preferred</b>, or <b>Secondary</b>. This rule enables you to limit virtual machine objects to a selected site or host in the stretched cluster.</p> <p>Default value is None.</p>

**Table 3-1. Storage Policy Rules (Continued)**

Capability	Description
Failure tolerance method	<p>Specifies whether the data replication method optimizes for Performance or Capacity. If you select <b>RAID-1 (Mirroring) - Performance</b>, vSAN uses more disk space to place the components of objects but provides better performance for accessing the objects. If you select <b>RAID-5/6 (Erasure Coding) - Capacity</b>, vSAN uses less disk space, but the performance is reduced. You can use RAID 5 by applying the <b>RAID-5/6 (Erasure Coding) - Capacity</b> attribute to clusters with four or more fault domains, and set the <b>Primary level of failures to tolerate</b> to 1. You can use RAID 6 by applying the <b>RAID-5/6 (Erasure Coding) - Capacity</b> attribute to clusters with six or more fault domains, and set the <b>Primary level of failures to tolerate</b> to 2.</p> <p>In stretched clusters with <b>Secondary level of failures to tolerate</b> configured, this rule is applies only to the <b>Secondary level of failures to tolerate</b>.</p> <p>For more information about RAID 5 or RAID 6, see <a href="#">Using RAID 5 or RAID 6 Erasure Coding</a>.</p>
Number of disk stripes per object	<p>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.</p> <p>Default value is 1. Maximum value is 12.</p> <p>Do not change the default striping value.</p> <p>In a hybrid environment, the disk stripes are spread across magnetic disks. In the case of an all-flash configuration, the striping is across flash devices that make up the capacity layer. Make sure that your vSAN environment has sufficient capacity devices present to accommodate the request.</p>
Flash read cache reservation	<p>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.</p> <p>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.</p> <p>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.</p> <p>Default value is 0%. Maximum value is 100%.</p> <p><b>Note</b> By default, vSAN 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.</p> <p>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.</p>

**Table 3-1. Storage Policy Rules (Continued)**

Capability	Description
Force provisioning	<p>If the option is set to <b>Yes</b>, the object is provisioned even if the <b>Primary level of failures to tolerate</b>, <b>Number of disk stripes per object</b>, and <b>Flash read cache reservation</b> 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.</p> <p>The default <b>No</b> is acceptable for most production environments. vSAN fails to provision a virtual machine when the policy requirements are not met, but it successfully creates the user-defined storage policy.</p>
Object space reservation	<p>Percentage of the logical size of the virtual machine disk (vmdk) object that must be reserved, or thick provisioned when deploying virtual machines.</p> <p>Default value is 0%. Maximum value is 100%.</p>
Disable object checksum	<p>If the option is set to <b>No</b>, the object calculates checksum information to ensure the integrity of its data. If this option is set to <b>Yes</b>, the object does not calculate checksum information.</p> <p>vSAN 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, vSAN repairs the data or reports the error.</p> <p>If a checksum mismatch is detected, vSAN automatically repairs the data by overwriting the incorrect data with the correct data. Checksum calculation and error-correction are performed as background operations.</p> <p>The default setting for all objects in the cluster is <b>No</b>, which means that checksum is enabled.</p>
IOPS limit for object	<p>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.</p> <p>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 <b>IOPS limit for object</b> is set to 0, IOPS limits are not enforced.</p> <p>vSAN allows the object to double the rate of the IOPS limit during the first second of operation or after a period of inactivity.</p>

When working with virtual machine storage policies, you must understand how the storage capabilities affect the consumption of storage capacity in the vSAN cluster. For more information about designing and sizing considerations of storage policies, see "Designing and Sizing a vSAN Cluster" in *Administering VMware vSAN*.

## Host Pinning

The vSAN Host Pinning storage policy enables you to store a single copy of data on a VM's local host.

The vSAN Host Pinning storage policy adapts the efficiency and resiliency of vSAN for next-generation, shared-nothing applications. When you use this policy, vSAN maintains a single copy of the data, which is stored on the local host running the VM. This policy is offered as a deployment choice for Big Data (Hadoop, Spark), NoSQL, and other such applications that maintain data redundancy at the application layer.

vSAN Host Pinning has specific requirements and guidelines that require VMware validation to ensure proper deployment. The vSAN Host Pinning policy must be applied to all of the VMs in the cluster, and cannot be combined with other policies on the same cluster. vSAN Encryption and Deduplication cannot be used with the vSAN Host Pinning policy. vSphere DRS and HA options must be turned off to prevent automated movement of VMs.

Administrators interested in this feature must contact VMware to file an intent to deploy request. VMware will evaluate the request to ensure that your deployment meets the requirements before approving for support and production use. VMware shall not support any deployment that does not have explicit approval. For more information, contact your VMware representative.

## View vSAN Storage Providers

Enabling vSAN automatically configures and registers a storage provider for each host in the vSAN cluster.

vSAN storage providers are built-in software components that communicate datastore capabilities to vCenter Server. A storage capability typically is 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 vSAN storage providers report a set of underlying storage capabilities to vCenter Server. They also communicate with the vSAN layer to report the storage requirements of the virtual machines. For more information about storage providers, see the *vSphere Storage* documentation.

vSAN registers a separate storage provider for each host in the vSAN 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 Navigate to vCenter Server.
- 2 Click the **Configure** tab, and click **Storage Providers**.

The storage providers for vSAN 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 vSAN. To remove or unregister the vSAN storage providers, remove corresponding hosts from the vSAN cluster and then add the hosts back. Make sure that at least one storage provider is active.

---

## About the vSAN Default Storage Policy

vSAN requires that the virtual machines deployed on the vSAN 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 vSAN Default Storage Policy is assigned to the virtual machine.

The default policy contains vSAN rule sets and a set of basic storage capabilities, typically used for the placement of virtual machines deployed on vSAN datastores.

**Table 3-2. vSAN 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 the read cache	0
Object space reservation	0
	<b>Note</b> Setting the Object space reservation to zero means that the virtual disk is thin provisioned, by default.
Force provisioning	No

You can review the configuration settings for the default virtual machine storage policy when you navigate to the **VM Storage Policies > vSAN 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 Storage Policy for vSAN Using vSphere Client](#).

When you assign a user-defined storage policy to a datastore, vSAN 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 vSAN datastore.

## Characteristics

The following characteristics apply to the vSAN Default Storage Policy.

- The vSAN default storage policy is assigned to all virtual machine objects if you do not assign any other vSAN policy when you provision a virtual machine. The **VM Storage Policy** text box 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 vSAN Default Storage Policy with **Force provisioning** set to **Yes**.

---

- The vSAN default policy only applies to vSAN datastores. You cannot apply the default storage policy to non-vSAN datastores, such as NFS or a VMFS datastore.

- Because the default virtual machine storage policy is compatible with any vSAN datastore in the vCenter Server, you can move your virtual machine objects provisioned with the default policy to any vSAN 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 vSAN enabled cluster that contains at least one host. Typically you do not edit the settings of the default storage policy.
- You cannot edit the name and description of the default policy, or the vSAN 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 vSAN.

## Assign a Default Storage Policy to vSAN 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 vSAN datastore meets the requirements of virtual machines in the vSAN cluster.

### Procedure

- 1 Navigate to the vSAN datastore.
- 2 Click **Configure**.
- 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 vSAN datastore.

You can choose from a list of storage policies that are compatible with the vSAN datastore, such as the vSAN Default Storage Policy and user-defined storage policies that have vSAN 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 Storage Policy for vSAN Using vSphere Client](#).

## Define a Storage Policy for vSAN Using vSphere Client

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 vSAN datastore.

The screenshot shows the 'Create VM Storage Policy' dialog box with the 'vSAN' tab selected. The dialog is divided into a left sidebar with five steps: 1 Name and description, 2 Policy structure, 3 vSAN (selected), 4 Storage compatibility, and 5 Review and finish. The main area shows the 'Advanced Policy Rules' tab with the following settings:

- Availability: (tab)
- Advanced Policy Rules: (tab)
- Tags: (tab)
- Number of disk stripes per object: 1 (dropdown)
- IOPS limit for object: 0
- Object space reservation: Thin provisioning (dropdown)
- Flash read cache reservation (%): 0
- Disable object checksum:
- Force provisioning:

At the bottom right, there are three buttons: CANCEL, BACK, and NEXT.

### Prerequisites

- Verify that the vSAN storage provider is available. See [View vSAN Storage Providers](#).
- Required privileges: **Profile-driven storage.Profile-driven storage view** and **Profile-driven storage.Profile-driven storage update**

### Procedure

- 1 Navigate to **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, select Enable rules for "vSAN" storage, and click **Next**.

6 On the vSAN page, define the policy rule set.

a On the Availability tab, define the **Site disaster tolerance** and **Failures to tolerate**.

Availability options define the rules for Primary and Secondary level of failures to tolerate, Data locality, and Failure tolerance method.

- **Site disaster tolerance** defines the type of site failure tolerance used for virtual machine objects.
- **Failures to tolerate** defines the number of host and device failures that a virtual machine object can tolerate, and the data replication method.

For example, if you choose **Dual site mirroring** and **2 failures - RAID-6 (Erasure Coding)**, vSAN configures the following policy rules:

- Primary level of failures to tolerate: 1
  - Secondary level of failures to tolerate: 2
  - Data locality: None
  - Failure tolerance method: RAID-5/6 (Erasure Coding) - Capacity
- b On the Advanced Policy Rules tab, define advanced policy rules, such as number of disk stripes per object and IOPS limits.
- c On the Tags tab, click **Add Tag Rule**, and define the options for your tag rule.

Make sure that the values you provide are within the range of values advertised by storage capabilities of the vSAN datastore.

7 Click **Next**.

8 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 vSAN datastore meets the requirements set in the storage policy and that it appears on the list of compatible datastores.

9 On the Review and finish page, review the policy settings, and click **Finish**.

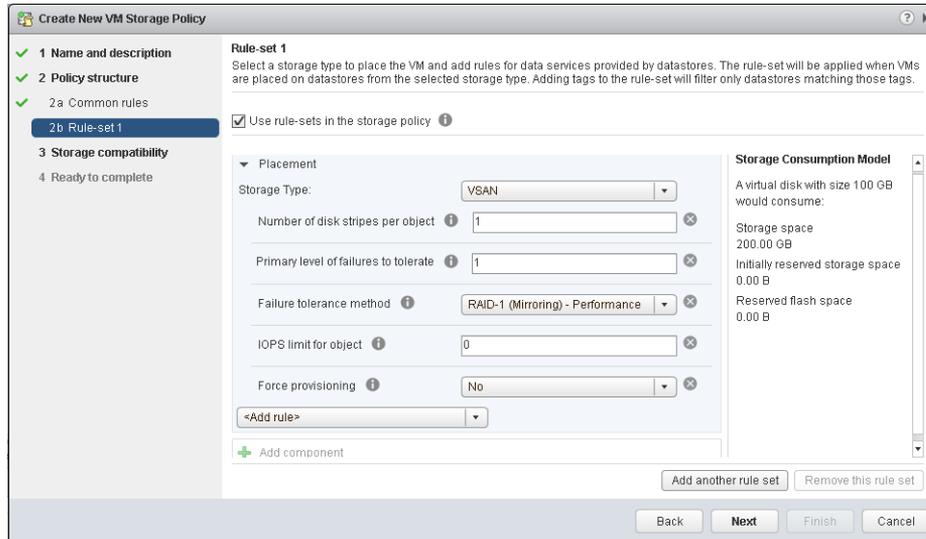
The new policy is added to the list.

#### What to do next

Assign this policy to a virtual machine and its virtual disks. vSAN places the virtual machine objects according to the requirements specified in the policy. For information about applying the storage policies to virtual machine objects, see the *vSphere Storage* documentation.

# Define a Storage Policy for vSAN Using the vSphere Web Client

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 vSAN datastore.



## Prerequisites

- Verify that the vSAN storage provider is available. See [View vSAN Storage Providers](#).
- 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 menu.

The page expands as you add rules for the vSAN datastore.

- c Select a rule from the **Add rule** drop-down menu.

Make sure that the values you provide are within the range of values advertised by storage capabilities of the vSAN datastore.

From the Storage Consumption model, you can review the virtual disk size available and the corresponding cache and capacity requirements, including the reserved storage space your virtual machines might potentially consume when you apply the 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 vSAN 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. vSAN places the virtual machine objects according to the requirements specified in the policy. For information about applying the storage policies to virtual machine objects, see the *vSphere Storage* documentation.

# Expanding and Managing a vSAN Cluster

# 4

After you have set up your vSAN cluster, you can add hosts and capacity devices, remove hosts and devices, and manage failure scenarios.

This chapter includes the following topics:

- [Expanding a vSAN Cluster](#)
- [Working with Maintenance Mode](#)
- [Managing Fault Domains in vSAN Clusters](#)
- [Using the vSAN iSCSI Target Service](#)
- [Migrate a Hybrid vSAN Cluster to an All-Flash Cluster](#)
- [Power off a vSAN Cluster](#)

## Expanding a vSAN Cluster

You can expand an existing vSAN cluster by adding hosts or adding devices to existing hosts, without disrupting any ongoing operations.

Use one of the following methods to expand your vSAN cluster.

- Add new ESXi hosts to the cluster that are configured using the supported cache and capacity devices. See [Add a Host to the vSAN Cluster](#). When you add a device or add a host with capacity, vSAN does not automatically distribute data to the newly added device. To enable vSAN to distribute data to recently-added devices, you must manually rebalance the cluster by using the Ruby vSphere Console (RVC). See "Manual Rebalance" in *vSAN Monitoring and Troubleshooting*.
- Move existing ESXi hosts to the vSAN cluster by using host profile. See [Configuring Hosts Using Host Profile](#). 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 vSAN Host](#).

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 vSAN 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](#).

## Expanding vSAN Cluster Capacity and Performance

If your vSAN 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 adding disk groups. New disk groups require flash devices for the cache. For information about adding devices to disk groups, see [Add Devices to the Disk Group](#). 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 vSAN" in *Administering VMware vSAN*.
- 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 host. Or you can add one or more hosts with disk groups to produce the same performance impact after vSAN completes a proactive rebalance in the vSAN cluster.

Although compute-only hosts can exist in a vSAN cluster, and consume capacity from other hosts in the cluster, add uniformly configured hosts for efficient operation. For best results, add hosts with cache and capacity devices to expand the cluster capacity. Although it is best to use the same or similar devices in your disk groups, any device listed on the vSAN HCL is supported. Try to distribute capacity evenly across hosts and disk groups. For information about adding devices to disk groups, see [Add Devices to the Disk Group](#).

After you expand the cluster capacity, perform a manual rebalance to distribute resources evenly across the cluster. For more information, see "Manual Rebalance" in *vSAN Monitoring and Troubleshooting*.

## Use Quickstart to Add Hosts to a vSAN Cluster

If you configured your vSAN cluster through Quickstart, you can use the Quickstart workflow to add hosts and storage devices to the cluster.

When you add new hosts to the vSAN cluster, you can use the Cluster configuration wizard to complete the host configuration. For more information about Quickstart, see "Using Quickstart to Configure and Expand a vSAN Cluster" in *vSAN Planning and Deployment*.

---

**Note** If you are running vCenter Server on a host in the cluster, the host does not need to be placed into maintenance mode as you add it to a cluster using the Quickstart workflow. The host that contains the vCenter Server VM must be running ESXi 6.5 EP2 or later. The same host also can be running a Platform Services Controller. All other VMs on the host must be powered off.

---

### Prerequisites

The Quickstart workflow must be available for your vSAN cluster.

### Procedure

- 1 Navigate to the cluster in the vSphere Client.

- 2 Click the Configure tab, and select **Configuration > Quickstart**.
- 3 On the Add hosts card, click **Add** to open the Add hosts wizard.
  - a On the Add hosts page, enter information for new hosts, or click Existing hosts and select from hosts listed in the inventory.
  - b On the Host summary page, verify the host settings.
  - c On the Ready to complete page, click **Finish**.
- 4 On the Cluster configuration card, click **Configure** to open the Cluster configuration wizard.
  - a (Optional) On the vMotion traffic page, enter IP address information for vMotion traffic.
  - b On the Storage traffic page, enter IP address information for storage traffic.
  - c (optional) On the Claim disks page, select disks on each new host.
  - d (optional) On the Create fault domains page, move the new hosts into their corresponding fault domains.
 

For more information about fault domains, see [Managing Fault Domains in vSAN Clusters](#).
  - e On the Ready to complete page, verify the cluster settings, and click **Finish**.

## Add a Host to the vSAN Cluster

You can add an ESXi host to a running vSAN 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 vSAN 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 vSAN datastore with excessive virtual machine deployments.

### Procedure

- 1 Navigate to the vSAN cluster.
- 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 in to the host.
 

You can configure this option later by editing the Security Profile in host settings.

## 8 Review the settings and click **Finish**.

The host is added to the cluster.

### What to do next

Verify that the vSAN Disk Balance health check is green. If the Disk Balance health check issues a warning, perform a manual rebalance operation during off-peak hours. For more information, see "Manual Rebalance" in *vSAN Monitoring and Troubleshooting*.

## Configuring Hosts Using Host Profile

When you have multiple hosts in the vSAN cluster, you can use the profile of an existing vSAN host to configure the rest of the hosts in the vSAN cluster.

The host profile includes information about storage configuration, network configuration, and other characteristics of the host. If you are planning to create a cluster with many hosts, such as 8, 16, 32, or 64 hosts, use the host profile feature. Host profiles enable you to add more than one host at a time to the vSAN 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 vSAN 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 vSAN 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 vSAN 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.

- 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 vSAN cluster and its resources are accessible to the vSAN cluster. The host can also access all existing vSAN storage I/O policies in the vSAN cluster.

## Working with Maintenance Mode

Before you shut down, reboot, or disconnect a host that is a member of a vSAN cluster, you must put the host in maintenance mode.

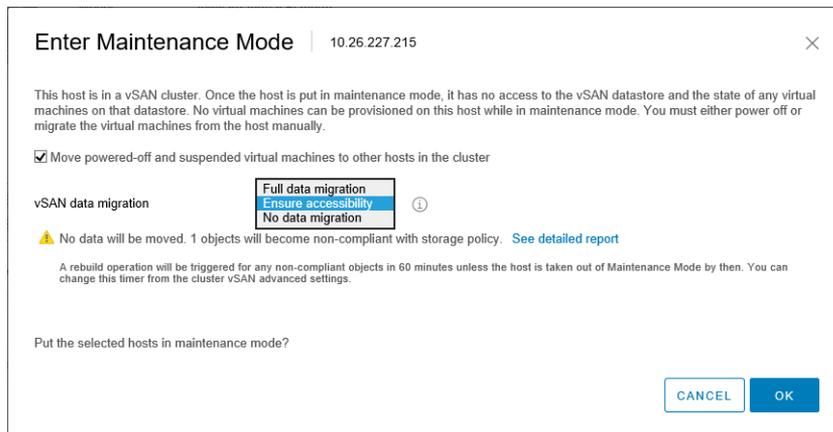
When working with maintenance mode, consider the following guidelines:

- 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 vSAN 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 might 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 the cluster has enough hosts and capacity available to meet the **Primary level of failures to tolerate** policy requirements.

- Verify that enough flash capacity exists on the remaining hosts to handle any flash read cache reservations. To analyze the current capacity use per host, and whether a single host failure might cause the cluster to run out of space and impact the cluster capacity, cache reservation, and cluster components, run the following RVC command: `vsan.whatif_host_failures`. 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 that 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 vSAN Cluster in Maintenance Mode

Before you shut down, reboot, or disconnect a host that is a member of a vSAN 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 accessibility** or **Full data migration**.

When any member host of a vSAN cluster enters maintenance mode, the cluster capacity is automatically reduced, because the member host no longer contributes capacity to the cluster.

Any vSAN iSCSI targets served by this host are transferred to other hosts in the cluster, and thus the iSCSI initiator are redirected to the new target owner.

### 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
<b>Ensure accessibility</b>	<p>This is the default option. When you power off or remove the host from the cluster, vSAN 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 remove the host from the cluster permanently.</p> <p>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. If a failure occurs while the host is in maintenance mode and the <b>Primary level of failures to tolerate</b> is set to 1, you might experience data loss in the cluster.</p> <p><b>Note</b> This is the only evacuation mode available if you are working with a three-host cluster or a vSAN cluster configured with three fault domains.</p>
<b>Full data migration</b>	<p>vSAN evacuates all data to other hosts in the cluster, maintains or fixes availability compliance for the affected components, and protects data when sufficient resources exist 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 that you migrate the virtual machines to another datastore and then place the host in maintenance mode.</p> <p>This evacuation mode results in the largest amount of data transfer and consumes the most time and resources. All the components on the local storage of the selected host are migrated elsewhere in the cluster. When the host enters maintenance mode, all virtual machines have access to their storage components and are still compliant with their assigned storage policies.</p> <p><b>Note</b> If a virtual machine object that has data on the host is not accessible and is not fully evacuated, the host cannot enter the maintenance mode.</p>
<b>No data migration</b>	<p>vSAN does not evacuate any data from this host. If you power off or remove the host from the cluster, some virtual machines might become inaccessible.</p>

A cluster with three fault domains has the same restrictions that a three-host cluster has, such as the inability to use **Full data migration** mode or to reprotect data after a failure.

### What to do next

You can track the progress of data migration in the cluster. See "Monitoring the Resynchronization Tasks in the vSAN Cluster" in *vSAN Monitoring and Troubleshooting*.

## Managing Fault Domains in vSAN Clusters

Fault domains enable you to protect against rack or chassis failure if your vSAN cluster spans across multiple racks or blade server chassis. You can create fault domains and add one or more hosts to each fault domain.

A fault domain consists of one or more vSAN hosts grouped according to their physical location in the data center. When configured, fault domains enable vSAN 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. When a virtual machine is configured with the **Primary level of failures to tolerate** set to 1 (PFTT=1), vSAN 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, vSAN 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), vSAN 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 to support PFTT=1. 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 vSAN Fault Domains" in *vSAN Planning and Deployment*.

Consider a scenario where you have a vSAN cluster with 16 hosts. The hosts are spread across four racks, that is, four hosts per rack. To tolerate an entire rack failure, create a fault domain for each rack. You can configure a cluster of such capacity with the **Primary level of failures to tolerate** set to 1. If you want the **Primary level of failures to tolerate** set to 2, configure five fault domains in the 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, configure fault domains of smaller sizes. Increasing the number of fault domains 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 vSAN 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 vSAN host to a fault domain. If you decide to use fault domains to protect the vSAN environment, consider creating equal sized fault domains.
- When moved to another cluster, vSAN hosts retain their fault domain assignments.
- When designing a fault domain, place a uniform number of hosts in each fault domain.

For guidelines about designing fault domains, see "Designing and Sizing vSAN Fault Domains" in *vSAN Planning and Deployment*.

- 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 vSAN 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, vSAN distributes protection components, such as witnesses and replicas of the virtual machine objects across different fault domains. As a result, the vSAN 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. vSAN 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 vSAN 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 vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Fault Domains**.
- 4 Click the **Create a new fault domain** icon (+).
- 5 Enter 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 **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 vSAN cluster.

### Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Fault Domains**.
- 4 Select the fault domain and click the **Move hosts into selected fault domains** 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 **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 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 vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Fault Domains**.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>a Expand an existing fault domain to view its member hosts.</li> <li>b Select the host that you want to move and click the <b>Move hosts out of fault domain</b> icon. You also can click and drag the host out of the fault domain.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>a Select the host that you want to move and click the <b>Move hosts out of fault domain</b> icon.</li> <li>b Click <b>Yes</b>.</li> </ol>

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 Host into Selected Fault Domain](#).

## Rename a Fault Domain

You can change the name of an existing fault domain in your vSAN cluster.

### Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.

### 3 Under vSAN, click **Fault Domains**.

Option	Description
vSphere Client	a Select the fault domain, and choose menu <b>Actions &gt; Rename</b> . b Enter a new fault domain name.
vSphere Web Client	a Select the fault domain, and click the <b>Rename selected fault domain</b> icon. b Enter a new fault domain name.

### 4 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 vSAN cluster.

### Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Fault Domains**.

Option	Description
vSphere Client	a Select the fault domain, and choose menu <b>Actions &gt; Delete</b> . b Click <b>Yes</b> to confirm.
vSphere Web Client	a Select the fault domain that you want to delete and click the <b>Remove selected fault domains</b> icon (✖). b Click <b>Yes</b> to confirm.

All hosts in the fault domain are removed and the selected fault domain is deleted from the vSAN cluster. Each host that is not part of a fault domain is considered to reside in its own single-host fault domain.

## Using the vSAN iSCSI Target Service

Use the iSCSI target service to enable hosts and physical workloads that reside outside the vSAN cluster to access the vSAN 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 vSAN cluster. vSAN 6.7 and later releases support Windows Server Failover Clustering (WSFC), so WSFC nodes can access vSAN iSCSI targets.

After you configure the vSAN iSCSI target service, you can discover the vSAN iSCSI targets from a remote host. To discover vSAN iSCSI targets, use the IP address of any host in the vSAN cluster, and the TCP port of the iSCSI target. To ensure high availability of the vSAN 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** vSAN iSCSI target service does not support other vSphere or ESXi clients or initiators, third-party hypervisors, or migrations using raw device mapping (RDMs).

---

vSAN iSCSI target service supports the following CHAP authentication methods:

<b>CHAP</b>	In CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target.
<b>Mutual CHAP</b>	In mutual CHAP authentication, an extra level of security enables the initiator to authenticate the target.

For more information about using the vSAN iSCSI target service, refer to the [iSCSI target usage guide](#).

## iSCSI Targets

You can add one or more iSCSI targets that provide storage blocks as logical unit numbers (LUNs). vSAN 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 vSAN storage policy to each LUN, and enable the iSCSI target service on a vSAN cluster. You can configure a storage policy to use as the default policy for the home object of the vSAN 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 vSAN cluster.

### Procedure

- 1 Navigate to the vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
<b>vSphere Client</b>	<ul style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Target Service</b>.</li> <li>b Click to <b>Enable</b> the vSAN iSCSI target service.</li> <li>c Edit the vSAN iSCSI target service configuration. You can select the default network, TCP port, and Authentication method at this time. You also can select a vSAN storage policy.</li> </ul>
<b>vSphere Web Client</b>	<ul style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Targets</b>.</li> <li>b Click the vSAN iSCSI target service <b>Edit</b> button.</li> <li>c Select the Enable vSAN iSCSI target service check box. You can select the default network, TCP port, and Authentication method at this time. You also can select a vSAN Storage Policy.</li> </ul>

## 3 Click **OK** or **Apply**.

### 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 vSAN cluster.
- 2 Click the **Configure** tab.

Option	Description
<b>vSphere Client</b>	<ul style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Target Service</b>.</li> <li>b Click the iSCSI Targets tab.</li> <li>c Click <b>Add</b>. The <b>New iSCSI Target</b> dialog box is displayed. If you leave the target IQN field blank, the IQN is generated automatically.</li> <li>d Enter a target alias. You also can edit the network, TCP port, and authentication method for this target.</li> </ul>
<b>vSphere Web Client</b>	<ul style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Targets</b>.</li> <li>b In the vSAN iSCSI Targets section, click the <b>Add a new iSCSI target</b> icon. The New iSCSI Target dialog box is displayed. The target IQN is generated automatically.</li> <li>c Enter a target alias. You also can edit the network, TCP port, and authentication method for this target.</li> <li>d (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.</li> </ul>

- 3 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 vSAN cluster.
- 2 Click the **Configure** tab.

Option	Description
<b>vSphere Client</b>	<ol style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Target Service</b>.</li> <li>b Click the iSCSI Targets tab, and select a target.</li> <li>c In the vSAN iSCSI LUNs section, click <b>Add</b>. The <b>Add LUN to Target</b> dialog box is displayed.</li> <li>d Enter the size of the LUN. The vSAN Storage Policy configured for the iSCSI target service is assigned automatically. You can assign a different policy to each LUN.</li> </ol>
<b>vSphere Web Client</b>	<ol style="list-style-type: none"> <li>a Under vSAN, click <b>iSCSI Targets</b>.</li> <li>b Select a target, and select the LUNs tab in the Target Details section of the page.</li> <li>c Click the <b>Add a new iSCSI LUN to the target</b> icon. The Add LUN to Target dialog box is displayed.</li> <li>d Enter the size of the LUN. The vSAN Storage Policy configured for the iSCSI target service is assigned automatically. You can assign a different policy to each LUN.</li> </ol>

- 3 Click **Add**.

## 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 vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
<b>vSphere Client</b>	<p>a Under vSAN, click <b>iSCSI Target Service</b>.</p> <p>b Click the Initiator Groups tab, and click the <b>Add a new iSCSI initiator group</b> (+) icon. The <b>New Initiator Group</b> dialog box is displayed.</p> <p>c Enter a name for the iSCSI initiator group.</p> <p>d (Optional) To add members to the initiator group, enter the IQN of each member. Use the following format to enter the member IQN:</p> <p><i>iqn.YYYY-MM.domain:name</i></p> <p>Where:</p> <ul style="list-style-type: none"> <li>■ YYYY = year, such as 2016</li> <li>■ MM = month, such as 09</li> <li>■ domain = domain where the initiator resides</li> <li>■ name = member name (optional)</li> </ul>
<b>vSphere Web Client</b>	<p>a Under vSAN, click <b>iSCSI Initiator Groups</b>.</p> <p>b In the vSAN iSCSI Initiator Groups section, click the <b>Add a new iSCSI initiator group</b> icon. The New vSAN iSCSI Initiator Group dialog box is displayed.</p> <p>c Enter a name for the iSCSI initiator group.</p> <p>d (Optional) To add members to the initiator group, enter the IQN of each member. Use the following format to enter the member IQN:</p> <p><i>iqn.YYYY-MM.domain:name</i></p> <p>Where:</p> <ul style="list-style-type: none"> <li>■ YYYY = year, such as 2016</li> <li>■ MM = month, such as 09</li> <li>■ domain = domain where the initiator resides</li> <li>■ name = member name (optional)</li> </ul>

## 3 Click **OK** or **Create**.

### 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 vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>Under vSAN, click <b>iSCSI Target Service</b>.</li> <li>Select the <b>Initiator Groups</b> tab.</li> <li>In the Accessible Targets section, click the <b>Add a new accessible target for iSCSI Initiator group (+)</b> icon. The <b>Add Accessible Targets</b> dialog box is displayed.</li> <li>Select a target from the list of available targets.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>Under vSAN, click <b>iSCSI Targets</b>.</li> <li>Select the Initiator Groups tab.</li> <li>In the Group Details section, select the Accessible Targets tab.</li> <li>Click the <b>Add a new accessible target for iSCSI Initiator group</b> icon. The Add Accessible Targets dialog box is displayed.</li> <li>On the Filter tab, select a target from the list of available targets. The Selected Objects tab displays the currently selected targets.</li> </ol>

## 3 Click **Add**.

# Monitor vSAN 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 vSAN iSCSI target service and created targets and LUNs.

### Procedure

- ◆ Browse to the vSAN cluster.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>Click <b>Monitor</b> and select <b>Virtual Objects</b>. iSCSI targets are listed on the page.</li> <li>Select a target and click <b>View Placement Details</b>. The Physical Placement shows where the data components of the target are located.</li> <li>Click <b>Group components by host placement</b> to view the hosts associated with the iSCSI data components.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>Click <b>Monitor</b> and select <b>vSAN</b>.</li> <li>Click <b>iSCSI Targets</b>. iSCSI targets and LUNs are listed at the top of the page.</li> <li>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.</li> <li>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 LUN are located. The Compliance Failures tab shows failed components.</li> </ol>

## Migrate a Hybrid vSAN Cluster to an All-Flash Cluster

You can migrate the disk groups in a hybrid vSAN cluster to all-flash disk groups.

The vSAN 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 vSAN cluster.
- 2 Remove the hybrid disk groups for each host in the cluster.
  - a Click the **Configure** tab.
  - b Under vSAN, 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 vSAN Cluster

You can power off a vSAN cluster to perform maintenance or upgrades.

### Prerequisites

If the vCenter Server VM is running on the vSAN 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 vSAN cluster.
 

If the vCenter Server is running on the vSAN cluster, the vCenter Server VM must be powered off last.
- 2 Place all ESXi hosts that compose the cluster in maintenance mode.
  - Right-click each host, and select **Maintenance Mode > Enter Maintenance Mode**.
  - Choose the **No data migration** option.
- 3 Power off the ESXi hosts.

# Device Management in a vSAN Cluster

# 5

You can perform various device management tasks in a vSAN cluster. You can create hybrid or all-flash disk groups, enable vSAN 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](#)
- [Working with Individual Devices](#)

## Managing Disk Groups and Devices

When you enable vSAN on a cluster, choose a disk-claiming mode to organize devices into groups.

vSAN 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 specify each host and each device to be used for the vSAN 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. The ratio depends on the requirements and workload of the cluster. For a hybrid cluster, consider using at least 10 percent of flash cache to consumed capacity ratio (not including replicas such as mirrors). For guidance on determining the cache ratio for all-flash clusters, refer to [Designing vSAN Disk groups – All Flash Cache Ratio Update](#).

The vSAN cluster initially contains a single vSAN datastore with zero bytes consumed.

As you create disk groups on each host and add cache and capacity devices, the size of the datastore increases according to the amount of physical capacity added by those devices. vSAN creates a single distributed vSAN 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 vSAN cluster, the local storage from that host is not added to the vSAN datastore automatically. You have to create a disk group and add the devices to the disk group to use the new storage from the new ESXi host.

---

### Claim Disks for the vSAN Cluster

You can select multiple devices from your hosts, and vSAN creates default disk groups for you.

When you add more capacity to the hosts or add new hosts with capacity, you can select the new devices to increase the capacity of the vSAN datastore. In an all-flash cluster, you can mark flash devices for use as capacity.

After vSAN has claimed devices, it creates the vSAN 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 vSAN 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 vSAN cluster.
- 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 **Create** or **OK** to confirm your selections.

The new disk group appears in the list.

## Claim Storage Devices for a vSAN Cluster

You can select a group of cache and capacity devices, and vSAN organizes them into default disk groups.

### Procedure

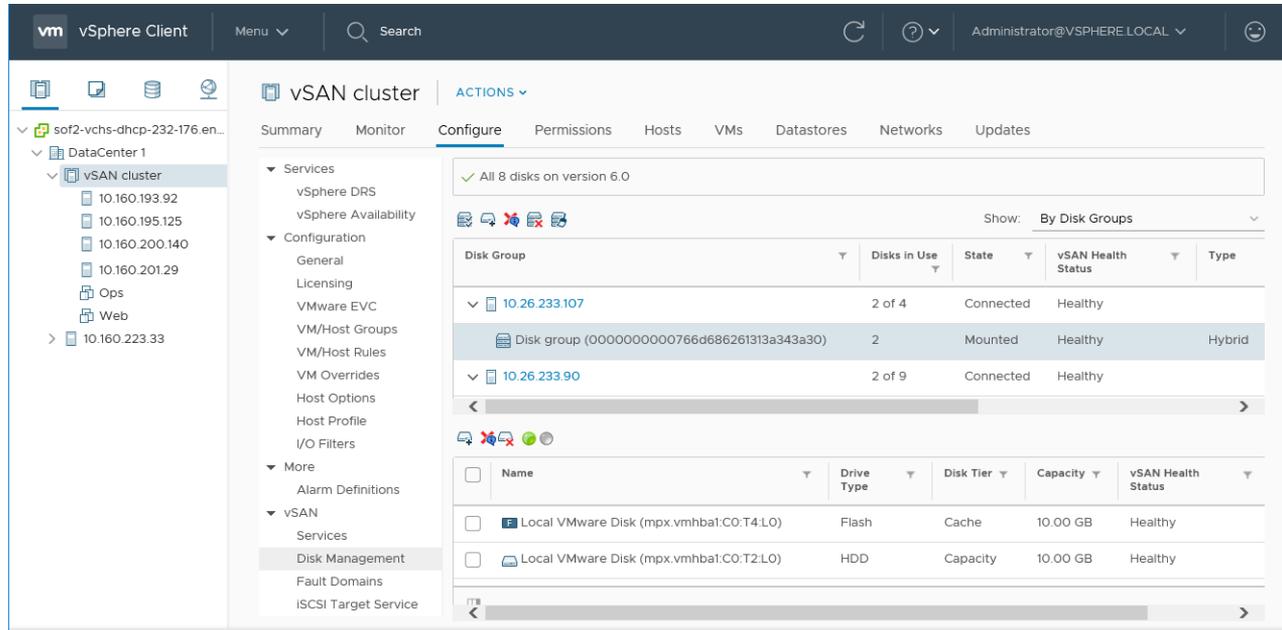
- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.
- 4 Click the **Claim Unused Disks for vSAN** icon ().
- 5 Select devices to add to the disk group.
  - For hybrid disk groups, each host that contributes storage must contribute one flash cache device and one or more HDD capacity devices. You can add only one 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, each host that contributes storage must contribute one flash cache device and one or more flash capacity devices. You can add only one cache device per disk group.
    - 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.

vSAN claims the devices that you selected and organizes them into default disk groups that support the vSAN datastore.

## Working with Individual Devices

You can perform various device management tasks in the vSAN 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 vSAN 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 vSAN cluster.
- 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](#). 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*.

### What to do next

Verify that the vSAN Disk Balance health check is green. If the Disk Balance health check issues a warning, perform a manual rebalance operation during off-peak hours. For more information, see "Manual Rebalance" in *vSAN Monitoring and Troubleshooting*.

## Remove Disk Groups or Devices from vSAN

You can remove selected devices from the disk group or an entire disk group.

Because removing unprotected devices might be disruptive for the vSAN datastore and virtual machines in the datastore, avoid removing devices or disk groups.

Typically, you delete devices or disk groups from vSAN 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 vSAN cluster.

Deleting a disk group permanently deletes the disk membership and 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 vSAN host in maintenance mode by selecting the **Full data migration** option or by selecting **Ensure data accessibility** when deleting a device or a disk group. If you select **No data migration** from the drop-down menu, your data might be at risk if a failure occurs during evacuation.

### Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.

#### 4 Remove a disk group or selected devices.

Option	Description
<b>Remove the Disk Group</b>	<ol style="list-style-type: none"> <li>Under Disk Groups, select the disk group to remove, and click the <b>Remove the disk group</b> icon ()</li> <li>Select a data evacuation mode.</li> </ol>
<b>Remove the Selected Device</b>	<ol style="list-style-type: none"> <li>Under Disk Groups, select the disk group that contains the device that you are removing.</li> <li>Under Disks, select the device to remove, and click the <b>Remove the selected disk(s) from the disk group</b> icon ()</li> <li>Select a data evacuation mode.</li> </ol>

You can move the evacuated data to another disk or disk group on the same host.

#### 5 Click **Yes** or **Delete** to confirm.

The data is evacuated from the selected devices or a disk group and is no longer available to vSAN.

## Recreate a Disk Group

When you recreate a disk group in the vSAN cluster, the existing disks are removed from the disk group, and the disk group is deleted. vSAN recreates the disk group with the same disks.

When you recreate a disk group on a vSAN cluster, vSAN manages the process for you. vSAN evacuates data from all disks in the disk group, removes the disk group, and creates the disk group with the same disks.

### Procedure

- 1 Navigate to the vSAN cluster in the vSphere Client.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.
- 4 Under Disk Groups, select the disk group to recreate, and click the **Recreate the disk group** icon ()

The Recreate Disk Group dialog box appears.

- 5 Select a data migration mode, and click **Recreate**.

All data residing on the disks is evacuated. The disk group is removed from the cluster, and recreated.

## Using Locator LEDs

You can use locator LEDs to identify the location of storage devices.

vSAN 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 vSAN 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 vSAN 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 vSAN cluster.
- 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 <b>Manage</b> tab and click <b>Storage &gt; Storage Devices</b> .
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 <b>Manage</b> tab and click <b>Storage &gt; Storage Devices</b> .

## 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.

Flash devices might not be recognized as flash 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 vSAN and you cannot use them in the vSAN cluster. Marking these devices as local flash makes them available to vSAN.

### 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 vSAN cluster.
- 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 as flash disk** icon ().
- 7 Click **Yes** to save your changes.

The Drive type for the selected devices appears 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 vSAN cluster.
- 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 as HDD disk** icon ().

- 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, vSAN 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 Navigate to the vSAN cluster.
- 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 as local** 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 vSAN does not claim the devices when it creates disk groups.

In vSAN, you cannot add shared devices to a disk group.

### Procedure

- 1 Navigate to the vSAN cluster.
- 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 **Mark as remote** icon.
- 7 Click **Yes** to confirm.

## Add a Capacity Device

You can add a capacity device to an existing vSAN 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 vSAN cluster.
- 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** or **Add**.

The device is added to the disk group.

## Remove Partition From Devices

You can remove partition information from a device so vSAN 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 vSAN use. VMware recommends adding clean devices to disk groups.

When you remove partition information from a device, vSAN 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 vSAN.

### Procedure

- 1 Navigate to the vSAN cluster.
- 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** icon ()

7 Click **OK** to confirm.

The device is clean and does not include any partition information.

# Increasing Space Efficiency in a vSAN Cluster

# 6

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 vSAN Space Efficiency](#)
- [Reclaiming Space with SCSI Unmap](#)
- [Using Deduplication and Compression](#)
- [Using RAID 5 or RAID 6 Erasure Coding](#)
- [RAID 5 or RAID 6 Design Considerations](#)

## Introduction to vSAN 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.

vSAN 6.7 Update 1 and later supports SCSI unmap commands that enable you to reclaim storage space that is mapped to a deleted vSAN object.

You can enable deduplication and compression on a vSAN cluster to eliminate duplicate data and reduce the amount of space required 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 each provide clearly defined space savings over RAID 1. Deduplication and compression can provide additional savings.

## Reclaiming Space with SCSI Unmap

vSAN 6.7 Update 1 and later supports SCSI UNMAP commands that enable you to reclaim storage space that is mapped to a deleted vSAN object.

Deleting or removing files frees space within the file system. This free space is mapped to a storage device until the file system releases or unmaps it. vSAN supports reclamation of free space, which is also called the unmap operation. You can free storage space in the vSAN datastore when you delete or migrate a VM, consolidate a snapshot, and so on.

Reclaiming storage space can provide higher host-to-flash I/O throughput and improve flash endurance.

vSAN also supports the SCSI UNMAP commands issued directly from a guest operating system to reclaim storage space. vSAN supports offline unmaps as well as inline unmaps. On Linux OS, offline unmaps are performed with the `fstrim(8)` command, and inline unmaps are performed when the `mount -o discard` command is used. On Windows OS, NTFS performs inline unmaps by default.

Unmap capability is disabled by default. To enable unmap on a vSAN cluster, use the following RVC command: `vsan.unmap_support -enable`

When you enable unmap on a vSAN cluster, you must power off and then power on all VMs. VMs must use virtual hardware version 13 or above to perform unmap operations.

## Using Deduplication and Compression

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

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. vSAN 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 vSAN 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 vSAN all-flash cluster or when you edit an existing vSAN all-flash cluster. For more information about creating and editing vSAN clusters, see "Enabling vSAN" in *vSAN Planning and Deployment*.

When you enable or disable deduplication and compression, vSAN performs a rolling reformat of every disk group on every host. Depending on the data stored on the vSAN datastore, this process might take a long time. 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.

---

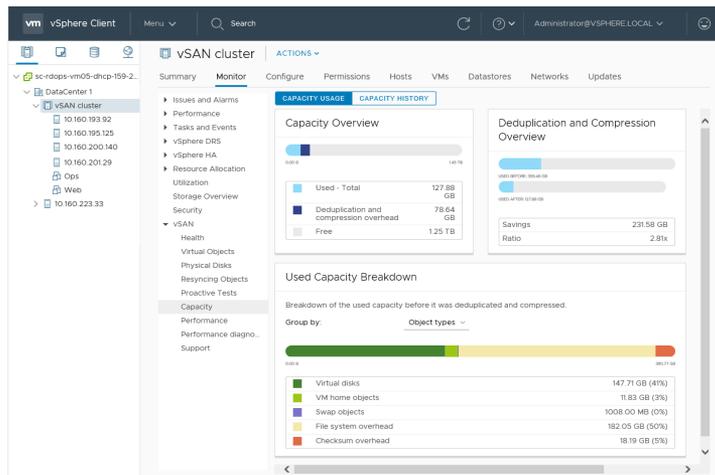
## How to Manage Disks in a Cluster with Deduplication and Compression

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 disk group to increase cluster storage capacity.
- When you add a disk group manually, add all 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 vSAN Capacity monitor.



You can view the Deduplication and Compression Overview when you monitor vSAN capacity in the vSphere 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 vSAN 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 vSAN 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.
- When you enable deduplication and compression on a vSAN cluster, all disk groups participate in data reduction through deduplication and compression.
- vSAN 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 must have either 0 percent or 100 percent object space reservations. Policies with 100 percent object space reservations are always honored, but can make deduplication and compression less efficient.

## Enable Deduplication and Compression on a New vSAN Cluster

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

### Procedure

- 1 Navigate to a new all-flash vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ul style="list-style-type: none"> <li>a Under vSAN, select <b>Services</b>, and click the <b>Configure</b> button.</li> <li>b On the <b>vSAN Capabilities</b> page of the Configure vSAN wizard, enable Deduplication and Compression.</li> <li>c (Optional) Select <b>Allow Reduced Redundancy</b>. If needed, vSAN reduces the protection level of your VMs while enabling Deduplication and Compression. For more details, see <a href="#">Reducing VM Redundancy for vSAN Cluster</a>.</li> </ul>
vSphere Web Client	<ul style="list-style-type: none"> <li>a Under vSAN, select <b>General</b>.</li> <li>b Click the <b>Configure vSAN</b> button.</li> <li>c Configure deduplication and compression on the cluster.               <ul style="list-style-type: none"> <li>1 On the <b>vSAN capabilities</b> page, select the <b>Enable</b> check box under Deduplication and Compression.</li> <li>2 Enable reduced redundancy for your VMs. See <a href="#">Reducing VM Redundancy for vSAN Cluster</a>.</li> </ul> </li> <li>d On the <b>Claim disks</b> page, specify which disks to claim for the vSAN cluster.               <ul style="list-style-type: none"> <li>1 Select a flash device to be used for capacity and click the <b>Claim for capacity tier</b> icon ().</li> <li>2 Select a flash device to be used as cache and click the <b>Claim for cache tier</b> icon (.</li> </ul> </li> </ul>

## 3 Complete your cluster configuration.

# Enable Deduplication and Compression on Existing vSAN Cluster

You can enable deduplication and compression by editing configuration parameters on an existing all-flash vSAN cluster.

### Prerequisites

Create an all-flash vSAN cluster.

### Procedure

- 1 Navigate to the vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>Services</b>.</li> <li>b Click the Deduplication and Compression <b>Edit</b> button.</li> <li>c Enable Deduplication and Compression.</li> <li>d (Optional) Select <b>Allow Reduced Redundancy</b>. If needed, vSAN reduces the protection level of your VMs, while enabling Deduplication and Compression. See <a href="#">Reducing VM Redundancy for vSAN Cluster</a>.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>a Under vSAN, select <b>General</b>.</li> <li>b In the vSAN is turned ON pane, click the <b>Edit</b> button.</li> <li>c Configure deduplication and compression on the cluster.               <ol style="list-style-type: none"> <li>1 Set deduplication and compression to <b>Enabled</b>.</li> <li>2 Enable reduced redundancy for your VMs. See <a href="#">Reducing VM Redundancy for vSAN Cluster</a>.</li> </ol> </li> </ol>

## 3 Click **Apply** or **OK** to save your configuration changes.

While enabling deduplication and compression, vSAN updates the on-disk format of each disk group of the cluster. To accomplish this change, vSAN 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 vSAN cluster.

When deduplication and compression are disabled on the vSAN 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 vSAN cluster.

## 2 Click the **Configure** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> <li>Under vSAN, select <b>Services</b>.</li> <li>Click the Deduplication and Compression <b>Edit</b> button.</li> <li>Disable Deduplication and Compression.</li> <li>(Optional) Select <b>Allow Reduced Redundancy</b>. If needed, vSAN reduces the protection level of your VMs, while enabling Deduplication and Compression. See <a href="#">Reducing VM Redundancy for vSAN Cluster</a>.</li> </ol>
vSphere Web Client	<ol style="list-style-type: none"> <li>Under vSAN, select <b>General</b>.</li> <li>In the vSAN is turned ON pane, click the <b>Edit</b> button.</li> <li>Disable deduplication and compression. <ol style="list-style-type: none"> <li>Set the disk claiming mode to <b>Manual</b>.</li> <li>Set deduplication and compression to <b>Disabled</b>.</li> </ol> </li> </ol>

## 3 Click **Apply** or **OK** to save your configuration changes.

While disabling deduplication and compression, vSAN changes the disk format on each disk group of the cluster. It evacuates data from the disk group, removes the disk group, and recreates it with a format that does not support 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 vSAN 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, vSAN 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 vSAN 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. vSAN restores full compliance and redundancy after the format conversion is completed.

## Adding or Removing Disks with Deduplication and Compression Enabled

When you add disks to a vSAN 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 from a cache tier, the entire disk group is removed. Removing a cache tier disk when deduplication and compression are enabled triggers data evacuation.
- Deduplication and compression are 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 vSAN 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 6-1. 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 vSAN 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 3 Using vSAN Policies](#).

## RAID 5 or RAID 6 Design Considerations

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

# Using Encryption on a vSAN Cluster



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

vSAN 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 is removed from the cluster.

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

vSAN 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 vSAN Encryption Works](#)
- [Design Considerations for vSAN Encryption](#)
- [Set Up the KMS Cluster](#)
- [Enable Encryption on a New vSAN Cluster](#)
- [Generate New Encryption Keys](#)
- [Enable vSAN Encryption on Existing vSAN Cluster](#)
- [vSAN Encryption and Core Dumps](#)

## How vSAN Encryption Works

When you enable encryption, vSAN encrypts everything in the vSAN 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.

vSAN 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. You can specify a password to 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 vSAN health service, under **Physical disks > Software state health**.

## Design Considerations for vSAN Encryption

Consider these guidelines when working with vSAN encryption.

- Do not deploy your KMS server on the same vSAN 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 vSAN 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 Client or using the `vm-support` command.

The password reencrypts 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 vSAN datastore.

Before you can encrypt the vSAN 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 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 vSAN cluster you plan to encrypt. If a failure occurs, hosts in the vSAN 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, and configure all KMS instances to synchronize keys among them. Use the method documented by your KMS vendor.
- 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.
- 2 Browse the inventory list and select the vCenter Server instance.
- 3 Click **Configure** and click **Key Management Servers**.
- 4 Click **Add**, specify the KMS information in the wizard, and click **Add**.

Option	Value
<b>KMS cluster</b>	Select <b>Create new cluster</b> for a new cluster. If a cluster exists, you can select that cluster.
<b>Cluster name</b>	Name for the KMS cluster. You can use this name to connect to the KMS if your vCenter Server instance becomes unavailable.
<b>Server alias</b>	Alias for the KMS. You can use this alias to connect to the KMS if your vCenter Server instance becomes unavailable.
<b>Server address</b>	IP address or FQDN of the KMS.
<b>Server port</b>	Port on which vCenter Server connects to the KMS.
<b>Proxy address</b>	Optional proxy address for connecting to the KMS.

Option	Value
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 Navigate to the vCenter Server.
- 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	<a href="#">Use the Root CA Certificate Option to Establish a Trusted Connection.</a>
Certificate	<a href="#">Use the Certificate Option to Establish a Trusted Connection.</a>
New Certificate Signing Request	<a href="#">Use the New Certificate Signing Request Option to Establish a Trusted Connection.</a>
Upload certificate and private key	<a href="#">Use the Upload Certificate and Private Key Option to Establish a Trusted Connection.</a>

### 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 Navigate to the vCenter Server.
- 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](#).

**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 Navigate to the vCenter Server.
- 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](#).

## 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 Navigate to the vCenter Server.
- 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](#).

## 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 Navigate to the vCenter Server.
- 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](#).

## 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 Navigate to the 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 Navigate to the vCenter Server.
- 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	<ol style="list-style-type: none"> <li>a Click <b>All Actions</b>, and select <b>Refresh KMS certificate</b>.</li> <li>b In the dialog box that appears, click <b>Trust</b>.</li> </ol>
Upload KMS certificate	<ol style="list-style-type: none"> <li>a Click <b>All Actions</b>, and select <b>Upload KMS Certificate</b>.</li> <li>b In the dialog box that appears, click <b>Upload file</b>, upload a certificate file, and click <b>OK</b>.</li> </ol>

## Enable Encryption on a New vSAN Cluster

You can enable encryption when you configure a new vSAN 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.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services** and click the Encryption **Edit** button.
- 4 On the **vSAN Services** dialog, enable **Encryption**, 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 Complete your cluster configuration.

Encryption of data at rest is enabled on the vSAN cluster. vSAN encrypts all data added to the vSAN 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 vSAN cluster.

- If you generate a new KEK, all hosts in the vSAN 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 reformat 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 vSAN host cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.
- 4 Click **Generate New Encryption Keys**.
- 5 To generate a new KEK, click **Apply**. The DEKs are re-encrypted with the new KEK.
  - To generate a new KEK and new DEKs, and re-encrypt all data in the vSAN cluster, select the following check box: **Also re-encrypt all data on the storage using new keys**.
  - If your vSAN 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 vSAN Encryption on Existing vSAN Cluster

You can enable encryption by editing the configuration parameters of an existing vSAN 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 vSAN host cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.
- 4 Click the Encryption **Edit** button.
- 5 On the vSAN Services dialog, enable **Encryption**, and select a KMS cluster.
- 6 (Optional) If the storage devices in your cluster contain sensitive data, select **Erase Disks Before Use**.

This setting directs vSAN to wipe existing data from the storage devices as they are encrypted. This option can increase the time to process each disk, so do not choose it unless you have unwanted data on the disks.

- 7 Click **Apply**.

A rolling reformat of all disk groups takes place as vSAN encrypts all data in the vSAN datastore.

## vSAN Encryption and Core Dumps

If your vSAN 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 vSAN Cluster

If encryption is enabled on a vSAN cluster, any core dumps in the vm-support package are encrypted. You can collect the package, 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 vSAN 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 using the Flex-based 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	<code>crypto-util envelope describe vmmcores.ve</code>
zdump file	<code>crypto-util envelope describe --offset 4096 zdumpFile</code>

- 3 Decrypt the core dump, depending on its type.

Option	Description
Monitor core dump	<code>crypto-util envelope extract vmmcores.ve vmmcores</code>
zdump file	<code>crypto-util envelope extract --offset 4096 zdumpEncrypted zdumpUnencrypted</code>

# Upgrading the vSAN Cluster

Upgrading vSAN 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 vSAN 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](#). For information about migrating and preparing your ESXi hosts for upgrade, see the *vSphere Upgrade* documentation.
- 3 Upgrade the vSAN 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 vSAN. See [Upgrade vSAN Disk Format Using RVC](#).

This chapter includes the following topics:

- [Before You Upgrade vSAN](#)
- [Upgrade the vCenter Server](#)
- [Upgrade the ESXi Hosts](#)
- [About the vSAN Disk Format](#)
- [Verify the vSAN Cluster Upgrade](#)
- [Using the RVC Upgrade Command Options](#)
- [vSAN Build Recommendations for vSphere Update Manager](#)

## Before You Upgrade vSAN

Plan and design your upgrade to be fail-safe. Before you attempt to upgrade vSAN, verify that your environment meets the vSphere hardware and software requirements.

## Upgrade Prerequisite

Consider the aspects that might 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 vSAN 6.7.1.

**Table 8-1. Upgrade Prerequisite**

Upgrade Prerequisites	Description
Software, hardware, drivers, firmware, and storage I/O controllers	Verify that vSAN 6.7.1 supports the software and hardware components, drivers, firmware, and storage I/O controllers that you plan on using. Supported items are listed on the VMware Compatibility Guide website at <a href="http://www.vmware.com/resources/compatibility/search.php">http://www.vmware.com/resources/compatibility/search.php</a> .
vSAN version	Verify that you are using the latest version of vSAN. You cannot upgrade from a beta version to vSAN 6.7.1. When you upgrade from a beta version, you must perform a fresh deployment of vSAN.
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 <i>vSphere Upgrade</i> documentation.
vSAN disk format	<p>Verify that you have enough capacity available to upgrade the disk format. If free space equal to the consumed capacity of the largest disk group is not available, with the space available on disk groups other than the disk groups that are being converted, you must choose <b>Allow reduced redundancy</b> as the data migration option.</p> <p>For example, the largest disk group in a cluster has 10 TB of physical capacity, but only 5 TB is being consumed. An extra 5 TB of spare capacity is needed elsewhere in the cluster, excluding the disk groups that are being migrated. When upgrading the vSAN disk format, verify that the hosts are not in maintenance mode. When any member host of a vSAN cluster enters maintenance mode, the cluster capacity is automatically reduced. 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 <a href="#">Place a Member of vSAN Cluster in Maintenance Mode</a>.</p>
vSAN hosts	<p>Verify that you have placed the vSAN hosts in maintenance mode and selected the <b>Ensure data accessibility</b> or <b>Evacuate all data</b> option.</p> <p>You can use the vSphere Update Manager for automating and testing the upgrade process. However, when you use vSphere Update Manager to upgrade vSAN, the default evacuation mode is <b>Ensure data accessibility</b>. When you use the <b>Ensure data accessibility</b> mode, your data is not protected, and if you encounter a failure while upgrading vSAN, you might experience unexpected data loss. However, the <b>Ensure data accessibility</b> mode is faster than the <b>Evacuate all data</b> mode, because you do not need to move all data to another host in the cluster. For information about various evacuation modes, see the <a href="#">Place a Member of vSAN Cluster in Maintenance Mode</a>.</p>
Virtual Machines	Verify that you have backed up your virtual machines.

## Recommendations

Consider the following recommendations when deploying ESXi hosts for use with vSAN:

- 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 vSAN is not claiming the device.
- When you boot a vSAN 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.
- To ensure your hardware meets the requirements for vSAN, refer to "Hardware Requirement for vSAN" in *vSAN Planning and Deployment*.

vSAN 6.5 and later enables you to adjust the boot size requirements for an ESXi host in a vSAN 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 vSAN 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 vSAN 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 vSAN 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 vSAN cluster upgrade is upgrading the ESXi hosts to use the current version.

If you have multiple hosts in the vSAN 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 vSAN you encounter a failure, you risk losing data. For information about working with evacuation modes, see [Place a Member of vSAN Cluster in Maintenance Mode](#)

For information about using vSphere Update Manager, see the documentation website at [https://www.vmware.com/support/pubs/vum\\_pubs.html](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 website 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 vSAN cluster.

- a Right-click the host 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 a 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 Manually restart your ESXi host.
  - a Navigate to the ESXi host in the 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 vSAN cluster, you can use vSphere Update Manager to upgrade the remaining hosts.

- 5 Exit maintenance mode.

#### What to do next

- 1 (Optional) Upgrade the vSAN disk format. See [Upgrade vSAN Disk Format Using RVC](#).
- 2 Verify the host license. In most cases, you must reapply your host license. 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 Client or vSphere Update Manager.

## About the vSAN Disk Format

The disk format upgrade is optional. Your vSAN 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 vSAN.

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 vSAN cluster. The disk group is then added back to vSAN 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, vSAN 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 8-2. Upgrade Progress**

Percentage of Completion	Description
0%-5%	<p>Cluster check. Cluster components are checked and prepared for the upgrade. This process takes a few minutes. vSAN verifies that no outstanding issues exist that can prevent the upgrade from completing.</p> <ul style="list-style-type: none"> <li>■ All hosts are connected.</li> <li>■ All hosts have the correct software version.</li> <li>■ All disks are healthy.</li> <li>■ All objects are accessible.</li> </ul>
5%-10%	<p>Disk group upgrade. vSAN performs the initial disk upgrade with no data migration. This process takes a few minutes.</p>
10%-15%	<p>Object realignment. vSAN 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.</p>
15%-95%	<p>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 load. A system at or near its I/O capacity transfers slowly.</p>
95%-100%	<p>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 <b>Allow reduced redundancy</b> option is selected.</p>

During the upgrade, you can monitor the upgrade process from the Resyncing Components page. See "Monitor the Resynchronization Tasks in the vSAN Cluster" in *vSAN Monitoring and Troubleshooting*. 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 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 0 (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, vSAN 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. vSAN 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*.

## Upgrading vSAN Disk Format Using vSphere Client

After you have finished upgrading the vSAN hosts, you can perform the disk format upgrade.

The screenshot shows the vSAN cluster configuration page in the vSphere Client. The left sidebar lists various configuration options, with 'vSAN' and 'Disk Management' selected. The main area displays a warning: '6 of 15 disks on older version' and 'Pre-check suggested before upgrading'. There are 'UPGRADE' and 'PRE-CHECK UPGRADE' buttons. Below this is a table showing disk groups and their status.

Disk Group	Disks in Use	State	vSAN Health Status
10.26.235.157	9 of 9	Connected	Healthy
Disk group (0000000000766d686261313a353a30)	3	Mounted	Healthy
Disk group (0000000000766d686261313a343a30)	3	Mounted	Healthy
10.26.235.159	6 of 6	Connected	Healthy
Disk group (0000000000766d686261313a353a30)	3	Mounted	Healthy

Name	Drive Type	Disk Tier
Local VMware Disk (mpx.vmhba1:C0:T5:L0)	Flash	Cache
Local VMware Disk (mpx.vmhba1:C0:T1:L0)	Flash	Capacit
Local VMware Disk (mpx.vmhba1:C0:T9:L0)	Flash	Capacit

3 items

**Note** If you enable encryption or deduplication and compression on an existing vSAN 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 vSAN Settings](#).

### 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 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 website 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 complete 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, do not place the hosts in maintenance mode. When any member host of a vSAN cluster enters maintenance mode, the member host no longer contributes capacity to the cluster. The cluster capacity is reduced and the cluster upgrade might fail.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. For information about vSAN resynchronization, see *vSphere Monitoring and Performance*.

## Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Disk Management**.
- 4 (Optional) 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** text box.

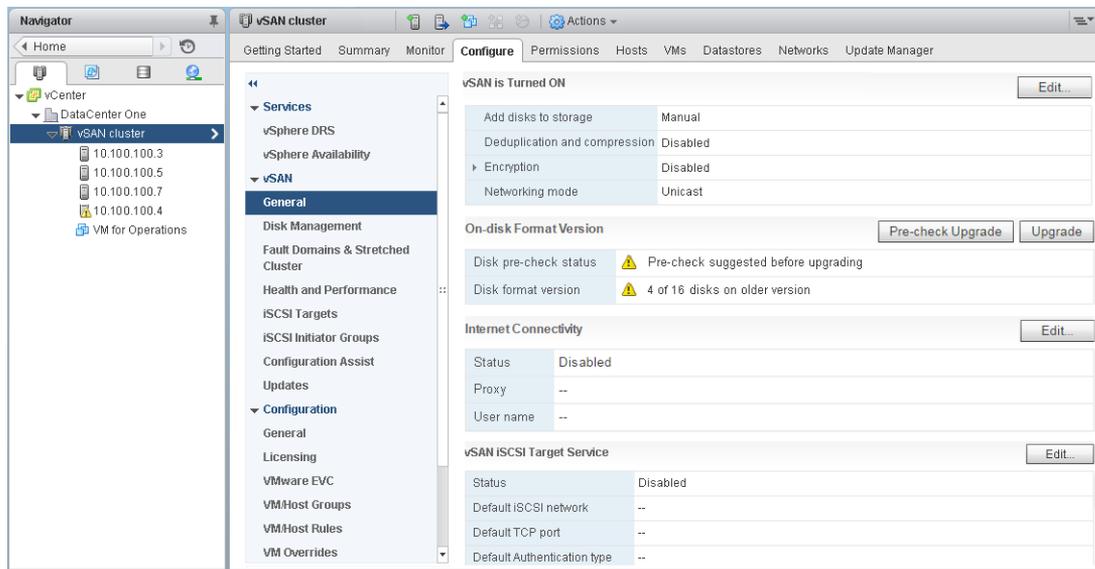
- 5 Click **Upgrade**.
- 6 Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN 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.

If a failure occurs during the upgrade, you can check the Resyncing Objects page. 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 vSAN Disk Format Using vSphere Web Client

After you have finished upgrading the vSAN hosts, you can perform the disk format upgrade.



**Note** If you enable encryption or deduplication and compression on an existing vSAN 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 vSAN Settings](#).

## 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 website 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 complete 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, do not place the hosts in maintenance mode. When any member host of a vSAN cluster enters maintenance mode, the member host no longer contributes capacity to the cluster. The cluster capacity is reduced and the cluster upgrade might fail.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster. See "Monitor the Resynchronization Tasks in the vSAN Cluster in *vSAN Monitoring and Troubleshooting*."

## Procedure

- 1 Navigate to the vSAN 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** text box.

- 5 Under **On-disk Format Version**, click **Upgrade**.
- 6 Click **Yes** on the Upgrade dialog box to perform the upgrade of the on-disk format.

vSAN 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** is 0 (zero).

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 vSAN Disk Format Using RVC

After you have finished upgrading the vSAN 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 vSAN cluster is 6.5 or later.
- Verify that the disks are in a healthy state from the Disk Management page. 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 website 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 complete 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 vSAN cluster enters maintenance mode, the available resource capacity in the cluster is reduced because the member host no longer contributes capacity to the cluster. The cluster upgrade might fail.
- Verify that there are no component rebuilding tasks currently in progress in the vSAN cluster by running the RVC `vsan.resync_dashboard` command.

### Procedure

- 1 Log in to your vCenter Server using RVC.
- 2 Run the following RVC command to view the disk status: `vsan.disks_stats /< vCenter IP address or hostname>/<data center name>/computers/<cluster name>`

For example: `vsan.disks_stats /192.168.0.1/BetaDC/computers/VSANCluster`

The command lists the names of all devices and hosts in the vSAN 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 following RVC command: `vsan.ondisk_upgrade <path to vsan cluster>`

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 following RVC command to verify that the object versions are upgraded to the new on-disk format: `vsan.obj_status_report`

## Verify the vSAN Disk Format Upgrade

After you finish upgrading the disk format, you must verify whether the vSAN cluster is using the new on-disk format.

### Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.

The current disk format version appears at the top of the page.

## Verify the vSAN Cluster Upgrade

The vSAN cluster upgrade is not complete until you have verified that you are using the latest version of vSphere and vSAN is available for use.

### Procedure

- 1 Navigate to the vSAN cluster.
- 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 vSAN 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 8-3. Upgrade Command Options**

Options	Description
<code>--hosts_and_clusters</code>	Use to specify paths to all host systems in the cluster or cluster's compute resources.
<code>--ignore-objects, -i</code>	Use to skip vSAN 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.
<code>--allow-reduced-redundancy, -a</code>	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. vSAN restores full compliance and redundancy after the upgrade is completed.
<code>--force, -f</code>	Use to enable force-proceed and automatically answer all confirmation questions.
<code>--help, -h</code>	Use to display the help options.

For information about using the RVC commands, see the *RVC Command Reference Guide*.

## vSAN Build Recommendations for vSphere Update Manager

vSAN generates system baselines and baseline groups for use with vSphere Update Manager. You can use these recommended baselines to update software, patches, and extensions for hosts in your vSAN cluster.

vSAN 6.6.1 and later generates automated build recommendations for vSAN clusters. vSAN combines information in the VMware Compatibility Guide and vSAN Release Catalog with information about the installed ESXi releases. These recommended updates provide the best available release to keep your hardware in a supported state.

System baselines for vSAN 6.7.1 and later also can include device driver and firmware updates. These updates support the ESXi software recommended for your cluster.

### vSAN System Baselines

vSAN build recommendations are provided through vSAN system baselines for Update Manager. These system baselines are managed by vSAN. They are read-only and cannot be customized.

vSAN generates one baseline group for each vSAN cluster. vSAN system baselines are listed in the **Baselines** pane of the Baselines and Groups tab. You can continue to create and remediate your own baselines.

vSAN system baselines can include custom ISO images provided by certified vendors. If hosts in your vSAN cluster have OEM-specific custom ISOs, then vSAN recommended system baselines can include custom ISOs from same vendor. Update Manager cannot generate a recommendation for custom ISOs not supported by vSAN. If you are running a customized software image that overrides the vendor name in the host's image profile, Update Manager cannot recommend a system baseline.

Update Manager automatically scans each vSAN cluster to check compliance against the baseline group. To upgrade your cluster, you must manually remediate the system baseline through Update Manager. You can remediate vSAN system baseline on a single host or on the entire cluster.

## vSAN Release Catalog

The vSAN release catalog maintains information about available releases, preference order for releases, and critical patches needed for each release. The vSAN release catalog is hosted on the VMware Cloud.

vSAN requires Internet connectivity to access the release catalog. You do not need to be enrolled in the Customer Experience Improvement Program (CEIP) for vSAN to access the release catalog.

If you do not have an Internet connection, you can upload the vSAN release catalog directly to the vCenter Server. In the vSphere Client, click **Configure > vSAN > Update**, and click **Upload from file** in the Release Catalog section. You can download the latest vSAN [release catalog](#).

Update Manager enables you to import storage controller firmware and drivers recommended for your vSAN cluster. Some storage controller vendors provide a software management tool that vSAN can use to update controller drivers and firmware. If the management tool is not present on ESXi hosts, you can download the tool.

## Working with vSAN Build Recommendations

Update Manager checks the installed ESXi releases against information in the Hardware Compatibility List (HCL) in the VMware Compatibility Guide. It determines the correct upgrade path for each vSAN cluster, based on the current vSAN Release Catalog. vSAN also includes the necessary drivers and patch updates for the recommended release in its system baseline.

vSAN build recommendations ensure sure that each vSAN cluster remains at the current hardware compatibility status or better. If hardware in the vSAN cluster is not included on the HCL, vSAN recommends an upgrade to the latest release, since it is no worse than the current state.

---

**Note** Update Manager uses the vSAN health service when performing remediation precheck for hosts in a vSAN cluster. vSAN health service is not available on hosts running ESXi 6.0 Update 1 or earlier. When Update Manager upgrades hosts running ESXi 6.0 Update 1 or earlier, the upgrade of the last host in the vSAN cluster might fail. If remediation failed because of vSAN health issues, you can still complete the upgrade. Use the vSAN health service to resolve health issues on the host, then take that host out of maintenance mode to complete the upgrade workflow.

---

The following examples describe the logic behind vSAN build recommendations.

- Example 1** A vSAN cluster is running 6.0 Update 2, and its hardware is included on the 6.0 Update 2 HCL. The HCL lists the hardware as supported up to release 6.0 Update 3, but not supported for 6.5 and later. vSAN recommends an upgrade to 6.0 Update 3, including the necessary critical patches for the release.
- Example 2** A vSAN cluster is running 6.0 Update 2, and its hardware is included on the 6.0 Update 2 HCL. The hardware is also supported on the HCL for release 6.5 Update 1. vSAN recommends an upgrade to release 6.5 Update 1.
- Example 3** A vSAN cluster is running 6.0 Update 2 and its hardware is not on the HCL for that release. vSAN recommends an upgrade to 6.5 Update 1, even though the hardware is not on the HCL for 6.5 Update 1. vSAN recommends the upgrade because the new state is no worse than the current state.

The recommendation engine runs periodically (once each day), or when the following events occur.

- Cluster membership changes. For example, when you add or remove a host.
- The vSAN management service restarts.
- A user logs in to My VMware (my.vmware.com) through a vSphere Client or RVC.
- An update is made to the VMware Compatibility Guide or the vSAN Release Catalog.

The vSAN Build Recommendation health check displays the current build that is recommended for the vSAN cluster. It also can warn you about any issues with the feature.

## System Requirements

Update Manager must be installed manually on Windows vCenter Server.

vSAN requires Internet access to update release metadata, to check the VMware Compatibility Guide, and to download ISO images from My VMware.

vSAN requires valid My VMware (my.vmware.com) credentials to download ISO images for upgrades. For hosts running 6.0 Update 1 and earlier, you must use RVC to enter the My VMware credentials. For hosts running later software, you can log in from the ESX Build Recommendation health check.

To enter My VMware credentials from RVC, run the following command: `vsan.login_iso_depot -u <username> -p <password>`