

Administering VMware vSAN

Update 3

VMware vSphere 7.0

VMware vSAN 7.0

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

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

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Copyright © 2015-2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

About Administering VMware vSAN 7

1 Introduction to vSAN 8

2 Configuring and Managing a vSAN Cluster 9

Configure a Cluster for vSAN Using the vSphere Client 9

Enable vSAN on an Existing Cluster 11

Disable vSAN 12

Edit vSAN Settings 13

View vSAN Datastore 14

Upload Files or Folders to vSAN Datastores 16

Download Files or Folders from vSAN Datastores 16

3 Using vSAN Policies 18

About vSAN Policies 18

View vSAN Storage Providers 22

About the vSAN Default Storage Policy 23

Change the Default Storage Policy for vSAN Datastores 24

Define a Storage Policy for vSAN Using vSphere Client 25

4 Expanding and Managing a vSAN Cluster 28

Expanding a vSAN Cluster 28

Expanding vSAN Cluster Capacity and Performance 29

Use Quickstart to Add Hosts to a vSAN Cluster 29

Add a Host to the vSAN Cluster 30

Configuring Hosts Using Host Profile 31

Sharing Remote Datastores with HCI Mesh 33

View Remote Datastores 35

Mount Remote Datastore 36

Unmount Remote Datastore 36

Monitor HCI Mesh 36

Working with Maintenance Mode 38

Check a Host's Data Migration Capabilities 39

Place a Member of vSAN Cluster in Maintenance Mode 40

Managing Fault Domains in vSAN Clusters 42

Create a New Fault Domain in vSAN Cluster 43

Move Host into Selected Fault Domain 44

Move Hosts out of a Fault Domain 44

- Rename a Fault Domain 44
- Remove Selected Fault Domains 45
- Tolerate Additional Failures with Fault Domain 45
- Using the vSAN iSCSI Target Service 46
 - Enable the iSCSI Target Service 47
 - Create an iSCSI Target 47
 - Add a LUN to an iSCSI Target 48
 - Resize a LUN on an iSCSI Target 49
 - Create an iSCSI Initiator Group 49
 - Assign a Target to an iSCSI Initiator Group 50
 - Disable the iSCSI Target Service 50
 - Monitor vSAN iSCSI Target Service 51
- vSAN File Service 51
 - Limitations and Considerations 53
 - Configure File Services 53
 - Edit vSAN File Service 59
 - Create a File Share 60
 - View File Shares 62
 - Access File Shares 62
 - Edit a File Share 64
 - Manage SMB File Share 65
 - Delete a File Share 65
 - vSAN Distributed File System Snapshot 66
 - Rebalance Workload on vSAN File Service Hosts 67
 - Reclaiming Space with Unmap 68
 - Upgrade File Service 68
 - Monitor Performance 69
 - Monitor Capacity 70
 - Monitor Health 70
- Migrate a Hybrid vSAN Cluster to an All-Flash Cluster 71
- Shutting Down and Restarting the vSAN Cluster 71
 - Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard 72
 - Restart the vSAN Cluster 73
 - Manually Shut Down and Restart the vSAN Cluster 73

5 Device Management in a vSAN Cluster 77

- Managing Disk Groups and Devices 77
 - Create a Disk Group on a vSAN Host 78
 - Claim Storage Devices for a vSAN Cluster 79
 - Claim Disks for vSAN Direct 80
- Working with Individual Devices 80

- Add Devices to the Disk Group 81
- Check a Disk or Disk Group's Data Migration Capabilities 82
- Remove Disk Groups or Devices from vSAN 82
- Recreate a Disk Group 83
- Using Locator LEDs 84
- Mark Devices as Flash 85
- Mark Devices as HDD 85
- Mark Devices as Local 86
- Mark Devices as Remote 86
- Add a Capacity Device 87
- Remove Partition From Devices 87

6 Increasing Space Efficiency in a vSAN Cluster 89

- Introduction to vSAN Space Efficiency 89
- Reclaiming Space with SCSI Unmap 89
- Using Deduplication and Compression 90
 - Deduplication and Compression Design Considerations 92
 - Enable Deduplication and Compression on a New vSAN Cluster 92
 - Enable Deduplication and Compression on Existing vSAN Cluster 93
 - Disable Deduplication and Compression 93
 - Reducing VM Redundancy for vSAN Cluster 94
 - Adding or Removing Disks with Deduplication and Compression Enabled 95
- Using RAID 5 or RAID 6 Erasure Coding 95
- RAID 5 or RAID 6 Design Considerations 96

7 Using Encryption in a vSAN Cluster 97

- vSAN Data-In-Transit Encryption 97
 - Enable Data-In-Transit Encryption on a vSAN Cluster 98
- vSAN Data-At-Rest Encryption 98
 - How Data-At-Rest Encryption Works 99
 - Design Considerations for Data-At-Rest Encryption 100
 - Set Up the Standard Key Provider 100
 - Enable Data-At-Rest Encryption on a New vSAN Cluster 106
 - Generate New Data-At-Rest Encryption Keys 107
 - Enable Data-At-Rest Encryption on Existing vSAN Cluster 108
- vSAN Encryption and Core Dumps 109

8 Upgrading the vSAN Cluster 113

- Before You Upgrade vSAN 114
- Upgrade the vCenter Server 116
- Upgrade the ESXi Hosts 116

About the vSAN Disk Format	117
Upgrading vSAN Disk Format Using vSphere Client	119
Upgrade vSAN Disk Format Using RVC	120
Verify the vSAN Disk Format Upgrade	122
About vSAN Object Format	122
Verify the vSAN Cluster Upgrade	123
Using the RVC Upgrade Command Options	123
vSAN Build Recommendations for vSphere Lifecycle Manager	124

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.

At VMware, we value inclusion. To foster this principle within our customer, partner, and internal community, we create content using inclusive language.

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*.

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:

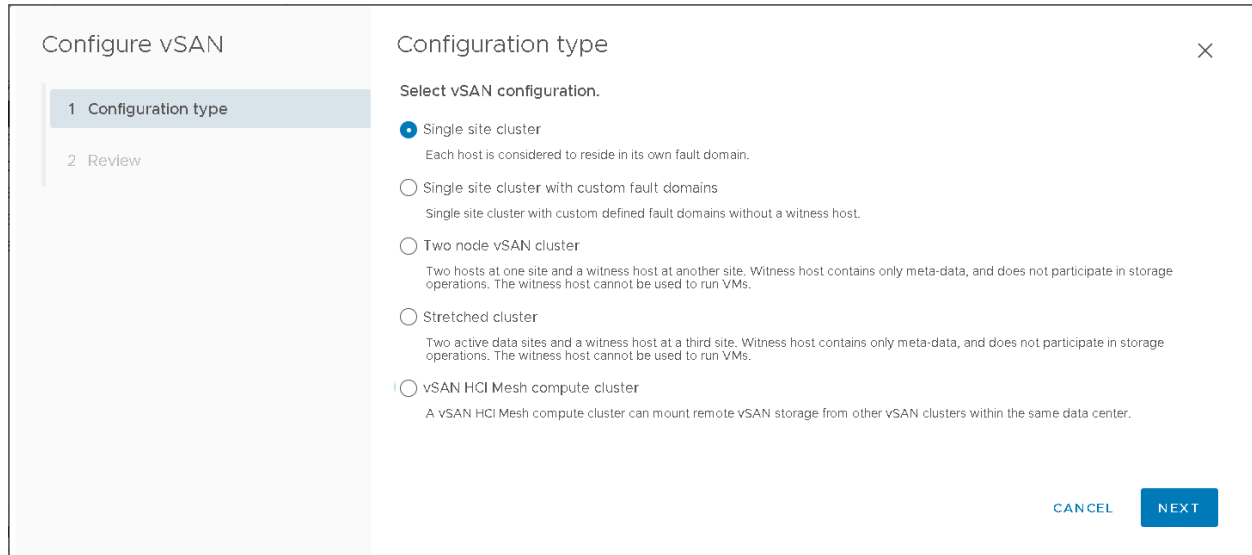
- [Configure a Cluster for vSAN Using the vSphere Client](#)
- [Enable vSAN on an Existing Cluster](#)
- [Disable vSAN](#)
- [Edit vSAN Settings](#)
- [View vSAN Datastore](#)
- [Upload Files or Folders to vSAN Datastores](#)
- [Download Files or Folders from vSAN Datastores](#)

Configure a Cluster for vSAN Using the vSphere Client

You can use the HTML5-based vSphere Client to configure your vSAN cluster.

Note You can use Quickstart to quickly create and configure a vSAN cluster. For more information, see "Using Quickstart to Configure and Expand a vSAN Cluster" in *vSAN Planning and Deployment* .

Note vSAN HCI Mesh compute clusters have limited configuration options.



Prerequisites

Verify that your environment meets all requirements. See "Requirements for Enabling vSAN" in *vSAN Planning and Deployment*.

Create a cluster and add hosts to the cluster before enabling and configuring vSAN.

Procedure

- 1 Navigate to an existing host cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.
- 4 Click **Configure vSAN** to open the Configure vSAN wizard.
- 5 Select the type of vSAN cluster to configure, and click **Next**.
 - Single site cluster. For more information, see "vSAN Deployment Options" in *vSAN Planning and Deployment*.
 - Single site cluster with custom fault domains.
 - Two node vSAN cluster.
 - Stretched cluster.
 - vSAN HCI Mesh compute cluster. For more information, see "Sharing Remote Datastores with HCI Mesh" in *Administering VMware vSAN*.
- 6 Configure the vSAN services to use, and click **Next**.

Configure data management features, including deduplication and compression, data-at-rest encryption, and data-in-transit encryption. For more details, see [Edit vSAN Settings](#).

- 7 Claim disks for the vSAN cluster, and click **Next**.

Each host requires at least one flash device in the cache tier, and one or more devices in the capacity tier. For more details, see "Managing Disk Groups and Devices" in *Administering VMware vSAN*.

- 8 Review the configuration, and click **Finish**.

Results

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

Claim disks or create disk groups. See "Managing Disk Groups and Devices" in *Administering VMware vSAN*.

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

Verify that the vSAN storage provider is registered.

Enable vSAN on an Existing Cluster

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

Prerequisites

Verify that your environment meets all requirements. See "Requirements for Enabling vSAN" in *vSAN Planning and Deployment*.

Note vSAN HCI Mesh compute clusters have limited configuration options.

Procedure

- 1 Navigate to an existing host cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.
- 4 Click **Configure vSAN**.
- 5 Select the type of vSAN cluster to configure, and click **Next**.
 - Single site cluster.
 - Single site cluster with custom fault domains.
 - Two node vSAN cluster.
 - Stretched cluster.
 - vSAN HCI Mesh compute cluster. For more information, see "Sharing Remote Datastores with HCI Mesh" in *Administering VMware vSAN*.

6 Configure the vSAN services to use, and click **Next**.

- Configure vSAN performance service. For more information, see "Monitoring vSAN Performance" in *vSAN Monitoring and Troubleshooting*.
- Enable File Service. For more information, see "vSAN File Service" in *Administering VMware vSAN*.
- Configure vSAN Network options. For more information, see "Designing the vSAN Network" in *vSAN Planning and Deployment*.
- Configure vSAN historical health service.
- Configure iSCSI target service. For more information, see "Using the vSAN iSCSI Target Service" in *Administering VMware vSAN*.
- Configure Data Management options, including deduplication and compression, data-at-rest encryption, and data-in-transit encryption.
- Configure capacity reservations and alerts. For more information, see "About Reserved Capacity" in *vSAN Monitoring and Troubleshooting*.
- Configure advanced options:
 - Object Repair Timer
 - Site Read Locality for stretched clusters
 - Thin Swap provisioning
 - Large Cluster Support for up to 64 hosts
 - Automatic Rebalance

7 Claim disks for the vSAN cluster, and click **Next**.

Each host requires at least one flash device in the cache tier, and one or more devices in the capacity tier. For more information, see "Managing Disk Groups and Devices" in *Administering VMware vSAN*.

8 Review the configuration, and click **Finish**.

Disable vSAN

You can turn off vSAN for a host cluster.

When you disable the vSAN cluster, all virtual machines and data services located on the vSAN datastore become inaccessible. If you have consumed storage on the vSAN cluster using the vSAN-Direct, then the vSAN Direct monitoring services, such as health checks, space reporting, and performance monitoring, are also disabled. 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.
- 3 Under vSAN, select **Services**.
- 4 Click **Turn Off vSAN**.
- 5 On the Turn Off vSAN dialog, confirm your selection.

Edit vSAN Settings

You can edit the settings of your vSAN cluster to configure data management features and enable services provided by the cluster.

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.

The screenshot shows the vSAN Services configuration page. The left sidebar contains a navigation menu with categories: Services, Configuration, Licensing, and vSAN. The 'vSAN Services' section is expanded, showing a list of services with their status and an 'EDIT' button. The 'Data Management' section is also expanded, showing a table of settings.

Service	Status	Action
Performance Service	Enabled	EDIT
File Service	Disabled	ENABLE
vSAN iSCSI Target Service	Disabled	ENABLE

Setting	Value	Action
Space efficiency	None	
Data-at-rest encryption	Disabled	
Key provider	--	
Disk wiping	Disabled	
Data-in-transit encryption	Disabled	
Rekey interval	--	

At the bottom of the Data Management section, there are buttons for 'EDIT' and 'GENERATE NEW ENCRYPTION KEYS'.


Procedure

- 1 Navigate to the vSAN host cluster.

- 2 Click the **Configure** tab.
 - a Under vSAN, select **Services**.
 - b Click the **Edit** or **Enable** button for the service you want to configure.
 - Configure vSAN performance service. For more information, see *Monitoring vSAN Performance* in *vSAN Monitoring and Troubleshooting*.
 - Enable File Service. For more information, see "vSAN File Service" in *Administering VMware vSAN*.
 - Configure vSAN Network options. For more information, see *Configuring vSAN Network* in *vSAN Planning and Deployment*.
 - Configure vSAN historical health service.
 - Configure iSCSI target service. For more information, see "vSAN iSCSI Target Service" in *Administering VMware vSAN*.
 - Configure Data Management options, including deduplication and compression, data-at-rest encryption, and data-in-transit encryption.
 - Configure capacity reservations and alerts. For more information, see "About Reserved Capacity" in *vSAN Monitoring and Troubleshooting*.
 - Configure advanced options:
 - Object Repair Timer
 - Site Read Locality for stretched clusters
 - Thin Swap provisioning
 - Large Cluster Support for up to 64 hosts
 - Automatic Rebalance
 - c Modify the settings to match your requirements.
- 3 Click **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.

Upload Files or Folders to vSAN Datastores

You can upload vmdk files to a vSAN datastore. You can also upload folders to a vSAN datastore. For more information about datastores, see *vSphere Storage*.

When you upload a vmdk file to a vSAN datastore, the following considerations apply:

- You can upload only stream-optimized vmdk files to a vSAN datastore. VMware stream-optimized file format is a monolithic sparse format compressed for streaming. If you want to upload a vmdk file that is not in stream-optimized format, then, before uploading, convert it to stream-optimized format using the `vmware-vdiskmanager` command-line utility. For more information, see *Virtual Disk Manager User's Guide*.
- When you upload a vmdk file to a vSAN datastore, the vmdk file inherits the default policy of that datastore. The vmdk does not inherit the policy of the VM from which it was downloaded. vSAN creates the objects by applying the `vsanDatastore` default policy, which is RAID -1. You can change the default policy of the datastore. See [Change the Default Storage Policy for vSAN Datastores](#).
- You must upload a vmdk file to VM home folder.

Procedure

- 1 Navigate to vSAN Datastore.
- 2 Click the **Files** tab.

Option	Description
Upload Files	<ol style="list-style-type: none"> a Select the target folder and click Upload Files. You see a message informing that you can upload vmdk files only in VMware stream-optimized format. If you try uploading a vmdk file in a different format, you see an internal server error message. b Click Upload. c Locate the item to upload on the local computer and click Open.
Upload Folders	<ol style="list-style-type: none"> a Select the target folder and click Upload Folder. You see a message informing that you can upload vmdk files only in VMware stream-optimized format. b Click Upload. c Locate the item to upload on the local computer and click Open.

Download Files or Folders from vSAN Datastores

You can download files and folders from a vSAN datastore. For more information about datastores, see *vSphere Storage*.

The vmdk files are downloaded as stream-optimized files with the filename `<vmdkName>_stream.vmdk`. VMware stream-optimized file format is a monolithic sparse format compressed for streaming.

You can convert a VMware stream-optimized vmdk file to other vmdk file formats using the `vmware-vdiskmanager` command-line utility. For more information, see *Virtual Disk Manager User's Guide*.

Procedure

1 Navigate to vSAN Datastore.

2 Click the **Files** tab and then click **Download**.

You see a message alerting you that vmdk files are downloaded from the vSAN datastores in VMware stream-optimized format with the filename extension `.stream.vmdk`.

3 Click **Download**.

4 Locate the item to download and then click **Download**.

Using vSAN Policies

3

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](#)
- [View vSAN Storage Providers](#)
- [About the vSAN Default Storage Policy](#)
- [Change the Default Storage Policy for vSAN Datastores](#)
- [Define a Storage Policy for vSAN Using vSphere 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 **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 **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 **Failures to tolerate**.

Table 3-1. Storage Policy Rules

Capability	Description
Failures to tolerate (FTT)	<p>Defines the number of host and device failures that a virtual machine object can tolerate. For n failures tolerated, each piece of data written is stored in $n+1$ places, including parity copies if using RAID 5 or RAID 6.</p> <p>If fault domains are configured, $2n+1$ 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>You can select a data replication method that optimizes for performance or capacity. RAID-1 (Mirroring) uses more disk space to place the components of objects but provides better performance for accessing the objects. RAID-5/6 (Erasure Coding) uses less disk space, but performance is reduced.</p> <hr/> <p>Note If you do not want vSAN to protect a single mirror copy of virtual machine objects, you can specify No data redundancy. 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 No data redundancy means that your data is unprotected, and you might lose data when the vSAN cluster encounters a device failure.</p> <hr/> <p>Note If you create a storage policy and you do not specify a value for FTT, 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> <hr/>
Site disaster tolerance	<p>In a stretched cluster, this rule defines the number of additional host failures that an object can tolerate after the number of failures defined by FTT is reached.</p> <p>None - standard cluster is the default value. For a stretched cluster, you can choose to keep data on the Preferred or Secondary site for host affinity.</p> <p>Host mirroring - 2 node cluster defines the number of additional failures that an object can tolerate after the number of failures defined by FTT is reached. vSAN performs object mirroring at the disk group level. Each data host must have at least three disk groups to use this rule.</p> <p>Site mirroring - stretched cluster defines the number of additional host failures that an object can tolerate after the number of failures defined by FTT is reached.</p> <hr/>

Table 3-1. Storage Policy Rules (continued)

Capability	Description
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. 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. For 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>Note 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>
Force provisioning	<p>If the option is set to Yes, the object is provisioned even if the Failures to tolerate, Number of disk stripes per object, and Flash read cache reservation policies specified in the storage policy cannot be satisfied by the datastore. Use this parameter in bootstrapping scenarios and during an outage when standard provisioning is no longer possible.</p> <p>The default No is acceptable for most production environments. vSAN fails to provision a virtual machine when the policy requirements are not met, but it successfully creates the user-defined storage policy.</p>

Table 3-1. Storage Policy Rules (continued)

Capability	Description
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. The following options are available:</p> <ul style="list-style-type: none"> ■ Thin provisioning (default) ■ 25% reservation ■ 50% reservation ■ 75% reservation ■ Thick provisioning
Disable object checksum	<p>If the option is set to No, the object calculates checksum information to ensure the integrity of its data. If this option is set to Yes, the object does not calculate checksum information.</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 No, 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 IOPS limit for object 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*.

How vSAN Manages Policy Changes

vSAN 6.7 Update 3 and later manages policy changes to reduce the amount of transient space consumed across the cluster. Transient capacity is generated when vSAN reconfigures objects for a policy change.

When you modify a policy, the change is accepted but not applied immediately. vSAN batches the policy change requests and performs them asynchronously, to maintain a fixed amount of transient space.

Policy changes are rejected immediately for non-capacity related reasons, such as changing a RAID5 policy to RAID6 on a five-node cluster.

You can view transient capacity usage in the vSAN Capacity monitor. To verify the status of a policy change on an object, use the vSAN health service to check the vSAN object health.

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 6.7 and later releases register only one vSAN Storage Provider for all the vSAN clusters managed by the vCenter Server using the following URL:

```
https://<VC fqdn>:<VC https port>/vsanHealth/vsanvp/version.xml
```

Verify that the storage providers are registered.

Procedure

- 1 Navigate to vCenter Server.
- 2 Click the **Configure** tab, and click **Storage Providers**.

Results

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
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
	Note 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. In some cases, when you scale up a cluster, you must modify the default storage policy to maintain compliance with the requirements of the [Service Level Agreement for VMware Cloud on AWS](#).

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.

Change the Default Storage Policy for vSAN Datastores

You can change the default storage policy for a selected vSAN datastore.

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 'Advanced Policy Rules' section is active, displaying the following settings:


- Number of disk stripes per object: 1
- IOPS limit for object: 0
- Object space reservation: Thin provisioning (Initially reserved storage space for 100 GB VM disk would be 0 B)
- Flash read cache reservation (%): 0 (Reserved cache space for 100GB VM disk would be 0 B)
- 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, and click **Next**.

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

Availability options define the rules for 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:

- 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 Storage Rules tab, define the encryption, space efficiency, and storage tier rules that can be used along with the HCI Mesh to distinguish the remote datastores.

- **Encryption services:** Defines the encryption rules for virtual machines that you deploy with this policy. You can choose one of the following options:
 - **Data-At-Rest encryption:** Encryption is enabled on the virtual machines.
 - **No encryption:** Encryption is not enabled on the virtual machines.
 - **No preference:** Makes the virtual machines compatible with both Data-At-Rest encryption and No encryption options.
- **Space Efficiency:** Defines the space saving rules for the virtual machines that you deploy with this policy. You can choose one of the following options:
 - **Deduplication and compression:** Enables both deduplication and compression on the virtual machines. Deduplication and compression are available only on all-flash disk groups. For more information, see [Deduplication and Compression Design Considerations](#).
 - **Compression only:** Enables only compression on the virtual machines. Compression is available only on all-flash disk groups. For more information, see [Deduplication and Compression Design Considerations](#).
 - **No space efficiency:** Space efficiency features are not enabled on the virtual machines. Choosing this option requires datastores without any space efficiency options to be turned on.
 - **No preference:** Makes the virtual machines compatible with all the options.

- **Storage tier:** Specifies the storage tier for the virtual machines that you deploy with this policy. You can choose one of the following options. Choosing the **No preference** option makes the virtual machines compatible with both hybrid and all flash environments.
 - **All flash**
 - **Hybrid**
 - **No preference**
- c On the Advanced Policy Rules tab, define advanced policy rules, such as number of disk stripes per object and IOPS limits.
- d 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 On the Storage compatibility page, review the list of datastores under the **COMPATIBLE** and **INCOMPATIBLE** tabs 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.
- 8 On the Review and finish page, review the policy settings, and click **Finish**.

Results

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](#)
- [Sharing Remote Datastores with HCI Mesh](#)
- [Working with Maintenance Mode](#)
- [Managing Fault Domains in vSAN Clusters](#)
- [Using the vSAN iSCSI Target Service](#)
- [vSAN File Service](#)
- [Migrate a Hybrid vSAN Cluster to an All-Flash Cluster](#)
- [Shutting Down and Restarting the 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 automatically distributes data to the newly added device. See "Configure Automatic 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. For more information See *vSAN Planning and Deployment*.
- 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 *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, the host cannot be placed into maintenance mode as you add it to a cluster using the Quickstart workflow. 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.
- No network configuration performed through the Quickstart workflow has been modified from outside of the Quickstart workflow.

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 **Launch** 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 **Launch** to open the Cluster configuration wizard.
 - a On the Configure the distributed switches page, enter networking settings for the new hosts.
 - b (optional) On the Claim disks page, select disks on each new host.
 - c (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](#).
 - d On the Ready to complete page, verify the cluster settings, and click **Finish**.

Add a Host to the vSAN Cluster

You can add ESXi hosts to a running vSAN cluster without disrupting any ongoing operations. The new 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 Hosts**. The Add hosts wizard appears.

Option	Description
New hosts	<ol style="list-style-type: none"> a Enter the host name or IP address. b Enter the user name and password associated with the host.
Existing hosts	<ol style="list-style-type: none"> a Select hosts that you previously added to vCenter Server.

- 3 Click **Next**.
- 4 View the summary information and click **Next**.
- 5 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.

For more information about vSAN cluster configuration and fixing problems, see "vSAN Cluster Configuration Issues" 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.

Sharing Remote Datastores with HCI Mesh

vSAN clusters can share their datastores with other vSAN clusters. You can provision VMs running on the local cluster use storage space on the remote datastore.

Use the Datastore Sharing view to monitor and manage remote datastores mounted on the local vSAN cluster. Each client vSAN cluster can mount remote datastores from server vSAN clusters located within the same data center managed by the vCenter Server. Each compatible vSAN cluster also can act as a server, and allow other vSAN clusters to mount its local datastores.

Mounting a remote datastore with HCI Mesh is a cluster-wide configuration. You can mount a remote datastore to a vSAN cluster, which is then mounted to all hosts in the cluster.

When you provision a new virtual machine, you can select a remote datastore that is mounted to the client cluster. Assign any compatible storage policy configured for the datastore.

Monitor views for capacity, performance, health, and placement of virtual objects show the status of remote objects and datastores.

HCI Mesh vSAN has the following design considerations:

- Clusters must be managed by the same vCenter Server and be located within the same data center.
- Clusters must be running 7.0 Update 1 or later.
- A vSAN cluster can serve its local datastore to up to five client vSAN clusters.
- A client cluster can mount up to five remote datastores from one or more vSAN server clusters.
- A single remote datastore can be mounted to up to 128 vSAN hosts, including hosts in the vSAN server cluster.
- All objects that make up a VM must reside on the same datastore.
- For vSphere HA to work with HCI Mesh, configure the following failure response for Datastore with APD: Power off and restart VMs.
- Client hosts that are not part of a cluster are not supported. You can configure a single host compute-only cluster, but vSphere HA does not work unless you add a second host to the cluster.

The following features are not supported with HCI Mesh:

- vSAN data-in-transit encryption
- vSAN stretched clusters
- vSAN 2-node clusters

The following configurations are not supported with HCI Mesh:

- Remote provisioning of vSAN file share, iSCSI volumes, or CNS persistent volumes. You can provision them on the local vSAN datastore, but not on any remote vSAN datastore.
- Air-gapped vSAN networks or clusters using multiple vSAN VMkernel ports
- vSAN communication over RDMA

HCI Mesh Compute-only Client

vSAN 7.0 Update 2 and later enables you to configure a non-vSAN cluster as an HCI Mesh client. The hosts in an HCI Mesh compute-only client cluster do not need local storage. They can mount remote datastores from a vSAN cluster located within the same data center.

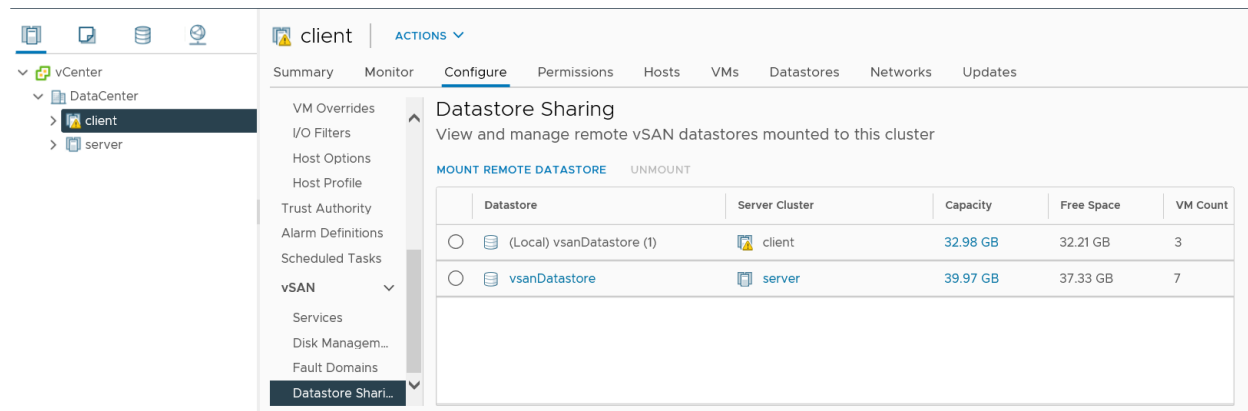
HCI Mesh compute-only clusters have the following design considerations:

- vSAN networking must be configured on the client hosts.
- No disk groups can be present on vSAN compute-only hosts.
- No vSAN data management features can be configured on the compute-only cluster.

When you configure a vSphere cluster for vSAN, you can specify it as an HCI Mesh compute cluster. You can mount a remote datastore, and monitor the capacity, health, and performance of the remote vSAN datastore.

View Remote Datastores

Use the Datastore Sharing page to view remote datastores mounted to the local vSAN cluster, and client clusters sharing the local datastore.



	Datastore	Server Cluster	Capacity	Free Space	VM Count
<input type="radio"/>	(Local) vsanDatastore (1)	client	32.98 GB	32.21 GB	3
<input type="radio"/>	vsanDatastore	server	39.97 GB	37.33 GB	7

Procedure

- 1 Navigate to the local vSAN cluster.
- 2 Click the Configure tab.
- 3 Under vSAN, click **Datastore Sharing**.

Results

This view lists information about each datastore mounted to the local cluster.

- Server cluster that hosts the datastore
- Capacity of the datastore
- Free space available
- Number of VMs using the datastore (number of VMs using the compute resources of the local cluster, but the storage resources of the server cluster)
- Client clusters that have mounted the datastore

What to do next

You can mount or unmount remote datastores from this page.

Mount Remote Datastore

You can mount one or more datastores from other vSAN clusters managed by the same vCenter Server.

Procedure

- 1 Navigate to the local vSAN cluster.
- 2 Click the Configure tab.
- 3 Under vSAN, click **Datastore Sharing**.
- 4 Click **Mount Remote Datastore**.
- 5 Select a datastore and click **Next**.
- 6 Check the datastore compatibility, and click **Finish**.

Results

The remote datastore is mounted to the local vSAN cluster.

What to do next

When you provision a VM, you can select the remote datastore as the storage resource. Assign a storage policy that is supported by the remote datastore.

Unmount Remote Datastore

You can unmount a remote datastore from a vSAN cluster.

If no virtual machines on the local cluster are using the remote vSAN datastore, you can unmount the datastore from your local vSAN cluster.

Procedure

- 1 Navigate to the local vSAN cluster.
- 2 Click the Configure tab.
- 3 Under vSAN, click **Datastore Sharing**.
- 4 Select a remote datastore, and click **Unmount**.
- 5 Click **Unmount** to confirm.

Results

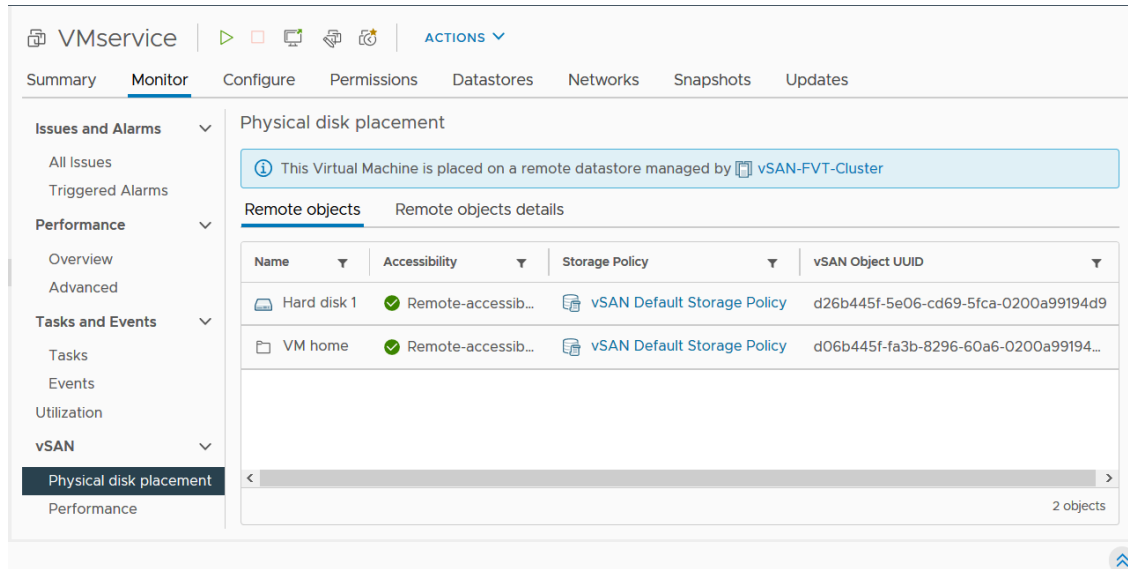
The selected datastore is unmounted from the local cluster.

Monitor HCI Mesh

You can use the vSphere Client to monitor the status of HCI Mesh operations.

vSAN capacity monitor notifies you when remote datastores are mounted to the cluster. You can select the remote datastore to view its capacity information.

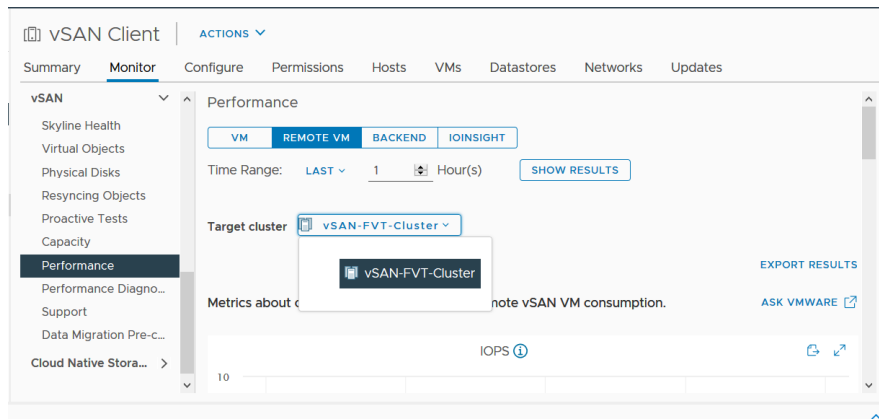
The Virtual Objects view shows the datastore where virtual objects reside. The Physical disk placement view for a VM located on a remote datastore shows information about its remote location.



vSAN health checks report on the status of HCI functions.

- Data > vSAN Object health check shows accessibility information of remote objects.
- Network > Server cluster partition check reports about network partitions between hosts in the client cluster and the server cluster.
- Network > Latency checks the latency between hosts in the client cluster and the server cluster.

vSAN cluster performance views include VM performance charts that display the VM level performance of the client cluster from the perspective of the remote cluster. You can select a remote datastore to view the performance.



You can run pro-active tests on remote datastores to verify VM creation and network performance. The VM creation test creates a VM on the remote datastore. The Network performance test checks the network performance between all hosts in the client cluster and all hosts the server clusters.

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.

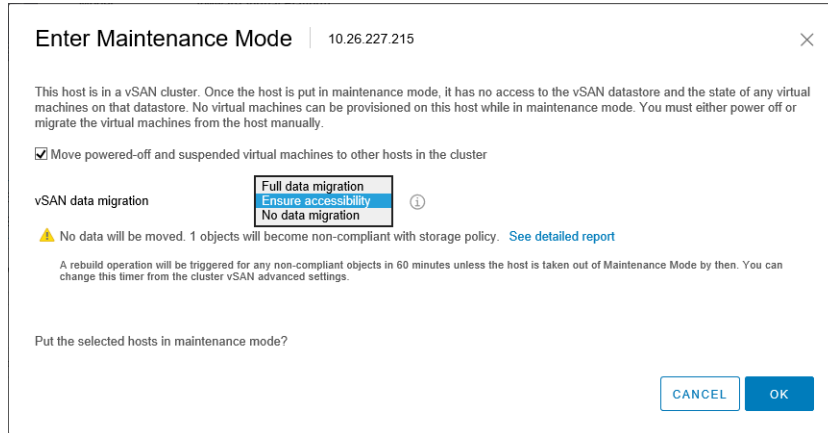
When working with maintenance mode, consider the following guidelines:

- When you place an ESXi host in maintenance mode, you must select a data evacuation mode, such as **Ensure accessibility** or **Full data migration**.
- 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 accessibility** mode is faster than the **Full data migration** mode because the **Ensure 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 accessibility** mode does not reprotect your data during failure and you might experience unexpected data loss.
- When you select the **Full data migration** mode, your data is automatically reprotected against a failure, if the resources are available and the **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 **Full data migration** 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 **Full data migration**. 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 **Full data migration** mode, verify that the cluster has enough hosts and capacity available to meet the **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.



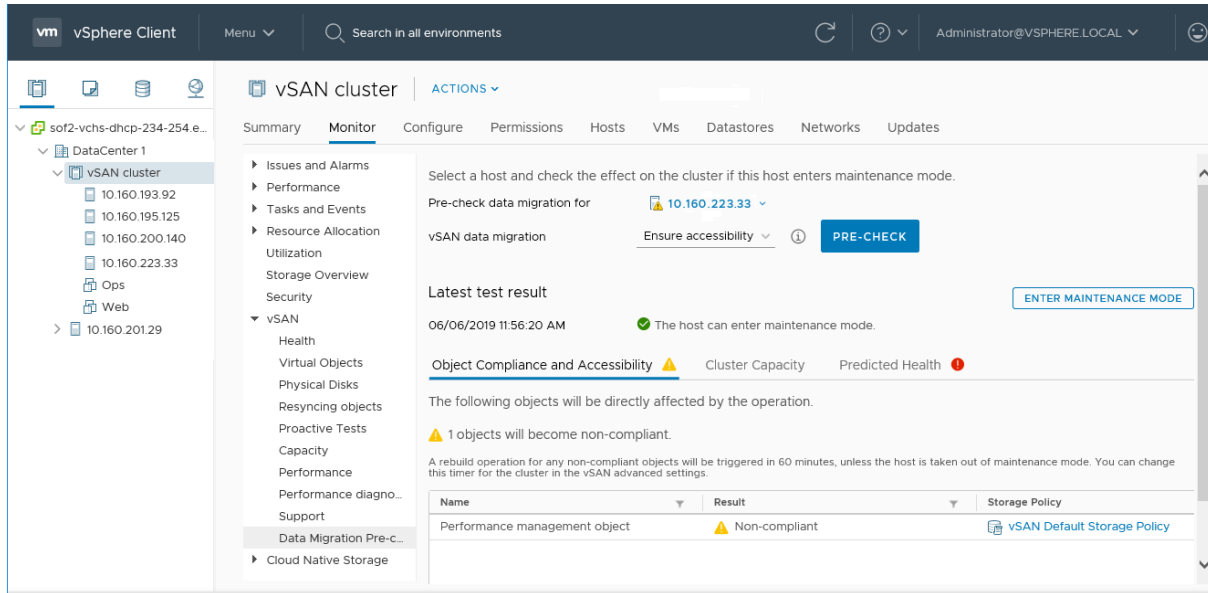
The Confirm Maintenance Mode dialog box provides information to guide your maintenance activities. You can view the impact of each data evacuation option.

- 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.

Check a Host's Data Migration Capabilities

Use the data migration pre-check to determine the impact of data migration options when placing a host into maintenance mode or removing it from the cluster.

Before you place a vSAN host into maintenance mode, run the data migration pre-check. The test results provide information to help you determine the impact to cluster capacity, predicted health checks, and any objects that will go out of compliance. If the operation will not succeed, pre-check provides information about what resources are needed.



Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the Monitor tab.
- 3 Under vSAN, click **Data Migration Pre-check**.
- 4 Select a host, a data migration option, and click **Pre-check**.

vSAN runs the data migration precheck tests.

- 5 View the test results.

The pre-check results show whether the host can safely enter maintenance mode.

- The Object Compliance and Accessibility tab displays objects that might have issues after the data migration.
- The Cluster Capacity tab displays the impact of data migration on the vSAN cluster before and after you perform the operation.
- The Predicted Health tab displays the health checks that might be affected by the data migration.

What to do next

If the pre-check indicates that you can place the host into maintenance mode, you can click **Enter Maintenance Mode** to migrate the data and place the host into maintenance mode.

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
Ensure accessibility	<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 Failures to tolerate is set to 1, you might experience data loss in the cluster.</p> <hr/> <p>Note 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>
Full data migration	<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> <hr/> <p>Note 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>
No data migration	<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. For more information see *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 **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 **Failures to tolerate** set to 1 ($FTT=1$), 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 **Failures to tolerate** set to N ($FTT=n$), 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 $FTT=1$. 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 **Failures to tolerate** set to 1. If you want the **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 plus icon. The New Fault Domain wizard opens.
- 5 Enter the fault domain name.
- 6 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.

7 Click **Create**.

The selected hosts appear in the fault domain. Each fault domain displays the used and reserved capacity information. This enables you to view the capacity distribution across 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 Click and drag the host that you want to add onto an existing fault domain.

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**.
 - a Click and drag the host from the fault domain to the Standalone Hosts area.
 - b Click **Move** to confirm.

Results

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**.
 - a Click the Actions icon on the right side of the fault domain, and choose **Edit**.
 - b Enter a new fault domain name.
- 4 Click **Apply** or **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**.
- 4 Click the Actions icon on the right side of the fault domain, and select **Delete**.
- 5 Click **Delete** to confirm.

Results

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.

Tolerate Additional Failures with Fault Domain

Fault domains in a vSAN cluster provides resilience and assures that the data is available even with failures based on policy. With failures to tolerate (FTT) set to 1, the object can tolerate a failure. However, a temporary failure followed by a permanent failure in a cluster can result in data loss.

An additional fault domain provides vSAN the ability to create a durability component without having additional FTTs for the object. vSAN triggers this extra component during planned and unplanned failures. Unplanned failures include network disconnect, disk failures, and host failures. Planned failures include Entering Maintenance Mode (EMM). For example, a 6 host cluster with RAID 6 object cannot create a durability component if there is a host failure.

vSAN ensures the data availability of the objects when the components go offline and comes back online unexpectedly based on the FTTs specified in the storage policy. During a failure, the writes of the failed component is redirected to the durability component. When the component recovers from the transient failure and comes back online, the durability component disappears and results in the resynchronization of the component.

Without the durability component in place, if there is a second permanent failure in the cluster and the mirror object is affected, the object data gets permanently lost even if the failure is resolved.

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:

CHAP

In CHAP authentication, the target authenticates the initiator, but the initiator does not authenticate the target.

Mutual CHAP

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 and click **Configure > vSAN > Services**.

2 On the vSAN iSCSI Target Service row, click **ENABLE**.

The Edit vSAN iSCSI Target Service wizard opens.

3 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.

4 Click the **Enable vSAN iSCSI Target service** slider to turn it on and then click **APPLY**.

Results

The vSAN iSCSI target service is enabled.

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.
 - a Under vSAN, click **iSCSI Target Service**.
 - b Click the iSCSI Targets tab.
 - c Click **Add**. The **New iSCSI Target** dialog box is displayed. If you leave the target IQN field blank, the IQN is generated automatically.
 - d Enter a target **Alias**.
 - e Select a **Storage policy**, **Network**, **TCP port**, and **Authentication** method.
 - f Select the **I/O Owner Location**. This feature is available only if you have configured vSAN cluster as a stretched cluster. It allows you to specify the site location for hosting the iSCSI target service for a target. This helps in avoiding the cross site iSCSI traffic. If you have set the policy as HFT>=1, then in the event of a site failure, the I/O owner location changes to the alternate site. After the site failure recovery, the I/O owner location automatically changes back to the original I/O owner location as per the configuration. You can select one of the following options to set the site location:
 - **Either**: Hosts the iSCSI target service either on Preferred or Secondary site.
 - **Preferred**: Hosts the iSCSI target service on the Preferred site.
 - **Secondary**: Hosts the iSCSI target service on the Secondary site.
- 3 Click **OK**.

Results

iSCSI target is created and listed under the vSAN iSCSI Targets section with the information such as IQN, I/O owner host, and so on.

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.
 - a Under vSAN, click **iSCSI Target Service**.
 - b Click the iSCSI Targets tab, and select a target.

- c In the vSAN iSCSI LUNs section, click **Add**. The **Add LUN to Target** dialog box is displayed.
- 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.

3 Click **Add**.

Resize a LUN on an iSCSI Target

Depending on your requirement, you can increase the size of an online LUN. Online resizing of the LUN is enabled only if all hosts in the cluster are upgraded to vSAN 6.7 Update 3 or later.

Procedure

- 1 In the vSphere Client, navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **iSCSI Target Service**.
- 4 Click the **iSCSI Targets** tab and select a target.
- 5 In the vSAN iSCSI LUNs section, select a LUN and click **Edit**. The Edit LUN dialog box is displayed.
- 6 Increase the size of the LUN depending on your requirement.
- 7 Click **OK**.

Create an iSCSI Initiator Group

You can create an iSCSI initiator group to provide access control for iSCSI targets. Only iSCSI initiators that are members of the initiator group can access the iSCSI targets.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
 - a Under vSAN, click **iSCSI Target Service**.
 - b Click the Initiator Groups tab, and click the **Add a new iSCSI initiator group (+)** icon. The **New Initiator Group** dialog box is displayed.

- c Enter a name for the iSCSI initiator group.
- d (Optional) To add members to the initiator group, enter the IQN of each member. Use the following format to enter the member IQN:

iqn.YYYY-MM.domain:name

Where:

- YYYY = year, such as 2016
- MM = month, such as 09
- domain = domain where the initiator resides
- name = member name (optional)

- 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.
 - a Under vSAN, click **iSCSI Target Service**.
 - b Select the **Initiator Groups** tab.
 - c In the Accessible Targets section, click the **Add a new accessible target for iSCSI Initiator group (+)** icon. The **Add Accessible Targets** dialog box is displayed.
 - d Select a target from the list of available targets.
- 3** Click **Add**.

Disable the iSCSI Target Service

You can disable the vSAN iSCSI target service. Disabling vSAN iSCSI target service does not delete the LUNs/Targets. If you wish to reclaim the space, delete the LUNs/targets manually before disabling vSAN iSCSI target service.

Prerequisites

Workloads running on iSCSI LUNs are stopped when you disable the iSCSI target service. Before disabling, ensure that there are no workloads running on iSCSI LUNs.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > Services**.
- 2 On the vSAN iSCSI Target Service row, click **EDIT**.
The Edit vSAN iSCSI Target Service wizard opens.
- 3 Click the **Enable vSAN iSCSI Target Service** slider to turn it off and click **Apply**.

Results

The vSAN iSCSI target service is disabled.

What to do next

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.

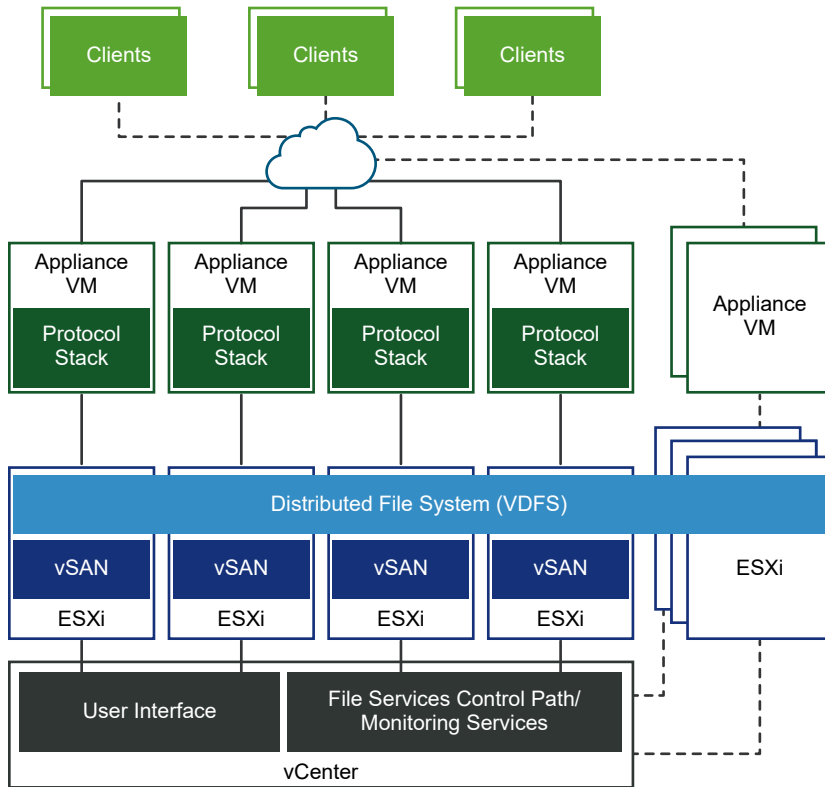
Procedure

- 1 Browse to the vSAN cluster.
- 2 Click **Monitor** and select **Virtual Objects**. iSCSI targets are listed on the page.
- 3 Select a target and click **View Placement Details**. The Physical Placement shows where the data components of the target are located.
- 4 Click **Group components by host placement** to view the hosts associated with the iSCSI data components.

vSAN File Service

Use the vSAN file service to create file shares in the vSAN datastore that client workstations or VMs can access. The data stored in a file share can be accessed from any device that has access rights.

vSAN File Service is a layer that sits on top of vSAN to provide file shares. It currently supports SMB, NFSv3, and NFSv4.1 file shares. vSAN File Service comprises of vSAN Distributed File System (vDFS) which provides the underlying scalable filesystem by aggregating vSAN objects, a Storage Services Platform which provides resilient file server end points and a control plane for deployment, management, and monitoring. File shares are integrated into the existing vSAN Storage Policy Based Management, and on a per-share basis. vSAN file service brings in capability to host the file shares directly on the vSAN cluster.



When you configure vSAN file service, vSAN creates a single VDFS distributed file system for the cluster which will be used internally for management purposes. A file service VM (FSVM) is placed on each host. The FSVMs manage file shares in the vSAN datastore. Each FSVM contains a file server that provides both NFS and SMB service.

A static IP address pool should be provided as an input while enabling file service workflow. One of the IP addresses is designated as the primary IP address. The primary IP address can be used for accessing all the shares in the file services cluster with the help of SMB and NFSv4.1 referrals. A file server is started for every IP address provided in the IP pool. A file share is exported by only one file server. However, the file shares are evenly distributed across all the file servers. To provide computing resources that help manage access requests, the number of IP addresses must be equal to the number of hosts in the vSAN cluster.

vSAN file service supports stretched clusters and two-node clusters. A two-node cluster should have two data node servers in the same location or office, and the witness in a remote or shared location.

Limitations and Considerations

Consider the following when configuring vSAN File Service:

- With vSAN 7.0 U3, File Service VMs are powered off and but no longer deleted when the vSAN cluster enters maintenance mode.
- vSAN 7.0 Update 3 supports two-node configurations and stretched clusters.
- vSAN 7.0 Update 3 supports 64 file servers in a 64 host setup.
- vSAN 7.0 Update 3 supports 100 file shares.
- In releases prior to vSAN 7.0 Update 3, when a host enters maintenance mode, the Protocol Stack container moves to another FSVM. The FSVM on the host that entered maintenance mode is deleted. After the host exits maintenance mode, a new FSVM is provisioned.

File Service VMs are powered off and deleted when the vSAN cluster enters maintenance mode, and recreated when the host exits maintenance mode.

- vSAN File Services VM (FSVM) docker internal network may overlap with the customer network without warning or reconfiguration.

There is known conflict issue if the specified file service network overlaps with the docker internal network (172.17.0.0/16). This causes routing problem for the traffic to the correct endpoint.

As a workaround, specify a different file service network so that it does not overlap with the docker internal network (172.17.0.0/16).

Configure File Services

You can configure the File Services, which enable you to create file shares on your vSAN datastore. You can enable vSAN File Services on a regular vSAN cluster, a vSAN stretched cluster, or a vSAN ROBO cluster.

Prerequisites

Ensure that the following are configured before enabling the vSAN File Services:

- Active Directory (AD) domain if you are planning to create an SMB file share or an NFSv4.1 file share with the Kerberos security.
- A static IP address to use as the single point of access to vSAN file shares. For best performance, the number of IP addresses must be equal to the number of hosts in the vSAN cluster.

Note For the file servers, vSAN File Services support only the IPV4 addresses.

- The static IP addresses should be part of the Forward lookup and Reverse lookup zones in the DNS server.
- All the static IP addresses should be from the same subnet.

- vSAN File Services are supported on DVS version 6.6.0 or later. Create a dedicated port group for vSAN File Services in the DVS.
- To enable vSAN File Services, an additional 4 cores CPU and 10 GB physical memory are required for each host.
- MacLearning and Forged Transmits are enabled as part of the vSAN File Services enablement process for a provided DVS port group.

For standard switches, the Promiscuous Mode and Forged Transmits are enabled as part of the vSAN File Services enablement process.

If NSX-based networks are being used, ensure that MacLearning is enabled for the provided network entity from the NSX admin console, and all the hosts and File Services nodes are connected to the desired NSX-T network.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > Services**.
- 2 On the File Service row, click **Enable**.
The Configure File Service wizard opens.
- 3 Review the checklist on the Introduction page, and click **Next**.

- 4 In the File service agent page, select one of the following options to download the OVF file.

Option	Description
<p>Automatic approach</p>	<p>This option lets the system search and download the OVF.</p> <hr/> <p>Note</p> <ul style="list-style-type: none"> ■ Ensure that you have configured the proxy and firewall so that vCenter can access the following website and download the appropriate JSON file. https://download3.vmware.com/software/VSANOVF/FsOvfMapping.json For more information about configuring the vCenter DNS, IP address, and proxy settings, see <i>vCenter Server Appliance Configuration</i>. ■ If an OVF is already downloaded and available, then following the options are available: <ul style="list-style-type: none"> ■ Use current OVF: Lets you use the OVF that is already available. ■ Automatically load latest OVF: Lets the system search and download the latest OVF.
<p>Manual approach</p>	<p>This option allows you to browse and select an OVF that is already available on your local system.</p> <hr/> <p>Note If you select this option, you should upload all the following files:</p> <ul style="list-style-type: none"> ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x_OVF10.mf ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-x_OVF10.cert ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-x-system.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-cloud-components.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-log.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x_OVF10.ovf

- 5 In the Domain page, enter the following information and click **Next**:
- **File service domain:** The domain name should have minimum two characters. The first character should be an alphabet or a number. The remaining characters can include an alphabet, a number, an underscore (_), a period (.), a hyphen (-).
 - **DNS servers:** Enter a valid DNS server to ensure the proper configuration of File Services.
 - **DNS suffixes:** Provide the DNS suffix that is used with the file services. All other DNS suffixes from where the clients can access these file servers should also be included. File Services does not support DNS domain with single label, such as "app",

"wiz", "com" and so on. A domain name given to file services should be of the format `thisdomain.registerdrootdnsname`. DNS name and suffix must adhere to the best practices detailed in <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/selecting-the-forest-root-domain>.

- **Directory Service:** Configure an Active Directory domain to vSAN File Services for authentication. If you are planning to create an SMB file share or an NFSv4.1 file share with Kerberos authentication, then you must configure an AD domain to vSAN File Services.

Enter appropriate values in the following text boxes to configure the Active Directory domain to vSAN File Services:

Option	Description
AD domain	Fully qualified domain name joined by the file server.
Organizational unit (Optional)	Contains the computer account that the vSAN File Services creates. In an organization with complex hierarchies, create the computer account in a specified container by using a forward slash mark to denote hierarchies (for example, <code>organizational_unit/inner_organizational_unit</code>).
	Note By default, the vSAN File Services create the computer account in the Computers container.

Option	Description
AD username	<p>User name to be used for connecting and configuring the Active Directory service.</p> <p>This user name authenticates the active directory on the domain. A domain user authenticates the domain controller and creates vSAN File Services computer accounts, related SPN entries, and DNS entries (when using Microsoft DNS). As a best practice, create a dedicated service account for the file services.</p> <p>This user should have the following privileges in the organizational unit:</p> <ul style="list-style-type: none"> ■ Create and delete Computer Objects. ■ Read and Write ms-DS-PrincipleName. ■ Read and Write uPNSuffixes. ■ (Optional) Add/Update DNS entries
Password	<p>Password for the user name of the Active Directory on the domain. vSAN File Services use the password to authenticate to AD and to create the vSAN File Services computer account.</p>

Note

- vSAN File Services does not support the following:
 - Read-Only Domain Controllers (RODC) for joining domains because the RODC cannot create machine accounts. As a security best practice, a dedicated org unit should be pre-created in the Active Directory and the user name mentioned here should be controlling this organization.
 - Disjoint namespace.
 - Spaces in organizational units (OUs) names.
 - Multi domain and Single Active Directory Forest environments.
- Only English characters are supported for Active Directory user name.
- Only single AD domain configuration is supported. However, the file servers can be put on a valid DNS subdomain. For example, an AD domain with the name `example.com` can have file server FQDN as `name1.eng.example.com`.
- Pre-created computer objects for file servers are not supported. Make sure that the user provided here have sufficient privilege over the organizational unit.
- vSAN File Services update the DNS records for the file servers if the Active Directory is also used as a DNS server and the user has sufficient permission to update the DNS records. vSAN File Services also has a Health Check to indicate if the forward and reverse lookups for file servers are working properly. However, if there are other proprietary solutions used as DNS servers, the Vi admin should update these DNS records.

6 In the Networking page, enter the following information, and click **Next**:

- Network
- Protocol
- Subnet mask
- Gateway

7 In the IP Pool page, enter the following information, select a **Primary IP**, and then click **Next**.

- **IP address**
- **DNS name**
- **Affinity site**: This option is available if you are configuring vSAN file service on a stretched cluster. This option allows you to configure the placement of the file server on **Preferred** or **Secondary** site. This helps in reducing the cross-site traffic latency. The default value is **Either**, which indicates that no site affinity rule is applied to the file server.

Note If your cluster is a ROBO cluster, ensure that the Affinity site value is set to **Either**.

In a site failure event, the file server affiliated to that site fails over to the other site. The file server fails back to the affiliated site when it is recovered. Configure more file servers to one site if more workloads can be expected from a certain site.

Note If the file server contains SMB file shares, then it does not failback automatically even if the site failure is recovered.

Consider the following while configuring the IP addresses and DNS names:

- To ensure proper configuration of File Services, the IP addresses you enter in the IP Pool page should be static addresses and the DNS server should have records for those IP addresses. For best performance, the number of IP addresses must be equal to the number of hosts in the vSAN cluster.
- You can enter up to 32 IP addresses.
- You can use the following options to automatically fill the IP address and DNS server name text boxes:

AUTO FILL: This option is displayed after you enter the first IP address in the IP address text box. Click the AUTO FIL option to automatically fill the remaining fields with sequential IP addresses, based on the subnet mask and gateway address of the IP address that you have provided in the first row. You can edit the auto filled IP addresses.

LOOK UP DNS: This option is displayed after you enter the first IP address in the IP address text box. Click the LOOK UP DNS option to automatically retrieve the FQDN corresponding to the IP addresses in the IP address column.

Note

- All valid rules apply for the FQDNs. For more information, see <https://tools.ietf.org/html/rfc953>.
 - The first part of the FQDN, also known as NetBIOS Name, should not have more than 15 characters.
-

The FQDNs are automatically retrieved only under the following conditions:

- You should have entered a valid DNS server in the Domain page.
- The IP addresses entered in the IP Pool page should be static addresses and the DNS server should have records for those IP addresses.

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

Results

The OVF is downloaded and deployed. The file services domain is created and the vSAN file services is enabled. File servers are started with the IP addresses that were assigned during the vSAN File Services configuration process.

- The OVF is downloaded and deployed.
- The file services domain is created and the vSAN file services is enabled.
- The file servers are started with the IP addresses that were assigned during the vSAN File Services configuration process.
- A File Services VM (FSVM) is placed on each host.

Note The FSVMs are managed by the vSAN File Services. Do not perform any operation on the FSVMs.

Edit vSAN File Service

You can edit and reconfigure the settings of a vSAN File Service.

Prerequisites

- If you are upgrading from vSAN 7.0 to 7.0 Update 1, you can create SMB and NFS Kerberos file shares. This requires configuring the Active Directory domain to vSAN File Service.
- If there are active shares, changing the Active Directory domain is not permitted as this action can disrupt the user permissions on the active shares.

- If your Active Directory password has been changed, then you can edit the Active Directory configuration settings and provide the new password.

Note This action might cause minor disruption to the inflight I/Os on the file shares.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > Services**.
- 2 On the File Service row, click **Edit**.

The Configure File Service wizard opens.

- 3 Make the appropriate configuration changes. You can make the following changes to the vSAN File Service configuration:

Page	Editable Fields
Domain	<p>You can edit the following domain-related information:</p> <ul style="list-style-type: none"> ■ File service domain ■ DNS servers ■ DNS suffixes ■ Directory service <p>Note Changing domain information is a disruptive action. It might require all clients to use new URLs to reconnect to the file shares.</p>
Networking	<p>You can edit the following networking-related information:</p> <ul style="list-style-type: none"> ■ Subnet mask ■ Gateway
IP Pool	<p>You can edit the static IP addresses and DNS names, except the primary IP address and DNS name.</p>

After making necessary changes, review the changes in the Review page and click **Finish**.

Results

The changes are applied to the vSAN File Service configuration.

Create a File Share

When the vSAN file service is enabled, you can create one or more file shares on the vSAN datastore. vSAN File Service does not support using these file shares as datastores on ESXi.

Prerequisites

If you are creating an SMB file share or a NFSv4.1 file share with Kerberos security, then ensure that you have configured vSAN File Service to an AD domain.

Considerations for Share Name and Usage

- Usernames with non-ascii characters can be used to access share data.

- Share names cannot exceed 80 characters and can contain English characters, numbers, and hyphen character. Every hyphen character must be preceded and followed by a number or alphabet. Consecutive hyphens are not allowed.
- For SMB type shares, file and directories can contain any Unicode compatible strings.
- For pure NFSv4 type shares, the file and directories can contain any UTF-8 compatible strings.
- For pure NFSv3 and NFSv3+NFSv4 shares file and directories can contain only ASCII compatible strings.
- Migrating any share data from older NFSv3 to new vSAN File Service shares with NFSv4 only requires conversion of all file and directories names to UTF-8 encoding. There are third part tools to achieve the same.

Procedure

1 Navigate to the vSAN cluster and click **Configure > vSAN > File Shares**.

2 Click **Add**.

The Create file share wizard opens.

3 In the General page, enter the following information and click **Next**.

- **Name:** Enter a name for the file share.
- **Protocol:** Select an appropriate protocol. vSAN File Service supports SMB and NFS file system protocols.

If you select the **SMB** protocol, you can also configure the SMB file share to accept only the encrypted data using the **Protocol encryption** option.

If you select the **NFS** protocol, you can configure the file share to support either **NFS 3**, **NFS 4**, or both **NFS 3 and NFS 4** versions. If you select **NFS 4** version, you can set either **AUTH_SYS** or **Kerberos** security.

Note SMB protocol and Kerberos security for NFS protocol can be configured only if the vSAN File Service is configured with Active Directory. For more information, see [Configure File Services](#).

- With SMB protocol, you can hide the files and folders that the share client user does not have permission to access using the **Access based enumeration** option.
- **Storage Policy:** Select an appropriate storage policy.
- **Affinity site:** This option is available if you are creating a file share on a stretched cluster. This option helps you place the file share on a file server that belongs to the site of your choice. Use this option when you prefer low latency while accessing the file share. The default value is **Either**, which indicates that the file share is placed on a site with less traffic on either preferred or secondary site.

- **Storage space quotas:** You can set the following values:
 - **Share warning threshold:** When the share reaches this threshold, a warning message is displayed.
 - **Share hard quota:** When the share reaches this threshold, new block allocation is denied.
 - **Labels:** A label is a key-value pair that helps you organize file shares. You can attach labels to each file share and then filter them based on their labels. A label key is a string with 1~250 characters. A label value is a string and the length of the label value should be less than 1k characters. vSAN File Service supports up to 5 labels per share.
- 4 The Net access control page, provides options to define access to the file share. Net access control options are available only for NFS shares. Select one of the following options and click **Next**.
- **No access:** Select this option to make the file share inaccessible from any IP address.
 - **Allow access from any IP:** Select this option to make the file share accessible from all IP addresses.
 - **Customize net access:** Select this option to define permissions for specific IP addresses. Using this option you can specify whether a particular IP address can access, make changes, or only read the file share. You can also enable or disable **Root squash** for each IP address. You can enter the IP addresses in the following formats:
 - A single IP address. For example, 123.23.23.123
 - IP address along with a subnet mask. For example, 123.23.23.0/8
 - A range by specifying a starting IP address and ending IP address separated by a hyphen (-). For example, 123.23.23.123-123.23.23.128
 - Asterisk (*) to imply all the clients.
- 5 In the Review page, review the settings, and then click **Finish**.
- A new file share is created on the vSAN datastore.

View File Shares

You can view the list of vSAN file shares.

To view the list of vSAN file shares, navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.

A list of vSAN file shares appears. For each file share, you can view information such as storage policy, hard quota, usage over quota, actual usage, and so on.

Access File Shares

You can access a file share from a host client.

Access NFS File Share

You can access a file share from a host client, using an operating system that communicates with NFS file systems. For RHEL-based Linux distributions, NFS 4.1 support is available in RHEL 7.3 and CentOS 7.3-1611 running kernel 3.10.0-514 or later. For Debian based Linux distributions, NFS 4.1 support is available in Linux kernel version 4.0.0 or later. All NFS clients must have unique hostnames for NFSv4.1 to work. You can use the Linux mount command with the Primary IP to mount a vSAN file share to the client. For example: `mount -t nfs4 -o minorversion=1,sec=sys <primary ip>:/vsanfs/<share name>`. NFSv3 support is available for RHEL-based and Debian based Linux distributions. You can use the Linux mount command to mount a vSAN file share to the client. For example: `mount -t nfs vers=3 <nfsv3_access_point> <localmount_point>`.

Example

Sample v41 commands for verifying the NFS file share from a host client:

```
[root@localhost ~]# mount -t nfs4 -o minorversion=1,sec=sys <primary ip address>:/vsanfs/
TestShare-0 /mnt/TestShare-0
[root@localhost ~]# cd /mnt/TestShare-0/
[root@localhost TestShare-0]# mkdir bar
[root@localhost TestShare-0]# touch foo
[root@localhost TestShare-0]# ls -l
total 0
drwxr-xr-x. 1 root root 0 Feb 19 18:35 bar
-rw-r--r--. 1 root root 0 Feb 19 18:35 foo
```

Access NFS Kerberos File Share

A Linux client accessing an NFS Kerberos share should have a valid Kerberos ticket.

Sample v41 commands for verifying the NFS Kerberos file share from a host client:

An NFS Kerberos share can be mounted using the following mount command:

```
[root@localhost ~]# mount -t nfs4 -o minorversion=1,sec=krb5/krb5i/krb5p <primary ip
address>:/vsanfs/TestShare-0 /mnt/TestShare-0
[root@localhost ~]# cd /mnt/TestShare-0/
[root@localhost TestShare-0]# mkdir bar
[root@localhost TestShare-0]# touch foo
[root@localhost TestShare-0]# ls -l
total 0
drwxr-xr-x. 1 root root 0 Feb 19 18:35 bar
-rw-r--r--. 1 root root 0 Feb 19 18:35 foo
```

Changing Ownership of a NFS Kerberos share

You must log in using the AD domain user name for changing the ownership of a share. The AD domain user name provided in the file service configuration acts as a sudo user for the Kerberos file share.

```
[root@localhost ~]# mount -t nfs4 -o minorversion=1,sec=sys <primary ip address>:/vsanfs/
TestShare-0 /mnt/TestShare-0
[fsadmin@ocalhost ~]# chown user1 /mnt/TestShare-0
```

```
[user1@localhost ~]# ls -l /mnt/TestShare-0
total 0
drwxr-xr-x. 1 user1 domain users 0 Feb 19 18:35 bar
-rw-r--r--. 1 user1 domain users 0 Feb 19 18:35 foo
```

Access SMB File Share

You can access an SMB file share from a Windows client.

Prerequisites

Ensure that the Windows client is joined to the Active Directory domain that is configured with vSAN File Service.

Procedure

- 1 Copy the SMB file share path using the following procedure:
 - a Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
List of all the vSAN file shares appears.
 - b Select the SMB file share that you want to access from the Windows client.
 - c Click **COPY PATH > SMB**.
The SMB file share path gets copied to your clipboard.
- 2 Log into the Windows client as a normal Active Directory domain user.
- 3 Access the SMB file share using path that you have copied.

Edit a File Share

You can edit the settings of a vSAN file share.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
List of all the vSAN file shares appears.
- 2 Select the file share that you want to modify and click **EDIT**.
- 3 In the Edit file share page, make appropriate changes to the file share settings and click **Finish**.

Results

The file share settings are updated.

Note vSAN does not allow file share protocol change between SMB and NFS.

Manage SMB File Share

vSAN File Service supports the shared folders snap-in for the Microsoft Management Console (MMC) for managing the SMB shares on the vSAN cluster.

You can perform the following tasks on vSAN File System SMB shares using the MMC tool:

- Manage Access Control List (ACL).
- Close open files.
- View active sessions.
- View open files.
- Close client connections.

Procedure

- 1 Copy the MMC Command using the following procedure:
 - a Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
List of all the vSAN file shares appears.
 - b Select the SMB file share that you want to manage from the Windows client using the MMC tool.
 - c Click **COPY MMC COMMAND**.
The MMC command gets copied to your clipboard.
- 2 Log into the Windows client as a file server admin user. You can configure a user as a file server admin user when you enable file service. A file service admin user has all the privileges on the file server.
- 3 In the search box on the taskbar, type Run, and then select **Run**.
- 4 In the Run box, run the MMC command that you have copied to access and manage the SMB share using the MMC tool.

Delete a File Share

You can delete a file share when you no longer need it. When you delete a file share, all the snapshots associated with that file share are also deleted.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
List of all the vSAN file shares appears.
- 2 Select the file share that you want to modify and click **DELETE**.
- 3 On the Delete file shares dialogue, click **DELETE**.

vSAN Distributed File System Snapshot

A snapshot provides a space-efficient and time-based archive of the data. It provides the ability to retrieve data from a file or a set of files in the event of accidental deletion of a file. A file system level snapshot provides you information about the files that have been changed and the changes made to the file. It provides you an automated file recovery service and it is more efficient compared to the traditional tape-based backup method. A snapshot on its own does not provide a full disaster recovery solution but it can be used by the third-party backup vendors to copy the changed files (incremental backup) to a different physical location.

vSAN File Services has a built-in feature that allows you to create a point-in-time image of the vSAN file share. When the vSAN File Service is enabled, you can create up to 32 snapshots per share. A vSAN file share snapshot is a file system snapshot that provides a point-in-time image of a vSAN file share.

Create a Snapshot

When the vSAN file service is enabled, you can create one or more snapshots that provide a point-in-time image of the vSAN file share. You can create a maximum of 32 snapshots per file share.

Prerequisites

You should have created a vSAN file share.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
A list of vSAN file shares appears.
- 2 Select the file share for which you want to create a snapshot and then click **SNAPSHOTS > NEW SNAPSHOT**.
Create new snapshot dialogue appears.
- 3 On the Create new snapshot dialogue, provide a name for the snapshot, and click **Create**.

Results

A point-in-time snapshot for the selected file share is created.

View a Snapshot

You can view the list of snapshots along with the information such as date and time of the snapshot creation, and its size.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.
A list of vSAN file shares appears.
- 2 Select a file share and click **SNAPSHOTS**.

Results

A list of snapshots for that file share appears. You can view information such as date and time of the snapshot creation, and its size.

Delete a Snapshot

You can delete a snapshot when you no longer need it.

Procedure

- 1 Navigate to the vSAN cluster and click **Configure > vSAN > File Service Shares**.

A list of vSAN file shares appears.

- 2 Select a file share and click **SNAPSHOTS**.

A list of snapshots of that belongs to the file share you have selected appears.

- 3 Select the snapshot that you want to delete and click **DELETE**.

Rebalance Workload on vSAN File Service Hosts

Skyline Health displays the workload balance health status for all the hosts that are part of the vSAN File Service Infrastructure.

If there is an imbalance in the workload of a host, you can correct it by rebalancing the workload.

Prerequisites

Procedure

- 1 Navigate to the vSAN cluster and then click **Monitor > vSAN > Skyline Health**.

- 2 Under Skyline Health, expand **File Service** and then click **Infrastructure Health**.

The Infrastructure Health tab displays a list of all the hosts that are part of the vSAN File Service infrastructure. For each host, the status of workload balance is displayed. If there is an imbalance in the workload of a host, an alert is displayed in the **Description** column.

- 3 Click **REMEDIATE IMBALANCE** and then **REBALANCE** to fix the imbalance.

Before proceeding with rebalancing, consider the following:

- During rebalancing, containers in the hosts with an imbalanced workload might be moved to other hosts. The rebalancing activity might also impact the other hosts in the cluster.
- During the rebalance process, the workloads running on NFS shares are not disrupted. However, the I/O to SMB shares located in the containers that have moved are disrupted.

Results

The host workload is balanced and the workload balance status turns green.

Reclaiming Space with Unmap

vSAN 6.7 Update 2 and later supports UNMAP commands that enable you to reclaim storage space that is mapped to deleted files in the vSAN Distributed File System (VDFS) created by the guest on the vSAN object.

Deleting or removing files and snapshots 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 VDFS when you delete file shares and snapshots, consolidate file shares and snapshots, and so on. You can unmap storage space when you delete files or snapshots

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.

Upgrade File Service

When you upgrade the file service, the upgrade is performed on a rolling basis. During the upgrade, the file server containers running on the virtual machines which are undergoing upgrade fails over to other virtual machines. The file shares remain accessible during the upgrade. During the upgrade, you might experience some interruptions while accessing the file shares.

Prerequisites

Ensure that the following are upgraded:

- ESXi Hosts
- vCenter Server
- vSAN disk format

Procedure

- 1 Navigate to the vSAN cluster and then click **Configure > vSAN > Services**.
- 2 Under vSAN Services, on the File Service row, click **CHECK UPGRADE**.

- 3 In the Upgrade File Service dialog box, select one of the following deployment options and then click **UPGRADE**.

Option	Action
Automatic approach	This is the default option. This option lets the system search and download the OVF. After the upgrade begins, you cannot cancel the task. Note vSAN requires internet connectivity for this option.
Manual approach	This option allows you to browse and select an OVF that is already available on your local system. After the upgrade begins, you cannot cancel the task. Note If you select this option, you should upload all the following files: <ul style="list-style-type: none"> ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x_OVF10.mf ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x_OVF10.cert ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-system.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-cloud-components.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x-log.vmdk ■ VMware-vSAN-File-Services-Appliance-x.x.x.x-x_OVF10.ovf

Monitor Performance

You can monitor the performance of NFS and SMB file shares.

Prerequisites

Ensure that vSAN Performance Service is enabled. If you are using the vSAN Performance Service for the first time, you see a message alerting you to enable it. For more information about vSAN Performance Service, see the *vSAN Monitoring and Troubleshooting Guide*.

Procedure

- 1 Navigate to the vSAN cluster and then click **Monitor > vSAN > Performance**.
- 2 Click the **FILE SHARE** tab.
- 3 Select one of the following options:

Option	Action
Time Range	<ul style="list-style-type: none"> ■ Select Last to select the number of hours for which you want to view the performance report. ■ Select CUSTOM to select the date and time for which you want to view the performance report. ■ Select SAVE to add the current setting as an option to the Time Range list.
File share	Select the file share for which you want to generate and view the performance report.

- 4 Click **SHOW RESULTS**.

Results

The throughput, IOPS, and latency metrics of the vSAN file service for the selected period are displayed.

For more information on vSAN Performance Graphs, see the VMware knowledge base article at <https://kb.vmware.com/s/article/2144493>.

Monitor Capacity

You can monitor the capacity for both native file shares and CNS-managed file shares.

Procedure

- 1 Navigate to the vSAN cluster and then click **Monitor > vSAN > Capacity**.
- 2 Click **CAPACITY USAGE** tab.
- 3 In the Usage breakdown before dedupe and compression section, expand **User objects**.

Results

The file share capacity information is displayed.

For more information about monitoring vSAN capacity, see the *vSAN Monitoring and Troubleshooting Guide*.

Monitor Health

You can monitor the health of both vSAN file service and file share objects.

View vSAN File Service Health

You can monitor the vSAN file service health.

Prerequisites

Ensure that vSAN Performance Service is enabled.

Procedure

- 1 Navigate to the vSAN cluster and then click **Monitor > vSAN**.
- 2 In the Skyline Health section, expand **File Service**.
- 3 Click the following file service health parameters to view the status.

Option	Action
Infrastructure health	Displays the file service infrastructure health status per ESXi host. For more information, click the Info tab.
File Server Health	Displays the file server health status. For more information, click the Info tab.
Share health	Displays the file service share health. For more information, click the Info tab.

Monitor File Share Objects Health

You can monitor the health of file share objects.

To view the file share object health, navigate to the vSAN cluster and then click **Monitor > vSAN > Virtual Objects**.

The device information such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts is displayed in the VIEW PLACEMENT DETAILS section.

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, click **...**, and then click **Remove**.
 - 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.

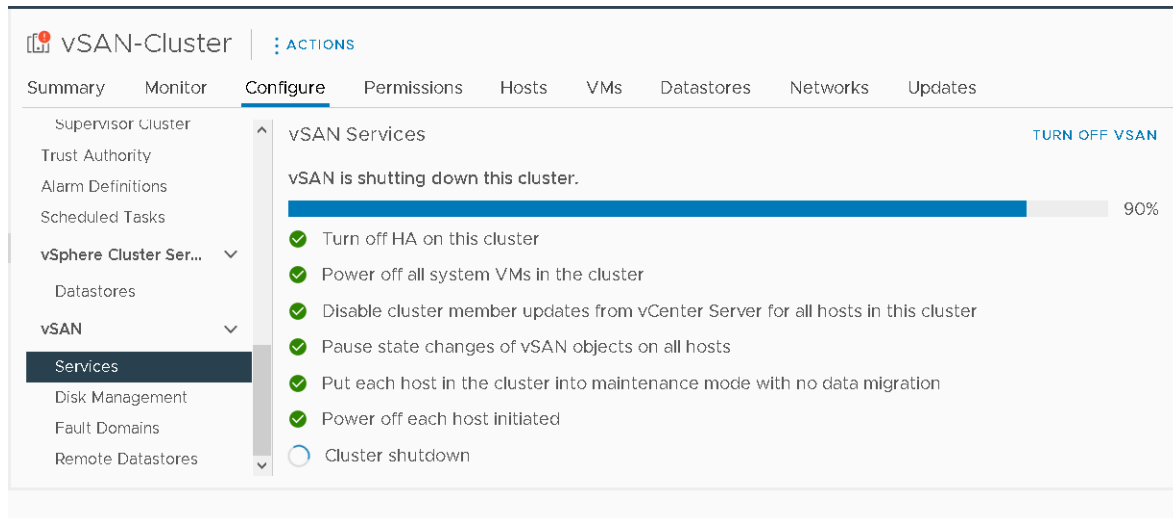
Verify that no partitions exist on the flash devices.
- 5 Create the all-flash disk groups on each host.

Shutting Down and Restarting the vSAN Cluster

You can shut down the entire vSAN cluster to perform maintenance or troubleshooting.

Use the Shutdown Cluster wizard to shutdown the vSAN cluster. The wizard performs the necessary steps and alerts you when it requires user action. You also can manually shut down the cluster, if necessary.

Note When you shut down a stretched cluster, the witness host remains active.



Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard

Use the Shutdown cluster wizard to gracefully shut down the vSAN cluster for maintenance or troubleshooting.

Note If you have a vSphere with Tanzu environment, you must follow the specified order when shutting down or starting up the components. For more information, see "Shut Down the vSphere with Tanzu Workload Domain" in the *VMware Validated Design Documentation*.

Procedure

- 1 Prepare the vSAN cluster for shutdown.
 - a Check the vSAN health service to confirm that the cluster is healthy.
 - b Power off all virtual machines (VMs) stored in the vSAN cluster, except for vCenter Server VMs, vCLS VMs and file service VMs. If vCenter Server is hosted in the vSAN cluster, do not power off the vCenter Server VM.
 - c If this is an HCI Mesh server cluster, power off all client VMs stored on the cluster. If the client cluster's vCenter Server VM is stored in this cluster, either migrate or power off the VM. Once this server cluster is shutdown, its shared datastore is inaccessible to clients.
 - d Verify that all resynchronization tasks are complete.

Click the **Monitor** tab and select **vSAN > Resyncing Objects**.

Note If any member hosts are in lockdown mode, add the host's root account to the security profile Exception User list. For more information, see Lockdown Mode in *vSphere Security*.

- 2 Right-click the vSAN cluster in the vSphere Client, and select menu **Shutdown cluster**.
You also can click **Shutdown Cluster** on the vSAN Services page.

- 3 On the Shutdown cluster wizard, verify that the Shutdown pre-checks are green checks. Resolve any issues that are red exclamations. Click **Next**.

If vCenter Server appliance is deployed on the vSAN cluster, the Shutdown wizard displays the vCenter Server notice. Note the IP address of the orchestration host, in case you need it during the cluster restart. Click **Next**.

- 4 Enter a reason for performing the shutdown, and click **Shutdown**.

The vSAN Services page changes to display information about the shutdown process.

- 5 Monitor the shutdown process.

vSAN performs the steps to shutdown the cluster, powers off the system VMs, and powers off the hosts.

Restart the vSAN Cluster

You can restart a vSAN cluster that is shut down for maintenance or troubleshooting.

Procedure

- 1 Power on the cluster hosts.

If the vCenter Server is hosted on the vSAN cluster, wait for vCenter Server to restart.

- 2 Right-click the vSAN cluster in the vSphere Client, and select menu **Restart cluster**.

You also can click **Restart Cluster** on the vSAN Services page.

- 3 On the Restart Cluster dialog, click **Restart**.

The vSAN Services page changes to display information about the restart process.

- 4 After the cluster has restarted, check the vSAN health service and resolve any outstanding issues.

Manually Shut Down and Restart the vSAN Cluster

You can manually shut down the entire vSAN cluster to perform maintenance or troubleshooting.

Use the Shutdown Cluster wizard unless your workflow requires a manual shut down. When you manually shut down the vSAN cluster, do not disable vSAN on the cluster.

Note If you have a vSphere with Tanzu environment, you must follow the specified order when shutting down or starting up the components. For more information, see "Shut Down the vSphere with Tanzu Workload Domain" in the *VMware Validated Design Documentation*.

Procedure

1 Shut down the vSAN cluster.

- a Check the vSAN health service to confirm that the cluster is healthy.
- b Power off all virtual machines (VMs) running in the vSAN cluster, if vCenter Server is not hosted on the cluster. If vCenter Server is hosted in the vSAN cluster, do not power off the vCenter Server VM.
- c Click the **Configure** tab and turn off HA. As a result, the cluster does not register host shutdowns as failures.
- d Verify that all resynchronization tasks are complete.

Click the **Monitor** tab and select **vSAN > Resyncing Objects**.

- e If vCenter Server is hosted on the vSAN cluster, power off the vCenter Server VM.
Make a note of the host that runs the vCenter Server VM. It is the host where you must restart the vCenter Server VM.
- f Disable cluster member updates from vCenter Server by running the following command on the ESXi hosts in the cluster. Ensure that you run the following command on all the hosts.

```
esxcfg-advcfg -s 1 /VSAN/IgnoreClusterMemberListUpdates
```

- g Log in to any host in the cluster other than the witness host.
- h Run the following command only on that host. If you run the command on multiple hosts concurrently, it may cause a race condition causing unexpected results.

```
python /usr/lib/vmware/vsan/bin/reboot_helper.py prepare
```

The command returns and prints the following:

```
Cluster preparation is done.
```

Note

- The cluster is fully partitioned after the successful completion of the command.
 - If you encounter an error, resolve the issue based on the error message and try enabling vCLS retreat mode again.
 - If there are unhealthy or disconnected hosts in the cluster, remove the hosts and retry the command.
-

- i Place all the hosts into maintenance mode with **No Action**. If the vCenter Server is powered off, use the following command to place the ESXi hosts into maintenance mode with **No Action**.

```
esxcli system maintenanceMode set -e true -m noAction
```

Perform this step on all the hosts.

To avoid the risk of data unavailability while using **No Action** at the same time on multiple hosts, followed by a reboot of multiple hosts, see the VMware knowledge base article at <https://kb.vmware.com/s/article/60424>. To perform simultaneous reboot of all hosts in the cluster using a built-in tool, see the VMware knowledge base article at <https://kb.vmware.com/s/article/70650>.

- j After all hosts have successfully entered maintenance mode, perform any necessary maintenance tasks and power off the hosts.

2 Restart the vSAN cluster.

- a Power on the ESXi hosts.

Power on the physical box where ESXi is installed. The ESXi host starts, locates the VMs, and functions normally.

If any hosts fail to restart, you must manually recover the hosts or move the bad hosts out of the vSAN cluster.

- b When all the hosts are back after powering on, exit all hosts from maintenance mode. If the vCenter Server is powered off, use the following command on the ESXi hosts to exit maintenance mode.

```
esxcli system maintenanceMode set -e false
```

Perform this step on all the hosts.

- c Log in to one of the hosts in the cluster other than the witness host.
- d Run the following command only on that host. If you run the command on multiple hosts concurrently, it may cause a race condition causing unexpected results.

```
python /usr/lib/vmware/vsan/bin/reboot_helper.py recover
```

The command returns and prints the following:

```
Cluster reboot/power-on is completed successfully!
```

- e Verify that all the hosts are available in the cluster by running the following command on each host.

```
esxcli vsan cluster get
```

- f Enable cluster member updates from vCenter Server by running the following command on the ESXi hosts in the cluster. Ensure that you run the following command on all the hosts.

```
esxcfg-advcfg -s 0 /VSAN/IgnoreClusterMemberListUpdates
```

- g Restart the vCenter Server VM if it is powered off. Wait for the vCenter Server VM to be powered up and running. To disable vCLS retreat mode, see the VMware knowledge base article at <https://kb.vmware.com/s/article/80472>.
- h Verify again that all the hosts are participating in the vSAN cluster by running the following command on each host.

```
esxcli vsan cluster get
```

- i Restart the remaining VMs through vCenter Server.
- j Check the vSAN health service and resolve any outstanding issues.
- k (Optional) If the vSAN cluster has vSphere Availability enabled, you must manually restart vSphere Availability to avoid the following error: `Cannot find vSphere HA master agent`.

To manually restart vSphere Availability, select the vSAN cluster and navigate to:

- 1 **Configure > Services > vSphere Availability > EDIT > Disable vSphere HA**
- 2 **Configure > Services > vSphere Availability > EDIT > Enable vSphere HA**

- 3 If there are unhealthy or disconnected hosts in the cluster, recover or remove the hosts from the vSAN cluster. Retry the above commands only after the vSAN health service shows all available hosts in the green state.

If you have a three-node vSAN cluster, the command `reboot_helper.py recover` cannot work in a one host failure situation. As an administrator, do the following:

- a Temporarily remove the failure host information from the unicast agent list.
- b Add the host after running the following command.

```
reboot_helper.py recover
```

Following are the commands to remove and add the host to a vSAN cluster:

```
#esxcli vsan cluster unicastagent remove -a <IP Address> -t node -u <NodeUuid>
```

```
#esxcli vsan cluster unicastagent add -t node -u <NodeUuid> -U true -a <IP Address> -p 12321
```

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.

Each disk group includes a single flash cache device. You can create multiple disk groups manually, and claim a flash cache device for each 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 vSAN Direct

Use vSAN Direct to enable stateful services to access raw, non-vSAN local storage through a direct path.

You can claim host-local devices for vSAN Direct, and use vSAN to manage and monitor those devices. On each local device, vSAN Direct creates an independent VMFS datastore and makes it available to your stateful application.

Each local vSAN Direct datastore appears as a vSAN-D datastore.

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.

Note Only the vSAN Data Persistence platform can consume vSAN Direct storage. The vSAN Data Persistence platform provides a framework for software technology partners to integrate with VMware infrastructure. Each partner must develop their own plug-in for VMware customers to receive the benefits of the vSAN Data Persistence platform. The platform is not operational until the partner solution running on top is operational. For more information, see *vSphere with Tanzu Configuration and Management*.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.
- 4 Click **Claim unused disks**.
- 5 Group by host.
- 6 Select disks to claim.
 - Select the flash device to use for the cache tier.
 - Select the disks to use for the capacity tier.
- 7 Click **Create** or **OK** to confirm your selections.

Results

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.

In this method, you select devices to create a disk groups for the vSAN cluster. You need one cache device and at least one capacity device for each disk group.

Note Only the vSAN Data Persistence platform can consume vSAN Direct storage. The vSAN Data Persistence platform provides a framework for software technology partners to integrate with VMware infrastructure. Each partner must develop their own plug-in for VMware customers to receive the benefits of the vSAN Data Persistence platform. The platform is not operational until the partner solution running on top is operational. For more information, see *vSphere with Tanzu Configuration and Management*.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.
- 4 Click **Claim Unused Disks**.
- 5 Select devices to add to disk groups.
 - 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 **Claim for cache tier**.
 - Select an HDD device to be used as capacity and click **Claim for capacity tier**.
 - Click **Create** or **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 **Claim for cache tier**.
 - Select a flash device to be used for capacity and click **Claim for capacity tier**.
 - Click **Create** or **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.

Claim Disks for vSAN Direct

You can claim local storage devices as vSAN Direct for use with the vSAN Data Persistence Platform.

Note Only the vSAN Data Persistence platform can consume vSAN Direct storage. The vSAN Data Persistence platform provides a framework for software technology partners to integrate with VMware infrastructure. Each partner must develop their own plug-in for VMware customers to receive the benefits of the vSAN Data Persistence platform. The platform is not operational until the partner solution running on top is operational. For more information, see *vSphere with Tanzu Configuration and Management*.

Procedure

- 1 In the vSphere Client, navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, click **Disk Management**.
- 4 Click **Claim Unused Disks**.
- 5 On the Claim Unused Disks wizard, select the vSAN Direct tab.
- 6 Select a device to claim, and select the checkbox to **Claim for vSAN Direct**.

Note Devices claimed for your vSAN cluster do not appear in the vSAN Direct tab.

- 7 Click **Create**.

Results


For each device you claim, vSAN creates a new vSAN Direct datastore.

What to do next

You can click the Datastores tab to display the vSAN Direct datastores in your cluster.

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. You can also add or remove disks that are claimed using the vSAN Direct.



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 Disks**.
- 5 Select the device that you want to add and click **Add**.

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*.

Check a Disk or Disk Group's Data Migration Capabilities

Use the data migration pre-check to determine the impact of data migration options when unmounting a disk or disk group, or removing it from the vSAN cluster.

Run the data migration pre-check before you unmount or remove a disk or disk group from the vSAN cluster. The test results provide information to help you determine the impact to cluster capacity, predicted health checks, and any objects that will go out of compliance. If the operation will not succeed, pre-check provides information about what resources are needed.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the Monitor tab.
- 3 Under vSAN, click **Data Migration Pre-check**.
- 4 Select a disk or disk group, choose a data migration option, and click **Pre-check**.

vSAN runs the data migration precheck tests.

- 5 View the test results.

The pre-check results show whether you can safely unmount or remove the disk or disk group.

- The Object Compliance and Accessibility tab displays objects that might have issues after the data migration.
- The Cluster Capacity tab displays the impact of data migration on the vSAN cluster before and after you perform the operation.
- The Predicted Health tab displays the health checks that might be affected by the data migration.

What to do next

If the pre-check indicates that you can unmount or remove the device, click the option to continue the operation.

Remove Disk Groups or Devices from vSAN

You can remove selected devices from the disk group, or you can remove 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

Run data migration pre-check on the device or disk group before you remove it from the cluster. For more information, see

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

- 5 Click **Yes** or **Remove** to confirm.

The data is evacuated from the selected devices or disk group.

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.
- 5 Click **...**, then click the **Recreate**.
The Recreate Disk Group dialog box appears.
- 6 Select a data migration mode, and click **Recreate**.

Results

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
Turn on LED	Enables locator LED on the selected storage device. You can enable locator LEDs from the Manage tab and click Storage > Storage Devices .
Turn off LED	Disables locator LED on the selected storage device. You can disable locator LEDs from the Manage tab and click Storage > Storage Devices .

Mark Devices as Flash

When flash devices are not automatically identified as flash by ESXi hosts, you can manually mark them as local flash devices.

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**.
- 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**.
- 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 disk**.
- 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**.
- 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 Disks** 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, select **Ineligible**.
- 6 Select a device from the list and click **Erase partitions**.
- 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 use deduplication and compression on a vSAN cluster to eliminate duplicate data and reduce the amount of space required to store data. Or you can use compression-only vSAN to reduce storage requirements without compromising server performance.

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 deleted files in the file system created by the guest on the 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 a higher host-to-flash I/O throughput and improve the flash endurance.

vSAN also supports the SCSI UNMAP commands issued directly from a guest operating system to reclaim storage space. vSAN supports offline unmaps and 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. Use compression-only vSAN for workloads that do not benefit from deduplication, such as online transactional processing.

Deduplication occurs inline when data is written back from the cache tier to the capacity tier. The deduplication algorithm uses a fixed block size and is applied within each disk group. Redundant copies of a block within the same disk group are deduplicated.

Deduplication and compression are enabled 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.

Note Compression-only vSAN is applied on a per-disk basis.

You can enable deduplication and compression when you create a 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.

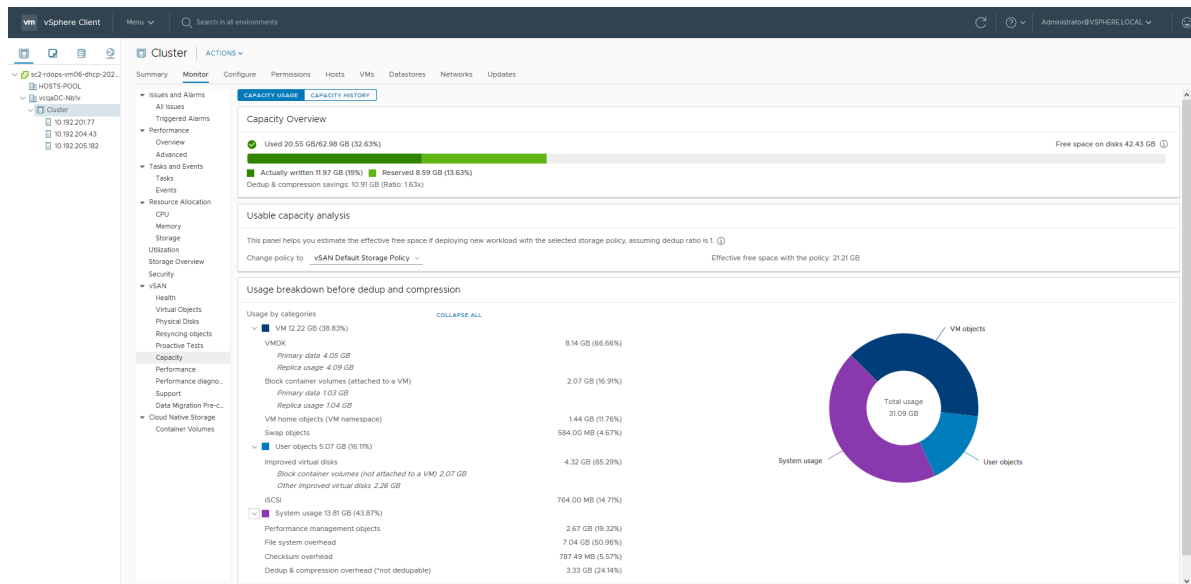
How to Manage Disks in a Cluster with Deduplication and Compression

Consider the following guidelines when managing disks in a cluster with deduplication and compression enabled. These guidelines do not apply to compression-only vSAN.

- Avoid adding disks to a disk group incrementally. For more efficient deduplication and compression, consider adding a disk group to increase the 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 Usage breakdown before dedup and compression in the vSAN Capacity monitor.



You can view the Usage breakdown before dedup and compression 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.

- 3 Under vSAN, select **Services**.
 - a Click to edit Space Efficiency.
 - b Select a space efficiency option: Deduplication and compression, or Compression only.
 - c (Optional) Select **Allow Reduced Redundancy**. If needed, vSAN reduces the protection level of your VMs while enabling Deduplication and Compression. For more details, see [Reducing VM Redundancy for vSAN Cluster](#).
- 4 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.
- 3 Under vSAN, select **Services**.
 - a Click to edit Space Efficiency.
 - b Select a space efficiency option: Deduplication and compression, or Compression only.
 - c (Optional) Select **Allow Reduced Redundancy**. If needed, vSAN reduces the protection level of your VMs while enabling Deduplication and Compression. For more details, see [Reducing VM Redundancy for vSAN Cluster](#).
- 4 Click **Apply** to save your configuration changes.

Results

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.
 - a Under vSAN, select **Services**.
 - b Click **Edit**.
 - c Disable Deduplication and Compression.
 - d (Optional) Select **Allow Reduced Redundancy**. If needed, vSAN reduces the protection level of your VMs, while disabling Deduplication and Compression. See [Reducing VM Redundancy for vSAN Cluster](#).
- 3 Click **Apply** or **OK** to save your configuration changes.

Results

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 **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	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 **Failures to tolerate** to 1. To use RAID 6, set **Failure tolerance method** to **RAID-5/6 (Erasure Coding) - Capacity** and **Failures to tolerate** to 2. RAID 5 or RAID 6 erasure coding does not support a **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 **Site disaster tolerance** setting.

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 in a vSAN Cluster

7

You can encrypt data-in transit in your vSAN cluster, and encrypt data-at-rest in your vSAN datastore.

vSAN can encrypt data in transit across hosts in the vSAN cluster. Data-in-transit encryption protects data as it moves around the vSAN cluster.

vSAN can encrypt data at rest in the vSAN datastore. Data-at-rest encryption protects data on storage devices, in case a device is removed from the cluster.

This chapter includes the following topics:

- [vSAN Data-In-Transit Encryption](#)
- [vSAN Data-At-Rest Encryption](#)

vSAN Data-In-Transit Encryption

vSAN can encrypt data in transit, as it moves across hosts in your vSAN cluster.

vSAN can encrypt data in transit across hosts in the cluster. When you enable data-in-transit encryption, vSAN encrypts all data and metadata traffic between hosts.

vSAN data-in-transit encryption has the following characteristics:

- vSAN uses AES-256 bit encryption on data in transit.
- vSAN data-in-transit encryption is not related to data-at-rest-encryption. You can enable or disable each one separately.
- Forward secrecy is enforced for vSAN data-in-transit encryption.
- Traffic between data hosts and witness hosts is encrypted.
- File service data traffic between the VDFS proxy and VDFS server is encrypted.
- vSAN file services inter-host connections are encrypted.

vSAN uses symmetric keys that are generated dynamically and shared between hosts. Hosts dynamically generate an encryption key when they establish a connection, and they use the key to encrypt all traffic between the hosts. You do not need a key management server to perform data-in-transit encryption.

Each host is authenticated when it joins the cluster, ensuring connections only to trusted hosts are allowed. When a host is removed from the cluster, its authentication certificate is removed.

vSAN data-in-transit encryption is a cluster-wide setting. When enabled, all data and metadata traffic is encrypted as it transits across hosts.

Enable Data-In-Transit Encryption on a vSAN Cluster

You can enable data-in-transit encryption by editing the configuration parameters of a vSAN cluster.

Procedure

- 1 Navigate to an existing cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services** and click the Data-In-Transit Encryption **Edit** button.
- 4 Click to enable **Data-In-Transit encryption**, and select a rekey interval.
- 5 Click **Apply**.

Results

Encryption of data in transit is enabled on the vSAN cluster. vSAN encrypts all data moving across hosts and file service inter-host connections in the cluster.

vSAN Data-At-Rest Encryption

vSAN can encrypt data at rest in your vSAN datastore.

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 datastore requires some preparation. After your environment is set up, you can enable data-at-rest encryption on your vSAN cluster.

Data-at-rest encryption requires an external Key Management Server (KMS) or a vSphere Native Key Provider. For more information about vSphere encryption, see *vSphere Security*.

You can use an external Key Management Server (KMS), the vCenter Server system, and your ESXi hosts to encrypt data in your vSAN cluster. 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.

How Data-At-Rest Encryption Works

When you enable data-at-rest 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**.

Encryption Key Persistence

In vSAN 7.0 Update 3 and later, data-at-rest encryption can continue to function even when the key server is temporarily offline or unavailable. With key persistence enabled, the ESXi hosts can persist the encryption keys even after a reboot.

Each ESXi host obtains the encryption keys initially and retains them in its key cache. If the ESXi host has a Trusted Platform Module (TPM), the encryption keys are persisted in the TPM across reboots. The host does not need to request encryption keys. Encryption operations can continue when the key server is unavailable, because the keys have persisted in the TPM.

Use the following commands to enable key persistence on a cluster host.

```
esxcli system settings encryption set --mode=TPM
```

```
esxcli system security keypersistence enable
```

For more information about encryption key persistence, see "Key Persistence Overview" in *vSphere Security*.

Using vSphere Native Key Provider

vSAN 7.0 Update 2 supports vSphere Native Key Provider. If your environment is set up for vSphere Native Key Provider, you can use it to encrypt virtual machines in your vSAN cluster. For more information, see "Configuring and Managing vSphere Native Key Provider" in *vSphere Security*.

vSphere Native Key Provider does not require an external Key Management Server (KMS). vCenter Server generates the Key Encryption Key and pushes it to the ESXi hosts. The ESXi hosts then generate Data Encryption Keys.

vSphere Native Key Provider can coexist with an existing key server infrastructure.

Design Considerations for Data-At-Rest Encryption

Consider these guidelines when working with data-at-rest 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. The witness host does not store customer data, only metadata, such as the size and UUID of vSAN object and components.

Note If the witness host is an appliance running on another cluster, you can encrypt the metadata stored on it. Enable data-at-rest encryption on the cluster that contains the witness host.

- Establish a policy regarding core dumps. Core dumps are encrypted because they can contain sensitive information. 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 Standard Key Provider

Use a standard key provider to distribute the keys that encrypt the vSAN datastore.

Before you can encrypt the vSAN datastore, you must set up a standard key provider 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 key provider.

The KMS must support the Key Management Interoperability Protocol (KMIP) 1.1 standard. See the *vSphere Compatibility Matrices* for details.

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 standard key provider when you add the first KMS instance. If you configure the key provider on two or more vCenter Servers, make sure you use the same key provider 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 key provider as a default. You can later change the default setting.
- After vCenter Server creates the first key provider, you can add KMS instances from the same vendor to the key provider, and configure all KMS instances to synchronize keys among them. Use the method documented by your KMS vendor.
- You can set up the key provider with only one KMS instance.
- If your environment supports KMS solutions from different vendors, you can add multiple key providers.

Prerequisites

- Verify that the Key Management 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 under Security, click **Key Providers**.

- 4 Click **Add Standard Key Provider**, enter key provider information, and click **Add Key Provider**.

You can click **Add KMS** to add more Key Management Servers.

- 5 Click **Trust**.

vCenter Server adds the key provider and displays the status as Connected.

Establish a Standard Key Provider Trusted Connection by Exchanging Certificates

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

Prerequisites

Add the standard key provider.

Procedure

- 1 Navigate to the vCenter Server.
- 2 Click **Configure** and select **Key Management Servers**.
- 3 Select the key provider.
The KMS for the key provider is displayed.
- 4 Select the KMS.
- 5 From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
- 6 Select the option appropriate for your server and follow the steps.

Option	See
vCenter Server Root CA certificate	Use the Root CA Certificate Option to Establish a Standard Key Provider Trusted Connection.
vCenter Server Certificate	Use the Certificate Option to Establish a Standard Key Provider Trusted Connection.
Upload certificate and private key	Use the Upload Certificate and Private Key Option to Establish a Standard Key Provider Trusted Connection.
New Certificate Signing Request	Use the New Certificate Signing Request Option to Establish a Standard Key Provider Trusted Connection.

Use the Root CA Certificate Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors 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 From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
- 5 Select **vCenter Root CA Certificate** and click **Next**.

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.

- 6 Copy the certificate to the clipboard or download the certificate as a file.
- 7 Follow the instructions from your KMS vendor to upload the certificate to their system.

Note Some KMS vendors 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 [Finish the Trust Setup for a Standard Key Provider](#).

Use the Certificate Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors 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 From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.

- 5 Select **vCenter Certificate** and click **Next**.

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.

- 6 Copy the certificate to the clipboard or download it as a file.
- 7 Follow the instructions from your KMS vendor to upload the certificate to the KMS.

What to do next

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Use the New Certificate Signing Request Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors 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 From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
- 5 Select **New Certificate Signing Request (CSR)** and click **Next**.
- 6 In the dialog box, copy the full certificate in the text box to the clipboard or download it as a file.

Use the **Generate new CSR** button in the dialog box only if you explicitly want to generate a CSR.
- 7 Follow the instructions from your KMS vendor to submit the CSR.
- 8 When you receive the signed certificate from the KMS vendor, click **Key Providers** again, select the key provider, and from the **Establish Trust** drop-down menu, select **Upload Signed CSR Certificate**.
- 9 Paste the signed certificate into the bottom text box or click **Upload File** and upload the file, and click **Upload**.

What to do next

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Use the Upload Certificate and Private Key Option to Establish a Standard Key Provider Trusted Connection

Some Key Management Server (KMS) vendors 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 From the **Establish Trust** drop-down menu, select **Make KMS trust vCenter**.
- 5 Select **KMS certificate and private key** and click **Next**.
- 6 Paste the certificate that you received from the KMS vendor into the top text box or click **Upload a File** to upload the certificate file.
- 7 Paste the key file into the bottom text box or click **Upload a File** to upload the key file.
- 8 Click **Establish Trust**.

What to do next

Finalize the trust relationship. See [Finish the Trust Setup for a Standard Key Provider](#).

Set the Default Key Provider

You must set the default key provider if you do not make the first key provider the default, or if your environment uses multiple key providers and you remove the default one.

Prerequisites

As a best practice, verify that the Connection Status in the **Key Providers** tab shows Connected and a green check mark.

Procedure

- 1 Navigate to the vCenter Server.
- 2 Click **Configure** and select **Key Management Servers**.
- 3 Select the key provider.
- 4 Click **Make Default**.

A confirmation dialog box appears.

5 Click **Make Default**.

The key provider displays as the current default.

Finish the Trust Setup for a Standard Key Provider

Unless the **Add Standard Key Provider** 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 **Upload KMS certificate** option.
- Upload a KMS leaf certificate or the KMS CA certificate to vCenter Server by using the **Make vCenter Trust KMS** 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 Select the KMS.
- 5 Select one of the following options from the **Establish Trust** drop-down menu.

Option	Action
Make vCenter Trust KMS	In the dialog box that appears, click Trust .
Upload KMS certificate	<ol style="list-style-type: none"> a In the dialog box that appears, either paste in the certificate, or click Upload a file and browse to the certificate file. b Click Upload.

Enable Data-At-Rest Encryption on a New vSAN Cluster

You can enable data-at-rest encryption when you configure a new vSAN cluster.

Prerequisites

- Required privileges:
 - **Host.Inventory.EditCluster**
 - **Cryptographer.ManageEncryptionPolicy**
 - **Cryptographer.ManageKMS**
 - **Cryptographer.ManageKeys**

- You must have configured a standard key provider 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 or key provider.

Note Use the **Wipe residual data** check box to erase residual data from devices before you enable vSAN encryption. Make sure that you deselect this check box, unless you want to wipe existing data from the storage devices when encrypting a cluster that contains VM data. That way it ensures that the unencrypted data no longer resides on the devices after enabling vSAN encryption. This setting is not necessary for new installations that do not have any VM data on the storage devices.

- 5 Complete your cluster configuration.

Results

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

Generate New Data-At-Rest Encryption Keys

You can generate new encryption keys for data at rest, 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 key provider 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 Data-At-Rest Encryption on Existing vSAN Cluster

You can enable data-at-rest 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 configured a standard key provider 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 or key provider.

- 6 (Optional) If the storage devices in your cluster contain sensitive data, select **Wipe residual data**.

This setting directs vSAN to erase 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**.

Results

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 data-at-rest 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 vCenter Server 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 Datastore

If data-at-rest 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 data-at-rest encryption is enabled for the vSAN datastore. 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 vSphere 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](#)
- [About vSAN Object Format](#)
- [Verify the vSAN Cluster Upgrade](#)
- [Using the RVC Upgrade Command Options](#)
- [vSAN Build Recommendations for vSphere Lifecycle 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.

Table 8-1. Upgrade Prerequisite

Upgrade Prerequisites	Description
Software, hardware, drivers, firmware, and storage I/O controllers	Verify that the new version of vSAN 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 http://www.vmware.com/resources/compatibility/search.php .
vSAN version	Verify that you are using the latest version of vSAN. You cannot upgrade from a beta version to the new vSAN. 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 Allow reduced redundancy 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 <i>Administering VMware vSAN</i> documentation.</p>

Table 8-1. Upgrade Prerequisite (continued)

Upgrade Prerequisites	Description
vSAN hosts	<p>Verify that you have placed the vSAN hosts in maintenance mode and selected the Ensure data accessibility or Evacuate all data option.</p> <p>You can use the vSphere Lifecycle Manager for automating and testing the upgrade process. However, when you use vSphere Lifecycle Manager to upgrade vSAN, the default evacuation mode is Ensure data accessibility. When you use the Ensure data accessibility mode, your data is not protected, and if you encounter a failure while upgrading vSAN, you might experience unexpected data loss. However, the Ensure data accessibility mode is faster than the Evacuate all data mode, because you do not need to move all data to another host in the cluster. For information about various evacuation modes, see the <i>Administering VMware vSAN</i> documentation.</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 *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 Lifecycle 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 Lifecycle 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.

The details and level of support for an upgrade to ESXi 7.0 depend on the host to be upgraded and the upgrade method that you use. Verify that the upgrade path from your current version of ESXi to the version to which you are upgrading, is supported. For more information, see the VMware Product Interoperability Matrices at http://www.vmware.com/resources/compatibility/sim/interop_matrix.php.

Instead of performing an in-place upgrade to vCenter Server, you can use a different machine for the upgrade. For detailed instructions and upgrade options, see the *vCenter Server 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.

You can upgrade the ESXi hosts in the vSAN cluster using:

- vSphere Lifecycle Manager - By using images or baselines, vSphere Lifecycle Manager enables you to upgrade ESXi hosts in the vSAN cluster. The default evacuation mode is **Ensure data accessibility**. If you use this mode, and while upgrading vSAN you encounter a failure, data can become inaccessible until one of the hosts is back online. For information about working with evacuation and maintenance modes, see [Working with Maintenance Mode](#). For more information about upgrades and updates, see the *Managing Host and Cluster Lifecycle* documentation.
- Esxcli command - You can use components, base images, and add-ons as new software deliverables to update or patch ESXi 7.0 hosts using the manual upgrade.

When you upgrade a vSAN cluster with configured fault domains, vSphere Lifecycle Manager upgrades a host within a single fault domain and then proceeds to the next host. This ensures that the cluster has the same vSphere versions running on all the hosts. When you upgrade a stretched cluster, vSphere Lifecycle Manager upgrades all the hosts from the preferred site and then proceeds to the host in the secondary site. This ensures that the cluster has the same vSphere versions running on all the hosts. For more information on the upgrading a stretched cluster, see the *Managing Host and Cluster Lifecycle* documentation.

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 detailed instructions and upgrade options, see the *VMware ESXi Upgrade* documentation.

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 Lifecycle 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"> ■ All hosts are connected. ■ All hosts have the correct software version. ■ All disks are healthy. ■ All objects are accessible.
5%-10%	<p>Disk group upgrade. vSAN performs the initial disk upgrade with no data migration. This process takes a few minutes.</p>

Table 8-2. Upgrade Progress (continued)

Percentage of Completion	Description
10%-15%	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.
15%-95%	Disk group removal and reformat. Each disk group is removed from the cluster, reformatted, and added back to the cluster. The time required for this process varies, depending on the megabytes allocated and the system load. A system at or near its I/O capacity transfers slowly.
95%-100%	Final object version upgrade. Object conversion to the new on-disk format and resynchronization is completed. The time required for this process varies, depending on the amount of space used and whether the Allow reduced redundancy option is selected.

During the upgrade, you can monitor the upgrade process from the Resyncing Components page. See *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 **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 **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.

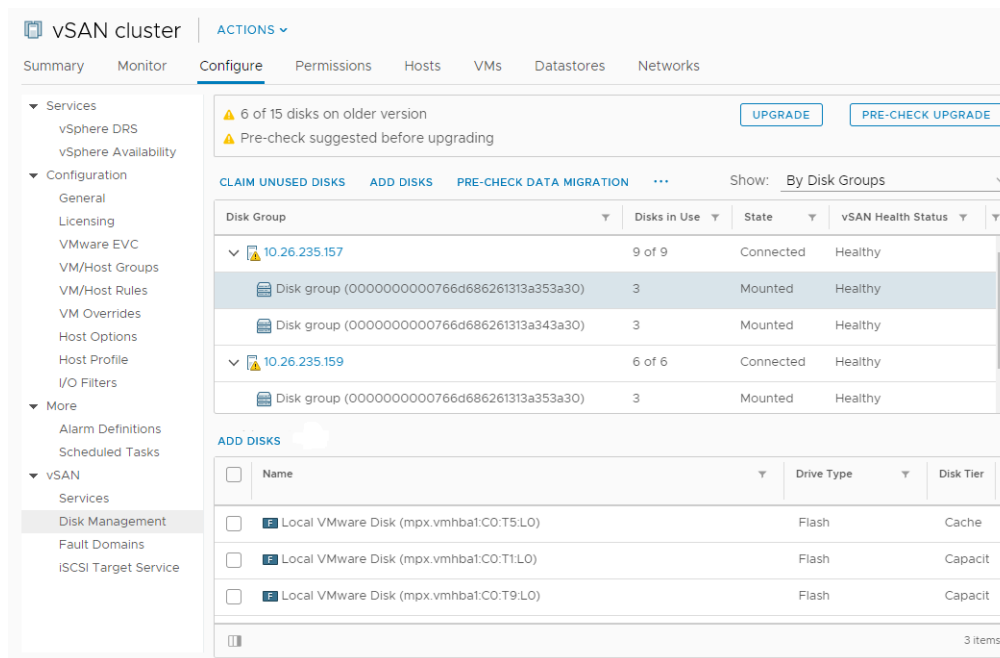
- When upgrading a three-host cluster or when upgrading a cluster with limited resources, allow the virtual machines to operate in a reduced redundancy mode. Run the RVC command with the option, `vsan.ondisk_upgrade --allow-reduced-redundancy`.

- 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 shows the navigation tree with 'vSAN' expanded to 'Disk Management'. The main content area displays a warning: '6 of 15 disks on older version' and 'Pre-check suggested before upgrading'. Below this, there are buttons for 'UPGRADE' and 'PRE-CHECK UPGRADE'. A table shows the 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

Below the table, there is an 'ADD DISKS' section with a table of local VMWare disks:

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

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.

Results

vSAN recreates 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 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 in the Disk Format Version column.

About vSAN Object Format

The operations space needed by vSAN to perform policy change or other such operations on an object created by vSAN 7.0 or earlier is the space used by a largest object in the cluster. This is typically difficult to plan for and hence the guidance was to keep 30 percent of free space in the cluster assuming that it is unlikely that the largest object in the cluster consumes more than 25 percent of the space and 5 percent of the space is reserved to make sure cluster does not become full due to policy changes. In vSAN 7.0U1 and later, all objects are created in a new format which allows the operations space needed by vSAN to perform policy change on an object if there is 255 GB per host for objects less than 8 TB and 765 GB per host for objects 8 TB or larger.

After a cluster is upgraded to vSAN 7.0 U1 or later from vSAN 7.0 or earlier release, the objects greater than 255 GB created with the older release must be rewritten in the new format before vSAN can provide the benefit of being able to perform operations on an object with the new free space requirements. A new object format health alert is displayed after an upgrade, if there are objects that must be fixed to the new object format and allows the health state to be remediated by starting a layout task to fix these objects. The health alert provides information on the

number of objects that must be fixed and the amount of data that will be rewritten. The cluster might experience a drop of about 20 percent in the performance while the relay layout task is in progress. The resync dashboard provides more accurate information about the amount of time this operation takes to complete.

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 Lifecycle Manager

vSAN generates system baselines and baseline groups that you can use with vSphere Lifecycle Manager. vSphere Lifecycle Manager in vSphere 7.0 includes the system baselines that Update Manager provided in earlier vSphere releases. It also includes new image management functionality for hosts running ESXi 7.0 and later.

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 to vSAN 7.0 also can include device driver and firmware updates. These updates support the ESXi software recommended for your cluster.

For vSAN 6.7.3 and later, you can choose to provide build recommendations for the current ESXi release only, or for the latest supported ESXi release. A build recommendation for the current release includes all patches and driver updates for the release.

In vSAN 7.0 and later, vSAN build recommendations include patch updates and applicable driver updates. To update firmware on vSAN 7.0 clusters, you must use an image through vSphere Lifecycle Manager.

vSAN System Baselines

vSAN build recommendations are provided through vSAN system baselines for vSphere Lifecycle 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 the same vendor. vSphere Lifecycle 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, vSphere Lifecycle Manager cannot recommend a system baseline.

vSphere Lifecycle 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 vSphere Lifecycle 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](#).

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

Working with vSAN Build Recommendations

vSphere Lifecycle 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 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 can recommend an upgrade to the latest release, since it is no worse than the current state.

Note vSphere Lifecycle 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 vSphere Lifecycle 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 6.7 Update 3. vSAN recommends an upgrade to release 6.7 Update 3.

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 6.7 Update 3, even though the hardware is not on the HCL for 6.7 Update 3. vSAN recommends the upgrade because the new state is no worse than the current state.

Example 4

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 6.7 Update 3 and selected baseline preference is patch-only. vSAN recommends an upgrade to 6.0 Update 3, including the necessary critical patches for the release.

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 [VMware Customer Connect](#) using a web browser 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

vSphere Lifecycle Manager is an extension service in vCenter Server 7.0 and later.

vSAN requires Internet access to update release metadata, to check the VMware Compatibility Guide, and to download ISO images from [VMware Customer Connect](#).

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

To enter **VMware Customer Connect** credentials from RVC, run the following command:

```
vsan.login_iso_depot -u <username> -p <password>
```