

vSAN Monitoring and Troubleshooting

Update 3

10 MAY 2021

VMware vSphere 6.7

VMware vSAN 6.7

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 © 2018-2021 VMware, Inc. All rights reserved. [Copyright and trademark information.](#)

Contents

	About vSAN Monitoring and Troubleshooting	5
1	Introduction to vSAN	6
2	Monitoring the vSAN Cluster	7
	Monitor vSAN Capacity	7
	Monitor Physical Devices	9
	Monitor Devices that Participate in vSAN Datastores	10
	Monitor Virtual Objects in the vSAN Cluster	10
	About vSAN Cluster Resynchronization	11
	Monitor the Resynchronization Tasks in the vSAN Cluster	12
	Throttle Resynchronization Activity in the vSAN Cluster	13
	About vSAN Cluster Rebalancing	14
	Configure Automatic Rebalance	14
	Using the vSAN Default Alarms	15
	View vSAN Default Alarms	15
	Using the VMkernel Observations for Creating Alarms	15
	Creating a vCenter Server Alarm for a vSAN Event	16
3	Monitoring vSAN Health	18
	About the vSAN Health Service	18
	Check vSAN Health	20
	Monitor vSAN from ESXi Host Client	21
	Proactive Tests	21
4	Monitoring vSAN Performance	22
	About the vSAN Performance Service	22
	Configure vSAN Performance Service	23
	Use Saved Time Range	24
	View vSAN Cluster Performance	25
	View vSAN Host Performance	26
	View vSAN VM Performance	29
	Using vSAN Performance Diagnostics	30
5	Handling Failures and Troubleshooting vSAN	32
	Uploading a vSAN Support Bundle	32
	Using Esxcli Commands with vSAN	33
	Using vsantop Command-Line Tool	36

vSAN Configuration on an ESXi Host Might Fail	36
Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly	37
vSAN Cluster Configuration Issues	37
Handling Failures in vSAN	38
Failure Handling in vSAN	38
Troubleshooting vSAN	46
Replacing Existing Hardware Components	51
Shutting Down and Restarting the vSAN Cluster	53

About vSAN Monitoring and Troubleshooting

vSAN Monitoring and Troubleshooting describes how to monitor and troubleshoot VMware vSAN® by using the vSphere Client, esxcli and RVC commands, and other tools.

Intended Audience

This manual is intended for anyone who wants to monitor vSAN operation and performance, or troubleshoot problems with a vSAN cluster. The information in this manual is written for experienced system administrators who are familiar with virtual machine technology and virtual datacenter operations. This manual assumes familiarity with VMware vSphere, including VMware ESXi, vCenter Server, and the vSphere Client.

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

For more information about vSAN features and how to configure a vSAN cluster, see *Administering VMware vSAN*.

vSphere Client and vSphere Web Client

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

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

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

Introduction to vSAN

1

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

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

Monitoring the vSAN Cluster

2

You can monitor the vSAN cluster and all the objects related to it.

You can monitor all of the objects in a vSAN environment, including hosts that participate in a vSAN cluster and the vSAN datastore. For more information about monitoring objects and storage resources in a vSAN cluster, see the *vSphere Monitoring and Performance* documentation.


This chapter includes the following topics:

- [Monitor vSAN Capacity](#)
- [Monitor Physical Devices](#)
- [Monitor Devices that Participate in vSAN Datastores](#)
- [Monitor Virtual Objects in the vSAN Cluster](#)
- [About vSAN Cluster Resynchronization](#)
- [About vSAN Cluster Rebalancing](#)
- [Using the vSAN Default Alarms](#)
- [Using the VMkernel Observations for Creating Alarms](#)

Monitor vSAN Capacity

You can monitor the capacity of the vSAN datastore, analyze usage, and view the capacity breakdown at the cluster level.

The cluster Summary page includes a summary of vSAN capacity. You also can view more detailed information in the Capacity monitor.



Procedure

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

Option	Description
vSphere Client	Under vSAN, click Capacity to view the vSAN capacity information.
vSphere Web Client	<ol style="list-style-type: none"> a Click vSAN. b Click Capacity to view the vSAN capacity information.

Results

- The Capacity Overview displays the storage capacity of the vSAN datastore, including total space, used space, free space, reserved space, and the space that is actually written or physically consumed on the vSAN disks. For clusters that have the deduplication and compression enabled, you can view the compression savings and the compression ratio.
- The Usable capacity analysis enables you to estimate the free space available based on the storage policy that you selected while keeping the deduplication ratio as 1.
- The Usage breakdown before dedup and compression displays the usage breakdown based on the categories such as VM usage, user objects, and system usage. You can view a graphical representation of the usage categories. Click the graphic to view the different usage categories.

Following are the different usage categories available:

Category	Description
VM usage	<p>Displays the following:</p> <ul style="list-style-type: none"> ■ VM home objects - Displays the VM namespace object. ■ Swap objects - Displays the VM swap files. ■ VMDK - Capacity consumed by VMDK objects that reside on the vSAN datastore that can be categorized as primary data and replica usage. Primary data includes the actual user data written into the physical disk which does not include any overhead. Replica usage displays the RAID overhead for the virtual disk. ■ VM memory snapshots - Memory snapshot file for VMs. ■ Block container volumes (attached to a VM) - Capacity consumed by the container objects that are attached to a VM.
User objects	<p>Displays iSCSI objects, block container volumes that are not attached to VM, user-created files, ISO files, and VM templates.</p>
System usage	<p>Displays the following:</p> <ul style="list-style-type: none"> ■ Performance management objects - Capacity consumed by objects created for storing performance metrics when you enable the performance service. ■ File system overhead - Overhead that the on-disk file system takes up on the capacity drives. ■ Checksum overhead - Overhead to store all the checksums. ■ Dedup & compression overhead - Overhead to get the benefits of deduplication and compression. This data is visible only if you enable deduplication and compression. ■ Transient space - Temporary space usage in a cluster.

When you enable deduplication and compression, it might take several minutes for capacity updates to be reflected in the Capacity monitor, as disk space is reclaimed and reallocated. For more information about deduplication and compression, see "Using Deduplication and Compression" in *Administering VMware vSAN*.

You can check the history of capacity usage in the vSAN datastore. Click **Capacity History**, select a time range, and click **Show Results**.

Monitor Physical Devices

You can monitor hosts, cache devices, and capacity devices used in the vSAN cluster.

Procedure

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

- 3 Click **Physical Disks** to review all hosts, cache devices, and capacity devices in the cluster. vSAN displays information about capacity devices, such as total capacity, used capacity, reserved capacity, functional status, physical location, and so on. The physical location is based on the hardware location of cache and capacity and devices on vSAN hosts.

Monitor Devices that Participate in vSAN Datastores

Verify the status of the devices that back up the vSAN datastore. You can check whether the devices experience any problems.

Procedure

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

You can view general information about the vSAN datastore, including capacity, capabilities, and the default storage policy.

- 4 Display information about local devices.
 - a Click **Disk Management** and select the disk group to display local devices in the table at the bottom of the page.
 - b Click **Capacity** to review information about the amount of capacity provisioned and used in the cluster, and also to review a breakdown of the used capacity by object type or by data type.

Monitor Virtual Objects in the vSAN Cluster

You can view the status of virtual objects in the vSAN cluster.

When one or more hosts are unable to communicate with the vSAN datastore, the information about virtual objects might not be displayed.

Procedure

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

- 3 Under **vSAN**, select **Virtual Objects** to view the corresponding virtual objects in the vSAN cluster.
 - a Select an object type in the **Affected inventory objects** area at the top of the page to display information about each object, such as health and availability, storage policy, and vSAN UUID.
 - b Select the check box on one of the virtual objects and click **View Placement details** to open the Physical Placement dialog. You can view device information, such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts.
 - c On the Physical Placement dialog, check the **Group components by host placement** check box to organize the objects by host and by disk.

Note At the cluster level, the Container Volumes filter displays detached container volumes. To view attached volumes, select the VM to which the container is attached.

About vSAN Cluster Resynchronization

You can monitor the status of virtual machine objects that are being resynchronized in the vSAN cluster.

When a hardware device, host, or network fails, or if a host is placed into maintenance mode, vSAN initiates resynchronization in the vSAN cluster. However, vSAN might briefly wait for the failed components to come back online before initiating resynchronization tasks.

The following events trigger resynchronization in the cluster:

- Editing a virtual machine (VM) storage policy. When you change VM storage policy settings, vSAN might initiate object recreation and subsequent resynchronization of the objects.

Certain policy changes might cause vSAN to create another version of an object and synchronize it with the previous version. When the synchronization is complete, the original object is discarded.

vSAN ensures that VMs continue to run, and resynchronization does not interrupt their operation. This process might require additional temporary capacity.

- Restarting a host after a failure.
- Recovering hosts from a permanent or long-term failure. If a host is unavailable for more than 60 minutes (by default), vSAN creates copies of data to recover the full policy compliance.
- Evacuating data by using the Full data migration mode before you place a host in maintenance mode.
- Exceeding the capacity threshold of a capacity device. Resynchronization is triggered when a capacity device in the vSAN cluster approaches or exceeds the threshold level of 80 percent.

If a VM is not responding due to latency caused by resynchronization, you can throttle the IOPS used for resynchronization.

Monitor the Resynchronization Tasks in the vSAN Cluster

To evaluate the status of objects that are being resynchronized, you can monitor the resynchronization tasks that are currently in progress.

Prerequisites

Verify that hosts in your vSAN cluster are running ESXi 6.5 or later.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Select the **Monitor** tab.
- 3 (Optional) On the vSphere Web Client, click **vSAN**.
- 4 Select **Resyncing objects**.
- 5 Track the progress of resynchronization of virtual machine objects.

The Object Repair Time defines the time vSAN waits before repairing a non-compliant object after placing a host in a failed state or maintenance mode. The default setting is 60 minutes. To change the setting, edit the Object Repair Timer (**Configure > vSAN > Services > Advanced Options**).

You can also view the following information about the objects that are resynchronized:

Objects	Description
Total resyncing objects	Total number of objects to be resynchronized in the vSAN cluster.
Bytes left to resync	Data (in bytes) that is remaining before the resynchronization is complete.

Objects	Description
ETA to compliance	<p>Estimated time left for the resynchronization to complete.</p> <p>The objects to be resynchronized are categorized as active, queued, and suspended. The objects that are actively synchronizing fall in the active category. The objects that are in the queue for resynchronization are the queued objects. The objects that were actively synchronizing but are now in the suspended state falls in the suspended category.</p>
Scheduled resyncing	<p>Remaining number of objects to be resynchronized.</p> <p>You can classify scheduled resyncing into two categories: scheduled and pending. The scheduled category displays the objects that are not resyncing because the delay timer has not expired.</p> <p>Resynchronization of objects starts once the timer expires. The pending category displays the objects with the expired delay timer that cannot be resynchronized.</p> <p>This can be due to insufficient resources in the current cluster or the vSAN FTT policy set on the cluster not being met.</p>

You can also view the resynchronization objects based on various filters such as **Intent** and **Status**. Using **Show first**, you can modify the view to display the number of objects.

Throttle Resynchronization Activity in the vSAN Cluster

You can reduce the number of IOPS used to perform resynchronization on disk groups in the vSAN cluster. Resynchronization throttling is a cluster-wide setting, and it is applied on a per disk group basis.

If VMs are not responding due to latency caused by resynchronization, you can throttle the number of IOPS used for resynchronization. Consider resynchronization throttling only if latencies are rising in the cluster due to resynchronization, or if resynchronization traffic is too high on a host.

Resynchronization throttling can increase the time required to complete resynchronization. Reprotection of non-compliant VMs might be delayed.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Select the **Monitor** tab.
- 3 (Optional) On the vSphere Web Client, click **vSAN**.
- 4 Select **Resyncing objects** and click **Resync Throttling**.
- 5 (Optional) Click **Show current resync traffic per host** to view resynchronization activity.
- 6 Select the **Enable throttling for resyncing objects traffic** check box.

7 Move the slider to set the throttle, as follows:

- Move the slider to the right to increase the amount of IOPS allowed for resynchronization.
- Move the slider to the left to decrease the amount of IOPS allowed for resynchronization.

A general rule is to throttle the IOPS by half and allow some time for the cluster to adapt. If further action is needed, throttle the IOPS by half again until the cluster stabilizes.

8 Click **Apply** or **OK**.

About vSAN Cluster Rebalancing

When any capacity device in your cluster reaches 80 percent full, vSAN automatically rebalances the cluster, until the space available on all capacity devices is below the threshold.

Cluster rebalancing evenly distributes resources across the cluster to maintain consistent performance and availability.

The following operations can cause disk capacity to reach 80% and initiate cluster rebalancing:

- Hardware failures occur on the cluster.
- vSAN hosts are placed in maintenance mode with the **Evacuate all data** option.
- vSAN hosts are placed in maintenance mode with **Ensure data accessibility** when objects assigned PFTT=0 reside on the host.

Note To provide enough space for maintenance and reprotection, and to minimize automatic rebalancing events in the vSAN cluster, consider keeping 30-percent capacity available at all times.

Configure Automatic Rebalance

vSAN automatically rebalances data on the disk groups by default. You can configure settings for automatic rebalancing.

Your vSAN cluster can become unbalanced due to uneven I/O patterns to devices, or when you add hosts or capacity devices. If the cluster becomes unbalanced, vSAN automatically rebalances the disks. This operation moves components from over-utilized disks to under-utilized disks.

You can enable or disable automatic rebalance, and configure the variance threshold for triggering an automatic rebalance. If any two disks in the cluster have a capacity variance that meets the rebalancing threshold, vSAN begins rebalancing the cluster.

Disk rebalancing can impact the I/O performance of your vSAN cluster. To avoid this performance impact, you can turn off automatic rebalance when peak performance is required.

Procedure

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

- 3 Under vSAN, select **Services**.
- 4 Click to edit **Advanced Options**.
- 5 Click to enable or disable Automatic Rebalance.

You can change the variance threshold to any percentage from 20 to 75.

What to do next

You can use the vSAN health service to check the disk balance. Expand the Cluster category, and select **vSAN Disk Balance**.

Using the vSAN Default Alarms

You can use the default vSAN alarms to monitor the cluster, hosts, and existing vSAN licenses.

The default alarms are automatically triggered when the events corresponding to the alarms are activated or if one or all the conditions specified in the alarms are met. You cannot edit the conditions or delete the default alarms. To configure alarms that are specific to your requirements, create custom alarms for vSAN. See [Creating a vCenter Server Alarm for a vSAN Event](#).

For information about monitoring alarms, events, and editing existing alarm settings, see the *vSphere Monitoring and Performance* documentation.

View vSAN Default Alarms

Use the default vSAN alarms to monitor your cluster, hosts, analyze any new events, and assess the overall cluster health.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click **Configure** and then click **Alarm Definitions**.
- 3 In the search box, type **vSAN** as the search term to display the alarms that are specific to vSAN.

Type vSAN Health Service Alarm to search for vSAN health service alarms.

The default vSAN alarms are displayed.

- 4 From the list of alarms, click each alarm to view the alarm definition.

Using the VMkernel Observations for Creating Alarms

VMkernel Observations (VOBs) are system events that you can use to set up vSAN alarms for monitoring and troubleshooting performance and networking issues in the vSAN cluster. In vSAN, these events are known as observations.

VMware ESXi Observation IDs for vSAN

Each VOB event is associated with an identifier (ID). Before you create a vSAN alarm in the vCenter Server, you must identify an appropriate VOB ID for the vSAN event for which you want to create an alert. You can create alerts in the VMware ESXi Observation Log file (`vobd.log`). For example, use the following VOB IDs to create alerts for any device failures in the cluster.

- `esx.problem.vob.vsan.lsom.diskerror`
- `esx.problem.vob.vsan.pdl.offline`

To review the list of VOB IDs for vSAN, open the `vobd.log` file located on your ESXi host in the `/var/log` directory. The log file contains the following VOB IDs that you can use for creating vSAN alarms.

Table 2-1. VOB IDs for vSAN

VOB ID	Description
<code>esx.audit.vsan.clustering.enabled</code>	The vSAN clustering service is enabled.
<code>esx.clear.vob.vsan.pdl.online</code>	The vSAN device has come online.
<code>esx.clear.vsan.clustering.enabled</code>	The vSAN clustering service is enabled.
<code>esx.clear.vsan.vsan.network.available</code>	vSAN has one active network configuration.
<code>esx.clear.vsan.vsan.vmknic.ready</code>	A previously reported vmknic has acquired a valid IP.
<code>esx.problem.vob.vsan.lsom.componentthreshold</code>	vSAN reaches the near node component count limit.
<code>esx.problem.vob.vsan.lsom.diskerror</code>	A vSAN device is in a permanent error state.
<code>esx.problem.vob.vsan.lsom.diskgrouplimit</code>	vSAN fails to create a disk group.
<code>esx.problem.vob.vsan.lsom.disklimit</code>	vSAN fails to add devices to a disk group.
<code>esx.problem.vob.vsan.lsom.diskunhealthy</code>	vSAN disk is unhealthy.
<code>esx.problem.vob.vsan.pdl.offline</code>	A vSAN device is offline.
<code>esx.problem.vsan.clustering.disabled</code>	vSAN clustering services are disabled.
<code>esx.problem.vsan.lsom.congestionthreshold</code>	vSAN device memory or SSD congestion has been updated.
<code>esx.problem.vsan.net.not.ready</code>	A vmknic is added to vSAN network configuration without a valid IP address. This happens when the vSAN network is not ready.
<code>esx.problem.vsan.net.redundancy.lost</code>	The vSAN network configuration does not have the required redundancy.
<code>esx.problem.vsan.no.network.connectivity</code>	vSAN does not have existing networking configuration, which is in use.
<code>esx.problem.vsan.vmknic.not.ready</code>	A vmknic is added to the vSAN network configuration without a valid IP address.

Creating a vCenter Server Alarm for a vSAN Event

You can create alarms to monitor events on the selected vSAN object, including the cluster, hosts, datastores, networks, and virtual machines.

Prerequisites

You must have the required privilege level of Alarms. `Create Alarm` or `Alarm.Modify Alarm`

Procedure

- 1 Select the vCenter Server object in the inventory that you want to monitor.
- 2 Click the **Configure** tab > **Alarm Definitions** > click the **+** icon.
- 3 Type a name and description for the new alarm.
- 4 From the **Monitor** drop-down menu, select the object on which you want to configure an alarm.
- 5 Click the **specific event occurring on this object for example VM Power On** and click **Next**.
- 6 Click **Triggers** to add a vSAN event that will trigger the alarm. The options on the Triggers page change depending on the type of activity you plan to monitor.
- 7 Click the **Add** icon (**+**).
- 8 Click in the **Event** column, and select an option from the drop-down menu.
- 9 Click in the **Status** column, and select an option from the drop-down menu.
- 10 (Optional) Configure additional conditions to be met before the alarm triggers.
 - a Click the **Add** icon to add an argument.
 - b Click in the **Argument** column, and select an option from the drop-down menu.
 - c Click in the **Operator** column, and select an option from the drop-down menu.
 - d Click in the **Value** column, and enter a value in the text field.

You can add more than one argument.
- 11 Click **Next**.

Results

You selected and configured alarm triggers.

Monitoring vSAN Health

3

You can check the overall health of the vSAN cluster, including hardware compatibility, networking configuration and operations, advanced vSAN configuration options, storage device health and virtual machine object health.

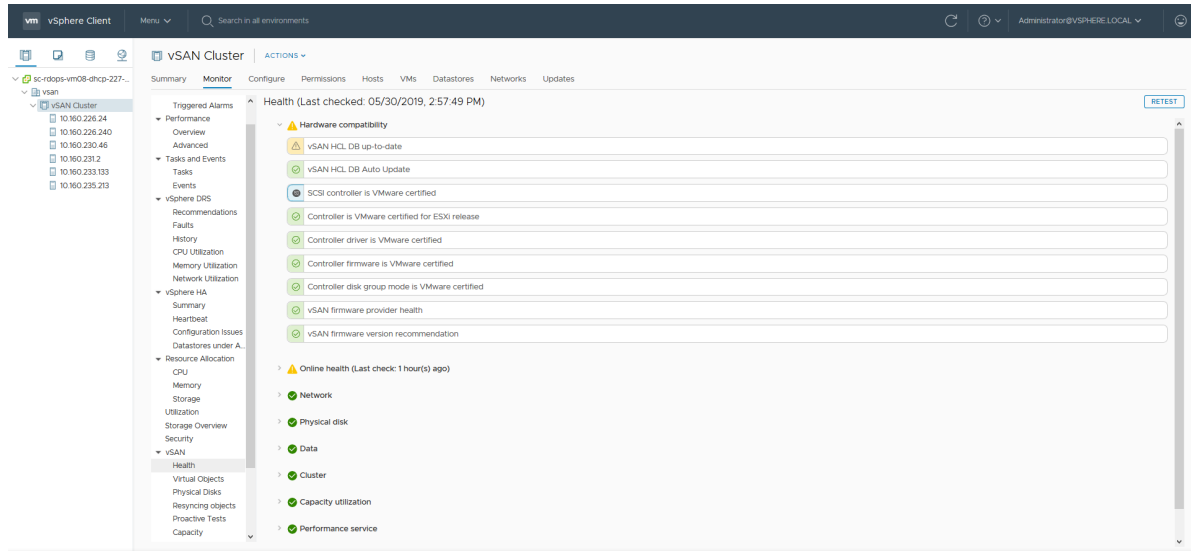
This chapter includes the following topics:

- [About the vSAN Health Service](#)
- [Check vSAN Health](#)
- [Monitor vSAN from ESXi Host Client](#)
- [Proactive Tests](#)

About the vSAN Health Service

Use the vSAN health service to monitor the health your vSAN cluster.

You can use the vSAN health checks to monitor the status of cluster components, diagnose issues, and troubleshoot problems. The health checks cover hardware compatibility, network configuration and operation, advanced vSAN configuration options, storage device health, and virtual machine objects.



The vSAN health checks are divided into categories. Each category contains individual health checks.

Table 3-1. vSAN Health Check Categories

Health Check Category	Description
Hardware Compatibility	Monitor the cluster components to ensure that they are using supported hardware, software, and drivers.
Performance Service	Monitor the health of vSAN performance service.
Network	Monitor vSAN network health.
Physical disk	Monitor the health of physical devices in the vSAN cluster.
Data	Monitor vSAN data health.
Cluster	Monitor vSAN cluster health.
Capacity utilization	Monitor vSAN cluster capacity.
Online health	Monitor vSAN cluster health and send to VMware’s analytics backend system for advanced analysis. You must participate in the Customer Experience Improvement Program to use online health checks.
vSAN Build Recommendation	Monitor vSAN build recommendations for vSphere Update Manager.
vSAN iSCSI target service	Monitor the iSCSI target service, including the network configuration and runtime status.
Encryption	Monitor vSAN encryption health.
Stretched cluster	Monitor the health of a stretched cluster, if applicable.
Hyperconverged cluster configuration compliance	Monitor the status of hosts and settings configured through the Quickstart workflow.

vSAN periodically retests each health check and updates the results. To run the health checks and update the results immediately, click the **Retest** button.

If you participate in the Customer Experience Improvement Program, you can run health checks and send the data to VMware for advanced analysis. Click the **Retest with Online health** button.

For more information about vSAN health checks, see *VMware Virtual SAN Health Check Plugin Guide*.

Monitoring vSAN Health on a Host

The ESXi host client is a browser-based interface for managing a single ESXi host. It enables you to manage the host when vCenter Server is not available. The host client provides tabs for managing and monitoring vSAN at the host level.

- The **vSAN** tab displays basic vSAN configuration.
- The **Hosts** tab displays the hosts participating in the vSAN cluster.
- The **Health** tab displays host-level health checks.

Check vSAN Health

You can view the status of vSAN health checks to verify the configuration and operation of your vSAN cluster.

Prerequisites

Procedure

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

Option	Description
vSphere Client	<p>Under vSAN, select Health to review the vSAN health check categories.</p> <ul style="list-style-type: none"> ■ If the Test Result column displays a Warning (yellow) or Failed (red) icon, expand the category to review the results of individual health checks. ■ Select an individual health check to view the detailed information. ■ In the Info section, you can click the Ask VMware button to open a knowledge base article that describes the health check and provides information about how to resolve the issue. ■ You can click Silence Alert on a health check, so it does not display any warnings or failures.
vSphere Web Client	<ol style="list-style-type: none"> a click vSAN. b Select Health to review the vSAN health check categories. <ul style="list-style-type: none"> ■ If the Test Result column displays Warning or Failed, expand the category to review the results of individual health checks. ■ Select an individual health check to view the detailed information at the bottom of the page. ■ You can click the Ask VMware button to open a knowledge base article that describes the health check and provides information about how to resolve the issue.

Monitor vSAN from ESXi Host Client

You can monitor vSAN health and basic configuration through the ESXi host client.

Procedure

- 1 Open a browser and enter the IP address of the host.
The browser redirects to the login page for the host client.
- 2 Enter the username and password for the host, and click **Login**.
- 3 In the host client navigator, click **Storage**.
- 4 In the main page, click the vSAN datastore to display the Monitor link in the navigator.
- 5 Click the tabs to view vSAN information for the host.
 - a Click the **vSAN** tab to display basic vSAN configuration.
 - b Click the **Hosts** tab to display the hosts participating in the vSAN cluster.
 - c Click the **Health** tab to display host-level health checks.
- 6 (Optional) On the **vSAN** tab, click **Edit Settings** to correct configuration issues at the host level. Select the values that match the configuration of your vSAN cluster.
Select the values that match the configuration of your vSAN cluster, and click **Save**.

Proactive Tests

You can initiate a health test on your vSAN cluster to verify that the cluster components are working as expected.

Run the VM creation test to verify the vSAN cluster health. Running the test creates a virtual machine on each host in the cluster. The test creates a VM and deletes it. If the VM creation and deletion tasks are successful, assume that the cluster components are working as expected and the cluster is functional.

Run the Network performance test to detect and diagnose connectivity issues, and to make sure the network bandwidth between the hosts supports the requirements of vSAN. The test is performed between the hosts in the cluster. It verifies that the network bandwidth between hosts, and reports a warning if the bandwidth is less than 850Mbps.

To access a proactive test, select your vSAN cluster in the vSphere Client, and click the Monitor tab. Click **vSAN > Proactive Tests**.

Monitoring vSAN Performance

4

You can monitor the performance of your vSAN cluster. Performance charts are available for clusters, hosts, physical disks, virtual machines, and virtual disks.

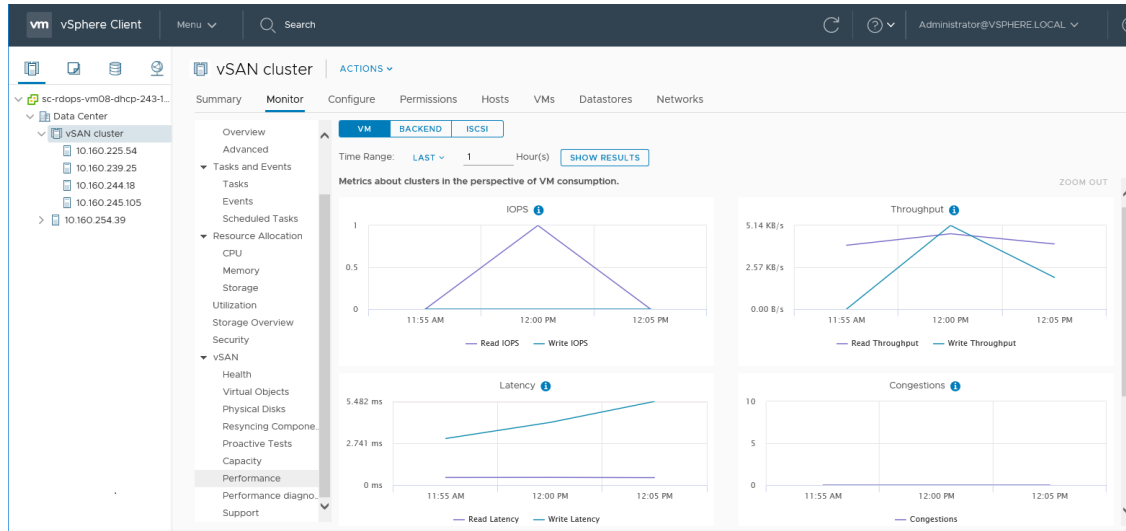
This chapter includes the following topics:

- [About the vSAN Performance Service](#)
- [Configure vSAN Performance Service](#)
- [Use Saved Time Range](#)
- [View vSAN Cluster Performance](#)
- [View vSAN Host Performance](#)
- [View vSAN VM Performance](#)
- [Using vSAN Performance Diagnostics](#)

About the vSAN Performance Service

You can use vSAN performance service to monitor the performance of your vSAN environment, and investigate potential problems.

The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.



When the vSAN performance service is turned on, the cluster summary displays an overview of vSAN performance statistics, including IOPS, throughput, and latency. You can view detailed performance statistics for the cluster, and for each host, disk group, and disk in the vSAN cluster. You also can view performance charts for virtual machines and virtual disks.

Configure vSAN Performance Service

Use the vSAN Performance Service to monitor the performance of vSAN clusters, hosts, disks, and VMs.

Note When you create a vSAN cluster in the vSphere Web Client, the Performance Service is disabled. You can enable and configure the Performance Service.

To support the Performance Service, vSAN uses a Stats database object to collect statistical data. The Stats database is a namespace object in the cluster's vSAN datastore.

Prerequisites

- All hosts in the vSAN cluster must be running ESXi 6.5 or later.
- Before you configure the vSAN Performance Service, make sure that the cluster is properly configured and has no unresolved health problems.

Procedure

- 1 Navigate to the vSAN cluster.

2 Click the **Configure** tab.

Option	Description
vSphere Client	<ul style="list-style-type: none"> a Under vSAN, select Services. b Click the Performance Service Edit button. c Click to enable vSAN Performance Service. d Select a storage policy for the Stats database object. e (Optional) Click to enable the verbose mode. This check box appears only after enabling vSAN Performance Service. When enabled, vSAN collects and saves the additional performance metrics to a Stats DB object. If you enable the verbose mode for more than 5 days, a warning message appears indicating that the verbose mode can be resource-intensive. Ensure that you do not enable it for a longer duration. f (Optional) Click to enable the network diagnostic mode. This check box appears only after enabling vSAN Performance Service. When enabled, vSAN collects and saves the additional network performance metrics to a RAM disk stats object. If you enable the network diagnostic mode for more than a day, a warning message appears indicating that the network diagnostic mode can be resource-intensive. Ensure that you do not enable it for a longer duration. g Click Apply.
vSphere Web Client	<ul style="list-style-type: none"> a Under vSAN, select Health and Performance. b Click the Performance Service Turn On button and click OK. c Click the Performance Service Edit Settings button. <ul style="list-style-type: none"> ■ Select a storage policy for the Stats database object. ■ (Optional) Click to enable the verbose mode. When enabled, vSAN collects and saves the additional performance metrics to a Stats DB object. If you enable the verbose mode for more than 5 days, a warning message appears indicating that the verbose mode can be resource-intensive. Ensure that you do not enable it for a longer duration. d Click OK.

Use Saved Time Range

You can select saved time ranges from the time range picker in performance views.

You can manually save a time range with customized name. When you run a storage performance test, the selected time range is saved automatically. You can save a time range for any of the performance views.

Prerequisites

- The vSAN performance service must be turned on.
- All hosts in the vSAN cluster must be running ESXi 6.5 or later.

Procedure

- 1 Navigate to the vSAN cluster.

- 2 Click the **Monitor** tab and click **Performance**.
- 3 Select any tab, such as **vSAN - Backend**. In the time range drop-down, select **Save**.
- 4 Enter a name for the selected time range.
- 5 Confirm your changes.

View vSAN Cluster Performance

You can use the vSAN cluster performance charts to monitor the workload in your cluster and determine the root cause of problems.

When the performance service is turned on, the cluster summary displays an overview of vSAN performance statistics, including vSAN IOPS, throughput, and latency. At the cluster level, you can view detailed statistical charts for virtual machine consumption and the vSAN back end.

Note To view iSCSI performance charts, all hosts in the vSAN cluster must be running ESXi 6.5 or later.

Prerequisites

The vSAN performance service must be turned on before you can view performance charts.

Procedure

- 1 Navigate to the vSAN cluster.

2 Click the **Monitor** tab.

Option	Description
vSphere Client	<ul style="list-style-type: none"> a Under vSAN, select Performance. b Select VM. Select a time range for your query. vSAN displays performance charts for clients running on the cluster, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster. c Select Backend. Select a time range for your query. vSAN displays performance charts for the cluster back-end operations, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster. d Select iSCSI and select an iSCSI target or LUN. Select a time range for your query. vSAN displays performance charts for iSCSI targets or LUNs, including IOPS, bandwidth, latency, and outstanding I/O.
vSphere Web Client	<ul style="list-style-type: none"> a Click Performance. b Select vSAN - Virtual Machine Consumption. Select a time range for your query. vSAN displays performance charts for clients running on the cluster, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster. c Select vSAN - Backend. Select a time range for your query. vSAN displays performance charts for the cluster back-end operations, including IOPS, throughput, latency, congestions, and outstanding I/Os. The statistics on these charts are aggregated from the hosts within the cluster. d Select vSAN - iSCSI and select an iSCSI target or LUN. Select a time range for your query. vSAN displays performance charts for iSCSI targets or LUNs, including IOPS, bandwidth, latency, and outstanding I/O.

3 Click **Refresh** or **Show Results** to update the display.

View vSAN Host Performance

You can use the vSAN host performance charts to monitor the workload on your hosts and determine the root cause of problems. You can view vSAN performance charts for hosts, disk groups, and individual storage devices.

When the performance service is turned on, the host summary displays performance statistics for each host and its attached disks. At the host level, you can view detailed statistical charts for virtual machine consumption and the vSAN back end, including IOPS, throughput, latency, and congestion. Additional charts are available to view the local client cache read IOPS and hit rate. At the disk group level, you can view statistics for the disk group. At the disk level, you can view statistics for an individual storage device.

Prerequisites

The vSAN performance service must be turned on before you can view performance charts.

To view the following performance charts, hosts in the vSAN cluster must be running ESXi 6.5 or later: Physical Adapters, VMkernel Adapters, VMkernel Adapters Aggregation, iSCSI, vSAN - Backend resync I/Os, resync IOPS, resync throughput, Disk Group resync latency.

Procedure

- 1 Navigate to the vSAN cluster, and select a host.

2 Click the **Monitor** tab.

Option	Description
vSphere Client	<ul style="list-style-type: none"> a Under vSAN, select Performance. b Select VM. Select a time range for your query. vSAN displays performance charts for clients running on the host, including IOPS, throughput, latency, congestions, and outstanding I/Os. c Select Backend. Select a time range for your query. vSAN displays performance charts for the host back-end operations, including IOPS, throughput, latency, congestions, outstanding I/Os, and resync I/Os. d Select Disks, and select a disk group. Select a time range for your query. vSAN displays performance charts for the disk group, including front end (Guest) IOPS, throughput, and latency, as well as overhead IOPS and latency. It also displays the read-cached hit rate, evictions, write-buffer free percentage, capacity and usage, cache disk destage rate, congestions, outstanding I/O, outstanding I/O size, delayed I/O percentage, delayed I/O average latency, internal queue IOPS, internal queue throughput, resync IOPS, resync throughput, and resync latency. e Select Physical Adapters, and select a NIC. Select a time range for your query. vSAN displays performance charts for the physical NIC (pNIC), including throughput, packets per second, and packets loss rate. f Select Host Network, and select a VMkernel adapter, such as vmk1. Select a time range for your query. vSAN displays performance charts for all network I/Os processed in the network adapters used by vSAN, including throughput, packets per second, and packets loss rate. g Select iSCSI. Select a time range for your query. vSAN displays performance charts for all the iSCSI services on the host, including IOPS, bandwidth, latency, and outstanding I/Os.
vSphere Web Client	<ul style="list-style-type: none"> a Click Performance. b Select vSAN - Virtual Machine Consumption. Select a time range for your query. vSAN displays performance charts for clients running on the host, including IOPS, throughput, latency, congestions, and outstanding I/Os. c Select vSAN - Backend. Select a time range for your query. vSAN displays performance charts for the host back-end operations, including IOPS, throughput, latency, congestions, outstanding I/Os, and resync I/Os. d Select vSAN - Disk Group, and select a disk group. Select a time range for your query. vSAN displays performance charts for the disk group, including front end (Guest) IOPS, throughput, and latency, as well as overhead IOPS and latency. It also displays the read-cached hit rate, evictions, write-buffer free percentage, capacity and usage, cache disk destage rate, congestions, outstanding I/O, outstanding I/O size, delayed I/O percentage, delayed I/O average latency, internal queue IOPS, internal queue throughput, resync IOPS, resync throughput, and resync latency. e Select vSAN - Disk, and select a disk. Select a time range for your query. vSAN displays performance charts for the disk, including a physical/firmware layer IOPS, throughput, and latency.

Option	Description
f	Select vSAN - Physical Adapters , and select a NIC. Select a time range for your query. vSAN displays performance charts for the physical NIC (pNIC), including throughput, packets per second, and packets loss rate.
g	Select vSAN - VMkernel Adapters , and select a VMkernel adapter, such as vmk1. Select a time range for your query. vSAN displays performance charts for all network I/Os processed in the network adapters used by vSAN, including throughput, packets per second, and packets loss rate.
h	Select vSAN - iSCSI . Select a time range for your query. vSAN displays performance charts for all the iSCSI services on the host, including IOPS, bandwidth, latency, and outstanding I/Os.

3 Click **Refresh** or **Show Results** to update the display.

View vSAN VM Performance

You can use the vSAN VM performance charts to monitor the workload on your virtual machines and virtual disks.

When the performance service is turned on, you can view detailed statistical charts for virtual machine performance and virtual disk performance. VM performance statistics cannot be collected during migration between hosts, so you might notice a gap of several minutes in the VM performance chart.

Note The performance service supports only virtual SCSI controllers for virtual disks. Virtual disks using other controllers, such as IDE, are not supported.

Prerequisites

The vSAN performance service must be turned on before you can view performance charts.

Procedure

1 Navigate to the vSAN cluster, and select a VM.

2 Click the **Monitor** tab.

Option	Description
vSphere Client	<ol style="list-style-type: none"> Under vSAN, select Performance. Select VM. Select a time range for your query. vSAN displays performance charts for the VM, including IOPS, throughput, and latency. Select Virtual Disk. Select a time range for your query. vSAN displays performance charts for the virtual disks, including IOPS, delayed normalized IOPS, virtual SCSI IOPS, virtual SCSI throughput, and virtual SCSI latency.
vSphere Web Client	<ol style="list-style-type: none"> Click Performance. Select vSAN - Virtual Machine Consumption. Select a time range for your query. vSAN displays performance charts for the VM, including IOPS, throughput, and latency. Select vSAN - Virtual Disk. Select a time range for your query. vSAN displays performance charts for the virtual disks, including IOPS, delayed normalized IOPS, virtual SCSI IOPS, virtual SCSI throughput, and virtual SCSI latency.

3 Click **Refresh** or **Show Results** to update the display.

Using vSAN Performance Diagnostics

You can use vSAN performance diagnostics to improve the performance of your vSAN cluster, and resolve performance issues.

The vSAN performance diagnostics tool analyzes previously run benchmarks gathered from the vSAN performance service. It can detect issues, suggest remediation steps, and provide supporting performance graphs for further insight.

The vSAN performance service provides the data used to analyze vSAN performance diagnostics. vSAN uses CEIP to send data to VMware for analysis.

Note Do not use vSAN performance diagnostics for general evaluation of performance on a production vSAN cluster.

Prerequisites

- The vSAN performance service must be turned on.
- vCenter Server requires Internet access to download ISO images and patches.
- You must participate in the Customer Experience Improvement Program (CEIP).

Procedure

- 1 Navigate to the vSAN cluster.

- 2 Click the **Monitor** tab.

Option	Description
vSphere Client	Under vSAN, select Performance Diagnostics .
vSphere Web Client	<ul style="list-style-type: none"> a Click vSAN. b Click Performance Diagnostics.

- 3 Select a benchmark goal from the drop-down menu.

You can select a goal based on the performance improvement that you want to achieve, such as maximum IOPS, maximum throughput, or minimum latency.

- 4 Select a time range for your query.

The default time range is the most recent hour. You can increase the range to include the last 24 hours, or define a custom time range within the last 90 days. If you used the HCIbench tool to run performance benchmark tests on the vSAN cluster, the time ranges of those tests appear in the drop-down menu.

- 5 Click **Show Results**.

Results

When you click **Show Results**, vSAN transmits performance data to the vSphere backend analytics server. After analyzing the data, the vSAN performance diagnostics tool displays a list of issues that might have affected the benchmark performance for the chosen goal.

You can click to expand each issue to view more details about each issue, such as a list of affected items. You also can click **See More** or **Ask VMware** to display a Knowledge Base article that describes recommendations to address the issue and achieve your performance goal.

Handling Failures and Troubleshooting vSAN

5

If you encounter problems when using vSAN, you can use troubleshooting topics. The topics help you understand the problem and offer you a workaround, when it is available.

This chapter includes the following topics:

- [Uploading a vSAN Support Bundle](#)
- [Using Esxcli Commands with vSAN](#)
- [Using vsantop Command-Line Tool](#)
- [vSAN Configuration on an ESXi Host Might Fail](#)
- [Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly](#)
- [vSAN Cluster Configuration Issues](#)
- [Handling Failures in vSAN](#)
- [Shutting Down and Restarting the vSAN Cluster](#)

Uploading a vSAN Support Bundle

You can upload a vSAN support bundle so VMware service personnel can analyze the diagnostic information.

VMware Technical Support routinely requests diagnostic information from your vSAN cluster when a support request is addressed. The support bundle is an archive that contains diagnostic information related to the environment, such as product specific logs, configuration files, and so on.

vSAN performs an automated upload of the support bundle, and does not allow you to review, obfuscate, or otherwise edit the contents of your support data prior to it being sent to VMware.

Note Data collected in the support bundle may be considered sensitive. If your support data contains regulated data, such as personal, health care, or financial data, you may want to avoid uploading the support bundle.

Procedure

- 1 Right-click the vSAN cluster in the vSphere Client.
- 2 Choose menu **vSAN > Upload support bundle...**
- 3 Enter your service request ID and a description of your issue.
- 4 Click **Upload**.

Using Esxcli Commands with vSAN

Use Esxcli commands to obtain information about vSAN and to troubleshoot your vSAN environment.

The following commands are available:

Command	Description
<code>esxcli vsan network list</code>	Verify which VMkernel adapters are used for vSAN communication.
<code>esxcli vsan storage list</code>	List storage disks claimed by vSAN.
<code>esxcli vsan cluster get</code>	Get vSAN cluster information.
<code>esxcli vsan health</code>	Get vSAN cluster health status.
<code>esxcli vsan debug</code>	Get vSAN cluster debug information.

The `esxcli vsan debug` commands can help you debug and troubleshoot the vSAN cluster, especially when vCenter Server is not available.

Use: `esxcli vsan debug {cmd} [cmd options]`

Debug commands:

Command	Description
<code>esxcli vsan debug disk</code>	Debug vSAN physical disks.
<code>esxcli vsan debug object</code>	Debug vSAN objects.
<code>esxcli vsan debug resync</code>	Debug vSAN resyncing objects.
<code>esxcli vsan debug controller</code>	Debug vSAN disk controllers.
<code>esxcli vsan debug limit</code>	Debug vSAN limits.
<code>esxcli vsan debug vmdk</code>	Debug vSAN VMDKs.

Example `esxcli vsan debug` commands:

```
esxcli vsan debug disk summary get
Overall Health: green
Component Metadata Health: green
Memory Pools (heaps): green
Memory Pools (slabs): green
```

```
esxcli vsan debug disk list
UUID: 52e1d1fa-af0e-0c6c-f219-e5e1d224b469
Name: mpx.vmhba1:C0:T1:L0
SSD: False
Overall Health: green
Congestion Health:
  State: green
  Congestion Value: 0
  Congestion Area: none
In Ccmds: true
In Vsi: true
Metadata Health: green
Operational Health: green
Space Health:
  State: green
  Capacity: 107365793792 bytes
  Used: 1434451968 bytes
  Reserved: 150994944 bytes
```

```
esxcli vsan debug object health summary get
```

Health Status	Number Of Objects
reduced-availability-with-no-rebuild-delay-timer	0
reduced-availability-with-active-rebuild	0
inaccessible	0
data-move	0
healthy	1
nonavailability-related-incompliance	0
nonavailability-related-reconfig	0
reduced-availability-with-no-rebuild	0

```
esxcli vsan debug object list
Object UUID: 47cbdc58-e01c-9e33-dada-020010d5dfa3
Version: 5
Health: healthy
Owner:
Policy:
  stripeWidth: 1
  CSN: 1
  spbmProfileName: vSAN Default Storage Policy
  spbmProfileId: aa6d5a82-1c88-45da-85d3-3d74b91a5bad
  forceProvisioning: 0
  cacheReservation: 0
  proportionalCapacity: [0, 100]
  spbmProfileGenerationNumber: 0
```

```
hostFailuresToTolerate: 1
```

```
Configuration:
```

```
RAID_1
```

```
Component: 47cbdc58-6928-333f-0c51-020010d5dfa3
```

```
Component State: ACTIVE, Address Space(B): 273804165120 (255.00GB),
Disk UUID: 52e95956-42cf-4d30-9cbe-763c616614d5, Disk Name: mpx.vmhba1..
Votes: 1, Capacity Used(B): 373293056 (0.35GB),
Physical Capacity Used(B): 369098752 (0.34GB), Host Name: sc-rdops...
```

```
Component: 47cbdc58-eebf-363f-cf2b-020010d5dfa3
```

```
Component State: ACTIVE, Address Space(B): 273804165120 (255.00GB),
Disk UUID: 52d11301-1720-9901-eb0a-157d68b3e4fc, Disk Name: mpx.vmh...
Votes: 1, Capacity Used(B): 373293056 (0.35GB),
Physical Capacity Used(B): 369098752 (0.34GB), Host Name: sc-rdops-vm..
```

```
Witness: 47cbdc58-21d2-383f-e45a-020010d5dfa3
```

```
Component State: ACTIVE, Address Space(B): 0 (0.00GB),
Disk UUID: 52bfd405-160b-96ba-cf42-09da8c2d7023, Disk Name: mpx.vmh...
Votes: 1, Capacity Used(B): 12582912 (0.01GB),
Physical Capacity Used(B): 4194304 (0.00GB), Host Name: sc-rdops-vm...
```

```
Type: vmnamespace
```

```
Path: /vmfs/volumes/vsan:52134fafd48ad6d6-bf03cb6af0f21b8d/New Virtual Machine
```

```
Group UUID: 00000000-0000-0000-0000-000000000000
```

```
Directory Name: New Virtual Machine
```

```
esxcli vsan debug controller list
```

```
Device Name: vmhba1
```

```
Device Display Name: LSI Logic/Symbios Logic 53c1030 PCI-X Fusion-MPT Dual Ult..
```

```
Used By VSAN: true
```

```
PCI ID: 1000/0030/15ad/1976
```

```
Driver Name: mptspi
```

```
Driver Version: 4.23.01.00-10vmw
```

```
Max Supported Queue Depth: 127
```

```
esxcli vsan debug limit get
```

```
Component Limit Health: green
```

```
Max Components: 750
```

```
Free Components: 748
```

```
Disk Free Space Health: green
```

```
Lowest Free Disk Space: 99 %
```

```
Used Disk Space: 1807745024 bytes
```

```
Used Disk Space (GB): 1.68 GB
```

```
Total Disk Space: 107365793792 bytes
```

```
Total Disk Space (GB): 99.99 GB
```

```
Read Cache Free Reservation Health: green
```

```
Reserved Read Cache Size: 0 bytes
Reserved Read Cache Size (GB): 0.00 GB
Total Read Cache Size: 0 bytes
Total Read Cache Size (GB): 0.00 GB
```

```
esxcli vsan debug vmdk list
Object: 50cbdc58-506f-c4c2-0bde-020010d5dfa3
Health: healthy
Type: vdisk
Path: /vmfs/volumes/vsan:52134fafd48ad6d6-bf03cb6af0f21b8d/47cbdc58-e01c-9e33-
dada-020010d5dfa3/New Virtual Machine.vmdk
Directory Name: N/A
```

```
esxcli vsan debug resync list
```

Object	Component	Bytes Left To Resync	GB Left To Resync
31cfdc58-e68d...	Component:23d1dc58...	536870912	0.50
31cfdc58-e68d...	Component:23d1dc58...	1073741824	1.00
31cfdc58-e68d...	Component:23d1dc58...	1073741824	1.00

Using vsantop Command-Line Tool

Use the command-line tool - vsantop - that runs on ESXi hosts to view the real time vSAN performance metrics. You can use this tool to monitor vSAN performance.

To display the different performance views and metrics in vsantop, enter the following commands:

Command	Description
^L	Redraw screen
Space	Update display
h or ?	Help; show this text
q	Quit
f/F	Add or remove fields
o/O	Change the order of displayed fields
s	Set the delay in seconds between updates
#	Set the number of instances to display
E	Change the selected entity type
L	Change the length of the field
l	Limit display to specific node id
.	Sort by column, same number twice to change sort order

vSAN Configuration on an ESXi Host Might Fail

In certain circumstances, the task of configuring vSAN on a particular host might fail.

Problem

An ESXi host that joins a vSAN cluster fails to have vSAN configured.

Cause

If a host does not meet hardware requirements or experiences other problems, vSAN might fail to configure the host. For example, insufficient memory on the host might prevent vSAN from being configured.

Solution

- 1 Place the host that causes the failure in Maintenance Mode.
- 2 Move the host out of the vSAN cluster.
- 3 Resolve the problem that prevents the host to have vSAN configured.
- 4 Exit Maintenance Mode.
- 5 Move the host back into the vSAN cluster.

Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly

When you use the **Check Compliance** button, a virtual machine object does not change its status from Not Compliant to Compliant even though vSAN resources have become available and satisfy the virtual machine profile.

Problem

When you use force provisioning, you can provision a virtual machine object even when the policy specified in the virtual machine profile cannot be satisfied with the resources available in the vSAN cluster. The object is created, but remains in the non-compliant status.

vSAN is expected to bring the object into compliance when storage resources in the cluster become available, for example, when you add a host. However, the object's status does not change to compliant immediately after you add resources.

Cause

This occurs because vSAN regulates the pace of the reconfiguration to avoid overloading the system. The amount of time it takes for compliance to be achieved depends on the number of objects in the cluster, the I/O load on the cluster and the size of the object in question. In most cases, compliance is achieved within a reasonable time.

vSAN Cluster Configuration Issues

After you change the vSAN configuration, vCenter Server performs validation checks for vSAN configuration. Validation checks are also performed as a part of a host synchronization process. If vCenter Server detects any configuration problems, it displays error messages.

Problem

Error messages indicate that vCenter Server has detected a problem with vSAN configuration.

Solution

Use the following methods to fix vSAN configuration problems.

Table 5-1. vSAN Configuration Errors and Solutions

vSAN Configuration Error	Solution
Host with the vSAN service enabled is not in the vCenter cluster	Add the host to the vSAN cluster. 1 Right-click the host, and select Move To . 2 Select the vSAN cluster and click OK .
Host is in a vSAN enabled cluster but does not have vSAN service enabled	Verify whether vSAN network is properly configured and enabled on the host. See <i>vSAN Planning and Deployment</i> .
vSAN network is not configured	Configure vSAN network. See <i>vSAN Planning and Deployment</i> .
Host cannot communicate with all other nodes in the vSAN enabled cluster	Might be caused by network isolation. See the <i>vSAN Planning and Deployment</i> documentation.
Found another host participating in the vSAN service which is not a member of this host's vCenter cluster.	Make sure that the vSAN cluster configuration is correct and all vSAN hosts are in the same subnet. See <i>vSAN Planning and Deployment</i> .

Handling Failures in vSAN

vSAN handles failures of the storage devices, hosts and network in the cluster according to the severity of the failure. You can diagnose problems in vSAN by observing the performance of the vSAN datastore and network.

Failure Handling in vSAN

vSAN implements mechanisms for indicating failures and rebuilding unavailable data for data protection.

Failure States of vSAN Components

In vSAN, components that have failed can be in absent or degraded state. According to the component state, vSAN uses different approaches for recovering virtual machine data.

vSAN also provides alerts about the type of component failure. See [Using the VMkernel Observations for Creating Alarms](#) and [Using the vSAN Default Alarms](#).

vSAN supports two types of failure states for components:

Table 5-2. Failure States of Components in vSAN

Component Failure State	Description	Recovery	Cause
Degraded	A component is in degraded state if vSAN detects a permanent component failure and assumes that the component is not going to recover to working state.	vSAN starts rebuilding the affected components immediately.	<ul style="list-style-type: none"> ■ Failure of a flash caching device ■ Magnetic or flash capacity device failure ■ Storage controller failure
Absent	A component is in absent state if vSAN detects a temporary component failure where the component might recover and restore its working state.	vSAN starts rebuilding absent components if they are not available within a certain time interval. By default, vSAN starts rebuilding absent components after 60 minutes.	<ul style="list-style-type: none"> ■ Lost network connectivity ■ Failure of a physical network adapter ■ ESXi host failure ■ Unplugged flash caching device ■ Unplugged magnetic disk or flash capacity device

Examine the Failure State of a Component

You can determine whether a component is in the absent or degraded failure state.

If a failure occurs in the cluster, vSAN marks the components for an object as absent or degraded based on the failure severity.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Monitor** tab, click **vSAN** and select **Virtual Objects**.

The home directories and virtual disks of the virtual machines in the cluster appear.

- 3 Select the check box on one of the virtual objects and click **View Placement details** to open the Physical Placement dialog. You can view device information, such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts.

If a failure has occurred in the vSAN cluster, the Placement and Availability is equal to Absent or Degraded.

Object States That Indicate Problems in vSAN

Examine the compliance status and the operational state of a virtual machine object to determine how a failure in the cluster affects the virtual machine.

Table 5-3. Object State

Object State Type	Description
Compliance Status	The compliance status of a virtual machine object indicates whether it meets the requirements of the assigned VM storage policy.
Operational State	<p>The operational state of an object can be healthy or unhealthy. It indicates the type and number of failures in the cluster.</p> <p>An object is healthy if an intact replica is available and more than 50 percent of the object's votes are still available.</p> <p>An object is unhealthy if an entire replica is not available or less than 50 percent of the object's votes are unavailable. For example, an object might become unhealthy if a network failure occurs in the cluster and a host becomes isolated.</p>

To determine the overall influence of a failure on a virtual machine, examine the compliance status and the operational state. If the operational state remains healthy although the object is noncompliant, the virtual machine can continue using the vSAN datastore. If the operational state is unhealthy, the virtual machine cannot use the datastore.

Examine the Health of an Object in vSAN

Use the vSphere Web Client to examine whether a virtual machine is healthy. A virtual machine is considered as healthy when a replica of the VM object and more than 50 percent of the votes for an object are available.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Monitor** tab, click **vSAN** and select **Virtual Objects**.
The home directories and virtual disks of the virtual machines in the cluster appear.
- 3 Select an object type in the **Affected inventory objects** area at the top of the page to display information about each object, such as health and availability, storage policy, and vSAN UUID.
If the inventory object is Unhealthy, the vSphere Web Client indicates the reason for the unhealthy state in brackets.

Examine the Compliance of a Virtual Machine in vSAN

Use the vSphere Web Client to examine whether a virtual machine object is compliant with the assigned VM storage policy.

Procedure

- 1 Examine the compliance status of a virtual machine.
 - a Browse to the virtual machine in the vSphere Web Client navigator.
 - b On the **Summary** tab, examine the value of the VM Storage Policy Compliance property under VM Storage Policies.

- 2 Examine the compliance status of the objects of the virtual machine.
 - a Navigate to the vSAN cluster.
 - b On the **Monitor** tab, click **vSAN** and select **Virtual Objects**.
 - c Select an object type in the **Affected inventory objects** area at the top of the page to display information about each object, such as health and availability, storage policy, and vSAN UUID.
 - d Select the check box on one of the virtual objects and click **View Placement details** to open the Physical Placement dialog. You can view device information, such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts.
 - e On the Physical Placement dialog, check the **Group components by host placement** check box to organize the objects by host and by disk.

Accessibility of Virtual Machines Upon a Failure in vSAN

If a virtual machine uses vSAN storage, its storage accessibility might change according to the type of failure in the vSAN cluster.

Changes in the accessibility occur when the cluster experiences more failures than the policy for a virtual machine object tolerates.

As a result from a failure in the vSAN cluster, a virtual machine object might become inaccessible. An object is inaccessible if a full replica of the object is not available because the failure affects all replicas, or when less than 50 percent of the object's votes are available.

According to the type of object that is inaccessible, virtual machines behave in the following ways:

Table 5-4. Inaccessibility of Virtual Machine Objects

Object Type	Virtual Machine State	Virtual Machine Symptoms
VM Home Namespace	<ul style="list-style-type: none"> ■ Inaccessible ■ Orphaned if vCenter Server or the ESXi host cannot access the .vmx file of the virtual machine. 	The virtual machine process might crash and the virtual machine might be powered off.
VMDK	Inaccessible	The virtual machine remains powered on but the I/O operations on the VMDK are not being performed. After a certain timeout passes, the guest operating system ends the operations.

Virtual machine inaccessibility is not a permanent state. After the underlying issue is resolved, and a full replica and more than 50 percent of the object's votes are restored, the virtual machine automatically becomes accessible again.

Storage Device is Failing in a vSAN Cluster

vSAN monitors the performance of each storage device and proactively isolates unhealthy devices. It detects gradual failure of a storage device and isolates the device before congestion builds up within the affected host and the entire vSAN cluster.

If a disk experiences sustained high latencies or congestion, vSAN considers the device as a dying disk, and evacuates data from the disk. vSAN handles the dying disk by evacuating or rebuilding data. No user action is required, unless the cluster lacks resources or has inaccessible objects.

Component Failure State and Accessibility

The vSAN components that reside on the magnetic disk or flash capacity device are marked as absent.

Behavior of vSAN

vSAN responds to the storage device failure in the following ways.

Parameter	Behavior
Alarms	An alarm is generated from each host whenever an unhealthy device is diagnosed. A warning is issued whenever a disk is suspected of being unhealthy.
Health check	The Disk operation health check issues a warning for the dying disk.
Health status	On the Disk Management page, the health status of the dying disk is listed as Unhealthy . When vSAN completes evacuation of data, the health status is listed as DyingDiskEmpty .
Rebuilding data	vSAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, vSAN starts the recovery process immediately because the components are marked as degraded. If resources are available, vSAN automatically reprotects the data.

If vSAN detects a disk with a permanent error, it makes a limited number of attempts to revive the disk by unmounting and mounting it.

Capacity Device Not Accessible in a vSAN Cluster

When a magnetic disk or flash capacity device fails, vSAN evaluates the accessibility of the objects on the device and rebuilds them on another host if space is available and the **Primary level of failures to tolerate** is set to 1 or more.

Component Failure State and Accessibility

The vSAN components that reside on the magnetic disk or flash capacity device are marked as degraded.

Behavior of vSAN

vSAN responds to the capacity device failure in the following ways.

Parameter	Behavior
Primary level of failures to tolerate	<p>If the Primary level of failures to tolerate in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, vSAN starts an automatic re-protection.</p> <p>If the Primary level of failures to tolerate is set to 0, a virtual machine object is inaccessible if one of the object's components resides on the failed capacity device.</p> <p>Restore the virtual machine from a backup.</p>
I/O operations on the capacity device	<p>vSAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.</p> <p>If vSAN determines that the object is available, all running I/O operations are resumed.</p>
Rebuilding data	<p>vSAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, vSAN starts the recovery process immediately because the components are marked as degraded.</p> <p>If resources are available, an automatic re-protect will occur.</p>

A Flash Caching Device Is Not Accessible in a vSAN Cluster

When a flash caching device fails, vSAN evaluates the accessibility of the objects on the disk group that contains the cache device, and rebuilds them on another host if possible and the **Primary level of failures to tolerate** is set to 1 or more.

Component Failure State and Accessibility

Both cache device and capacity devices that reside in the disk group, for example, magnetic disks, are marked as degraded. vSAN interprets the failure of a single flash caching device as a failure of the entire disk group.

Behavior of vSAN

vSAN responds to the failure of a flash caching device in the following way:

Parameter	Behavior
Primary level of failures to tolerate	<p>If the Primary level of failures to tolerate in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, vSAN starts an automatic re-protection.</p> <p>If the Primary level of failures to tolerate is set to 0, a virtual machine object is inaccessible if one of the object's components is on the failed disk group.</p>
I/O operations on the disk group	<p>vSAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.</p> <p>If vSAN determines that the object is available, all running I/O operations are resumed.</p>
Rebuilding data	<p>vSAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, vSAN starts the recovery process immediately because the components are marked as degraded.</p>

A Host Is Not Responding in a vSAN Cluster

If a host stops responding because of a failure or reboot of the host, vSAN waits for the host to recover before vSAN rebuilds the components on the host elsewhere in the cluster.

Component Failure State and Accessibility

The vSAN components that reside on the host are marked as absent.

Behavior of vSAN

vSAN responds to the host failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	<p>If the Primary level of failures to tolerate in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, vSAN starts an automatic re-protection.</p> <p>If the Primary level of failures to tolerate is set to 0, a virtual machine object is inaccessible if the object's components reside on the failed host.</p>
I/O operations on the host	<p>vSAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.</p> <p>If vSAN determines that the object is available, all running I/O operations are resumed.</p>
Rebuilding data	<p>If the host does not rejoin the cluster within 60 minutes, vSAN examines whether some of the other hosts in the cluster can satisfy the requirements for cache, space and placement rules for the objects on the inaccessible host. If such a host is available, vSAN starts the recovery process.</p> <p>If the host rejoins the cluster after 60 minutes and recovery has started, vSAN evaluates whether to continue the recovery or stop it and resynchronize the original components.</p>

Network Connectivity Is Lost in the vSAN Cluster

When the connectivity between the hosts in the cluster is lost, vSAN determines the active partition and rebuilds the components from the isolated partition on the active partition if the connectivity is not restored.

Component Failure State and Accessibility

vSAN determines the partition where more than 50 percent of the votes of an object are available. The components on the isolated hosts are marked as absent.

Behavior of vSAN

vSAN responds to a network failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	<p>If the Primary level of failures to tolerate in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, vSAN starts an automatic reprotection.</p> <p>If the Primary level of failures to tolerate is set to 0, a virtual machine object is inaccessible if the object's components are on the isolated hosts.</p>
I/O operations on the isolated hosts	<p>vSAN stops all running I/O operations for 5-7 seconds until it re-evaluates whether an object is still available without the failed component.</p> <p>If vSAN determines that the object is available, all running I/O operations are resumed.</p>
Rebuilding data	<p>If the host rejoins the cluster within 60 minutes, vSAN synchronizes the components on the host.</p> <p>If the host does not rejoin the cluster within 60 minutes, vSAN examines whether some of the other hosts in the cluster can satisfy the requirements for cache, space and placement rules for the objects on the inaccessible host. If such a host is available, vSAN starts the recovery process.</p> <p>If the host rejoins the cluster after 60 minutes and recovery has started, vSAN evaluates whether to continue the recovery or stop it and resynchronize the original components.</p>

A Storage Controller Fails in a vSAN Cluster

When a storage controller fails, vSAN evaluates the accessibility of the objects on the disk groups that are attached to the controller and rebuilds them on another host.

Symptoms

If a host contains a single storage controller and multiple disk groups, and all devices in all disk groups are failed, then you might assume that a failure in the common storage controller is the root cause. Examine the VMkernel log messages to determine the nature of the fault.

Component Failure State and Accessibility

When a storage controller fails, the components on the flash caching devices and capacity devices in all disk groups that are connected to the controller are marked as degraded.

If a host contains multiple controllers, and only the devices that are attached to an individual controller are inaccessible, then you might assume that this controller has failed.

Behavior of vSAN

vSAN responds to a storage controller failure in the following way:

Parameter	Behavior
Primary level of failures to tolerate	<p>If the Primary level of failures to tolerate in the VM storage policy is equal to or greater than 1, the virtual machine objects are still accessible from another ESXi host in the cluster. If resources are available, vSAN starts an automatic reprotection.</p> <p>If the Primary level of failures to tolerate is set to 0, a virtual machine object is inaccessible if the object's components reside on the disk groups that are connected to the storage controller.</p>
Rebuilding data	<p>vSAN examines whether the hosts and the capacity devices can satisfy the requirements for space and placement rules for the objects on the failed device or disk group. If such a host with capacity is available, vSAN starts the recovery process immediately because the components are marked as degraded.</p>

Stretched Cluster Site Fails or Loses Network Connection

A vSAN stretched cluster manages failures that occur due to the loss of a network connection between sites or the temporary loss of one site.

Stretched Cluster Failure Handling

In most cases, the stretched cluster continues to operate during a failure and automatically recovers after the failure is resolved.

Table 5-5. How Stretched Cluster Handles Failures

Type of Failure	Behavior
Network Connection Lost Between Active Sites	If the network connection fails between the two active sites, the witness host and the preferred site continue to service storage operations, and keep data available. When the network connection returns, the two active sites are resynchronized.
Secondary Site Fails or Loses Network Connection	If the secondary site goes offline or becomes isolated from the preferred site and the witness host, the witness host and the preferred site continue to service storage operations, and keep data available. When the secondary site returns to the cluster, the two active sites are resynchronized.
Preferred Site Fails or Loses Network Connection	If the preferred site goes offline or becomes isolated from the secondary site and the witness host, the secondary site continues storage operations if it remains connected to the witness host. When the preferred site returns to the cluster, the two active sites are resynchronized.
Witness Host Fails or Loses Network Connection	If the witness host goes offline or becomes isolated from the preferred site or the secondary site, objects become noncompliant but data remains available. VMs that are currently running are not affected.

Troubleshooting vSAN

Examine the performance and accessibility of virtual machines to diagnose problems in the vSAN cluster.

Verify Drivers, Firmware, Storage I/O Controllers Against the *VMware Compatibility Guide*

Use the vSAN Health Service to verify whether your hardware components, drivers, and firmware are compatible with vSAN.

Using hardware components, drivers, and firmware that are not compatible with vSAN might cause problems in the operation of the vSAN cluster and the virtual machines running on it.

The hardware compatibility health checks verify your hardware against the *VMware Compatibility Guide*. For more information about using the vSAN health service, see [Chapter 3 Monitoring vSAN Health](#).

Examining Performance in a vSAN Cluster

Monitor the performance of virtual machines, hosts, and the vSAN datastore to identify potential storage problems.

Monitor regularly the following performance indicators to identify faults in vSAN storage, for example, by using the performance charts in the vSphere Web Client:

- Datastore. Rate of I/O operations on the aggregated datastore.
- Virtual Machine. I/O operations, memory and CPU usage, network throughput and bandwidth.

You can use the vSAN performance service to access detailed performance charts. For information about using the performance service, see [Chapter 4 Monitoring vSAN Performance](#). For more information about using performance data in a vSAN cluster, see the *vSAN Troubleshooting Reference Manual*.

Network Misconfiguration Status in a vSAN Cluster

After you enable vSAN on a cluster, the datastore is not assembled correctly because of a detected network misconfiguration.

Problem

After you enable vSAN on a cluster, on the **Summary** tab for the cluster the Network Status for vSAN appears as `Misconfiguration detected`.

Cause

One or more members of the cluster cannot communicate because of either of the following reasons:

- A host in the cluster does not have a VMkernel adapter for vSAN.
- The hosts cannot connect each other in the network.

Solution

Join the members of the cluster to the same network. See *vSAN Planning and Deployment*.

Virtual Machine Appears as Noncompliant, Inaccessible or Orphaned in vSAN

The state of a virtual machine that stores data on a vSAN datastore appears as noncompliant, inaccessible or orphaned because of failures in the vSAN cluster.

Problem

A virtual machine on a vSAN datastore is in one of the following states that indicate a fault in the vSAN cluster.

- The virtual machine is non-compliant and the compliance status of some of its object is noncompliant. See [Examine the Compliance of a Virtual Machine in vSAN](#).
- The virtual machine object is inaccessible or orphaned. See [Examine the Failure State of a Component](#).

If an object replica is still available on another host, vSAN forwards the I/O operations of the virtual machine to the replica.

Cause

If the object of the virtual machine can no longer satisfy the requirement of the assigned VM storage policy, vSAN considers it noncompliant. For example, a host might temporarily lose connectivity. See [Object States That Indicate Problems in vSAN](#).

If vSAN cannot locate a full replica or more than 50 percent of the votes for the object, the virtual machine becomes inaccessible. If a vSAN detects that the `.vmx` file is not accessible because the VM Home Namespace is corrupted, the virtual machine becomes orphaned. See [Accessibility of Virtual Machines Upon a Failure in vSAN](#).

Solution

If the cluster contains enough resources, vSAN automatically recovers the corrupted objects if the failure is permanent.

If the cluster does not have enough resources to rebuild the corrupted objects, extend the space in the cluster. See *Administering VMware vSAN*.

Attempt to Create a Virtual Machine on vSAN Fails

When you try to deploy a virtual machine in a vSAN cluster, the operation fails with an error that the virtual machine files cannot be created.

Problem

The operation for creating a virtual machine fails with an error status: `Cannot complete file creation operation`.

Cause

The deployment of a virtual machine on vSAN might fail for several reasons.

- vSAN cannot allocate space for the virtual machine storage policies and virtual machine objects. Such a failure might occur if the datastore does not have enough usable capacity, for example, if a physical disk is temporarily disconnected from the host.
- The virtual machine has very large virtual disks and the hosts in the cluster cannot provide storage for them based on the placement rules in the VM storage policy

For example, if the **Primary level of failures to tolerate** in the VM storage policy is set to 1, vSAN must store two replicas of a virtual disk in the cluster, each replica on a different host. The datastore might have this space after aggregating the free space on all hosts in the cluster. However, no two hosts can be available in the cluster, each providing enough space to store a separate replica of the virtual disk.

vSAN does not move components between hosts or disks groups to free space for a new replica, even though the cluster might contain enough space for provisioning the new virtual machine.

Solution

- ◆ Verify the state of the capacity devices in the cluster.
 - a Navigate to the vSAN cluster.
 - b On the **Monitor** tab, click **vSAN** and select **Physical Disks**.
 - c Examine the capacity and health status of the devices on the hosts in the cluster.

Stretched Cluster Configuration Error When Adding a Host

Before adding new hosts to a stretched cluster, all current hosts must be connected. If a current host is disconnected, the configuration of the new host is incomplete.

Problem

After you add a host to a stretched cluster in which some hosts are disconnected, on the Summary tab for the cluster the Configuration Status for vSAN appears as `Unicast agent unset on host`.

Cause

When a new host joins a stretched cluster, vSAN must update the configuration on all hosts in the cluster. If one or more hosts are disconnected from the vCenter Server, the update fails. The new host successfully joins the cluster, but its configuration is incomplete.

Solution

Verify that all hosts are connected to vCenter Server, and click the link provided in the Configuration Status message to update the configuration of the new host.

If you cannot rejoin the disconnected host, remove the disconnected host from the cluster, and click the link provided in the Configuration Status message to update the configuration of the new host.

Stretched Cluster Configuration Error When Using RVC to Add a Host

If you use the RVC tool to add a host to a stretched cluster, the configuration of the new host is incomplete.

Problem

After you use the RVC tool to add a host to a stretched cluster, on the Summary tab for the cluster the Configuration Status for vSAN appears as `Unicast agent unset on host`.

Cause

When a new host joins a stretched cluster, vSAN must update the configuration on all hosts in the cluster. If you use the RVC tool to add the host, the update does not occur. The new host successfully joins the cluster, but its configuration is incomplete.

Solution

Verify that all hosts are connected to vCenter Server, and click the link provided in the Configuration Status message to update the configuration of the new host.

Cannot Add or Remove the Witness Host in a Stretched Cluster

Before adding or removing the witness host in a stretched cluster, all current hosts must be connected. If a current host is disconnected, you cannot add or remove the witness host.

Problem

When you add or remove a witness host in a stretched cluster in which some hosts are disconnected, the operation fails with an error status: `The operation is not allowed in the current state. Not all hosts in the cluster are connected to Virtual Center.`

Cause

When the witness host joins or leaves a stretched cluster, vSAN must update the configuration on all hosts in the cluster. If one or more hosts are disconnected from the vCenter Server, the witness host cannot be added or removed.

Solution

Verify all hosts are connected to vCenter Server, and retry the operation. If you cannot rejoin the disconnected host, remove the disconnected host from the cluster, and then you can add or remove the witness host.

Disk Group Becomes Locked

In an encrypted vSAN cluster, when communication between a host and the KMS is lost, the disk group can become locked if the host reboots.

Problem

vSAN locks a host's disk groups when the host reboots and it cannot get the KEK from the KMS. The disks behave as if they are unmounted. Objects on the disks become inaccessible.

You can view a disk group's health status on the Disk Management page in the vSphere Web Client. An Encryption health check warning notifies you that a disk is locked.

Cause

Hosts in an encrypted vSAN cluster do not store the KEK on disk. If a host reboots and cannot get the KEK from the KMS, vSAN locks the host's disk groups.

Solution

To exit the locked state, you must restore communication with the KMS and reestablish the trust relationship.

Replacing Existing Hardware Components

Under certain conditions, you must replace hardware components, drivers, firmware, and storage I/O controllers in the vSAN cluster.

In vSAN, you should replace hardware devices when you encounter failures or if you must upgrade your cluster.

Replace a Flash Caching Device on a Host

You should replace a flash caching device if you detect a failure or when you must upgrade it. Before you physically unplug a flash device from the host, you must manually remove the device from vSAN.

Caution If you decommission the flash caching device without removing it from vSAN first, vSAN uses smaller amount of cache than expected. As a result, the cluster performance becomes degraded.

When you replace a flash caching device, the virtual machines on the disk group become inaccessible and the components on the group are marked as degraded. See [A Flash Caching Device Is Not Accessible in a vSAN Cluster](#).

Prerequisites

- Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.
- If you upgrade the flash caching device, verify the following requirements:
 - If you upgrade the flash caching device, verify that the cluster contains enough space to migrate the data from the disk group that is associated with the flash device.
 - Place the host in maintenance mode.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Configure** tab, click **Disk Management** under vSAN.
- 3 Select the disk group that contains the device that you want to replace.
- 4 Select the flash caching device and click **Remove selected disk(s) from disk group**.

Results

After the flash caching device is deleted from the vSAN cluster, the cluster details reflect the current cluster capacity and configuration settings. vSAN discards the disk group memberships, deletes partitions, and removes stale data from all devices.

What to do next

- 1 Add a new device to the host.
The host automatically detects the device.
- 2 If the host is unable to detect the device, perform a device rescan.

Replace a Capacity Device

You should replace a flash capacity device or a magnetic disk if you detect a failure or when you upgrade it. Before you physically remove the device from the host, you must manually delete the device from vSAN.

When you unplug a capacity device without removing it from the vSAN cluster, the components on the disk are marked as absent. If the capacity device fails, the components on the disk are marked as degraded. When the number of failures of the object replica with the affected components exceeds the FTT value, the virtual machines on the disk become inaccessible. See [Capacity Device Not Accessible in a vSAN Cluster](#).

Note If your vSAN cluster uses deduplication and compression, you must remove the entire disk group from the cluster before you replace the device.

Prerequisites

- Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.

- If you upgrade the capacity device, verify the following requirements:
 - Verify that the cluster contains enough space to migrate the data from the capacity device.
 - Place the host in maintenance mode.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Configure** tab, click **Disk Management** under vSAN.
- 3 Select the disk group that contains the device that you want to replace.
- 4 Select the flash capacity device or magnetic disk, and click **Remove selected disk(s) from disk group**.

What to do next

- 1 Add a new device to the host.
The host automatically detects the device.

- 2 If the host is unable to detect the device, perform a device rescan.

Remove a Device from a Host by Using an ESXCLI Command

If you detect a failed storage device or if you upgrade a device, you can manually remove it from a host by using an ESXCLI command.

If you remove a flash caching device, vSAN deletes the disk group that is associated with the flash device and all its member devices.

Prerequisites

Verify that the storage controllers on the hosts are configured in passthrough mode and support the hot-plug feature.

If the storage controllers are configured in RAID 0 mode, see the vendor documentation for information about adding and removing devices.

Procedure

- 1 Open an SSH connection to the ESXi host.
- 2 To identify the device ID of the failed device, run this command and learn the device ID from the output.

```
esxcli vsan storage list
```

- 3 To remove the device from vSAN, run this command.

```
esxcli vsan storage remove -d device_id
```

What to do next

- 1 Add a new device to the host.
The host automatically detects the device.
- 2 If the host is unable to detect the device, perform a device rescan.

Shutting Down and Restarting the vSAN Cluster

When necessary, you can shut down the entire vSAN cluster.

If you plan to shut down the vSAN cluster, you do not need to manually disable vSAN on the cluster.

Procedure

- 1 Power off all virtual machines (VMs) running in the vSAN cluster, unless vCenter Server is running on the cluster.

If vCenter Server is hosted in the vSAN cluster, do not power off the vCenter Server VM.

- 2 Verify that all resynchronization tasks are complete.

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

- 3 If vCenter Server is hosted in the vSAN cluster, power off the vCenter Server VM.

The vSphere Client becomes unavailable.

- 4 Place the ESXi hosts in maintenance mode.

Connect to the host and use the host client or **esxcli** command to place the host into maintenance mode with no data migration (no action).

- 5 Power off the hosts after they have successfully entered maintenance mode.

- 6 Power on the ESXi hosts.

- a On the physical box where ESXi is installed, press the power button until the power-on sequence begins.

The ESXi host starts, locates its VMs, and functions normally.

After you power on the hosts, the vSAN cluster is automatically recreated.

You can ignore any misconfiguration status message if you did not make network configuration changes and the vSAN cluster was working as expected before you shut down the cluster. The message disappears after at least three hosts join the cluster.

- 7 Take the hosts out of maintenance mode.

- 8 Restart the VMs.