

vSAN Monitoring and Troubleshooting

Update 3

VMware vSphere 8.0

VMware vSAN 8.0

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

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

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

Copyright © 2018-2024 Broadcom. All Rights Reserved. The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, go to <https://www.broadcom.com>. All trademarks, trade names, service marks, and logos referenced herein belong to their respective companies. [Copyright and trademark information](#).

Contents

About vSAN Monitoring and Troubleshooting	5
1 What Is vSAN	6
2 Monitoring the vSAN Cluster	7
Monitor vSAN Capacity	7
Monitor Physical Devices in vSAN Cluster	12
Monitor Devices that Participate in vSAN Datastores	12
Monitor Virtual Objects in vSAN Cluster	13
Monitor Container Volumes in vSAN Cluster	14
About Reserved Capacity in vSAN Cluster	14
Configure Reserved Capacity for vSAN Cluster	16
About vSAN Cluster Resynchronization	18
Monitor the Resynchronization Tasks in vSAN Cluster	18
About vSAN Cluster Rebalancing	19
Configure Automatic Rebalance in vSAN Cluster	20
Using the vSAN Default Alarms	21
View vSAN Default Alarms	22
View vSAN Network Alarms	22
Using the VMkernel Observations for Creating vSAN Alarms	23
Creating a vCenter Server Alarm for a vSAN Event	24
3 Monitoring vSAN Skyline Health	26
About the vSAN Skyline Health	26
Check vSAN Skyline Health	28
Monitor vSAN from ESXi Host Client	29
Proactive Tests on vSAN Cluster	30
4 Managing Proactive Hardware	31
About Hardware Support Managers	31
Deploying and Configuring Hardware Support Managers	32
Registering Hardware Support Manager	32
Associating and Dissociating Hosts	32
Processing Hardware Failures	32
5 Monitoring vSAN Performance	34
About the vSAN Performance Service	34
Configure vSAN Performance Service	35

Use Saved Time Range in vSAN Cluster	36
View vSAN Cluster Performance	37
View vSAN Host Performance	39
View vSAN VM Performance	41
Use vSAN I/O Insight	42
View vSAN I/O Insight Metrics	42
Use vSAN I/O Trip Analyzer	43
View vSAN Performance Metrics for Support Cases	45
Using vSAN Performance Diagnostics	46
View vSAN Obfuscation Map	47

6 Handling Failures and Troubleshooting vSAN 48

Uploading a vSAN Support Bundle	48
Using Esxcli Commands with vSAN	49
Using vsantop Command-Line Tool	52
vSAN Configuration on an ESXi Host Might Fail	52
Not Compliant Virtual Machine Objects Do Not Become Compliant Instantly	53
vSAN Cluster Configuration Issues	53
Handling Failures in vSAN	54
Failure Handling in vSAN	54
Troubleshooting vSAN	63
Replacing Existing Hardware Components in vSAN Cluster	67

7 Shutting Down and Restarting the vSAN Cluster 73

Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard	73
Restart the vSAN Cluster	75
Manually Shut Down and Restart the vSAN Cluster	75

About vSAN Monitoring and Troubleshooting

vSAN Monitoring and Troubleshooting describes how to monitor and troubleshoot VMware vSAN® by using the vSphere Client.

In addition, *vSAN Monitoring and Troubleshooting* explains how to monitor and troubleshoot a vSAN cluster using esxcli and RVC commands, and other tools.

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

Intended Audience

This 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*.

What Is 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.

Read the following topics next:

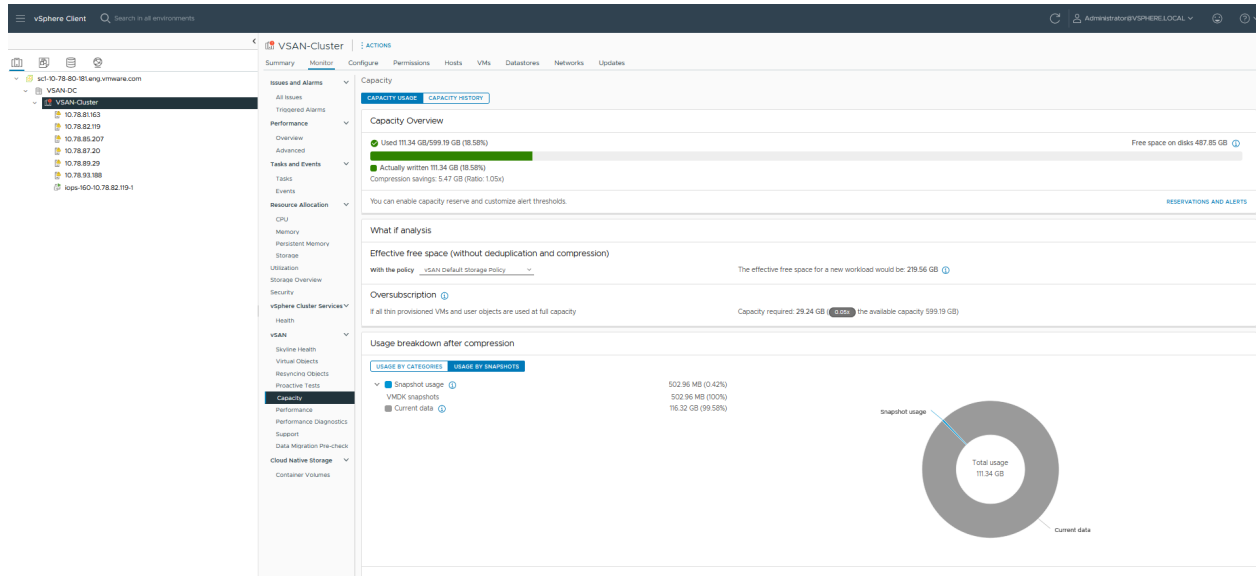
- [Monitor vSAN Capacity](#)
- [Monitor Physical Devices in vSAN Cluster](#)
- [Monitor Devices that Participate in vSAN Datastores](#)
- [Monitor Virtual Objects in vSAN Cluster](#)
- [Monitor Container Volumes in vSAN Cluster](#)
- [About Reserved Capacity in vSAN Cluster](#)
- [About vSAN Cluster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

Monitor vSAN Capacity

You can monitor the capacity of the vSAN datastore, vSAN Direct storage, and Persistent Memory (PMem) storage.

You can 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.
- 3 Under vSAN, click **Capacity** to view the vSAN capacity information.

Results

- The Capacity Overview displays the storage capacity of the vSAN datastore, including free space available, used space, and the space that is actually written and physically consumed on the vSAN disks. For clusters that have deduplication and compression enabled, you can view the deduplication and compression savings and the deduplication and compression ratio.

Note vSAN Express Storage Architecture (ESA) does not support deduplication.

Terms	Description
Free space	Total free space available in the cluster
Used space	Total written physical space
Actually written	Actually used capacity. This capacity is displayed when deduplication or compression are not enabled.
Compression savings	Space saved when data compression is enabled.
Object reserved	Includes the reservation for objects created with a policy that has specified object space reservation. This capacity is not actually used by the objects.
Reserved capacity	Includes the operations reserve and the host rebuild reserve.

- What if analysis enables you to estimate the effective free space while keeping the deduplication ratio as 1. **Effective free space** is an estimate of free space available based

on the selected storage policy. The effective free space typically is smaller than the free space available on the disks, due to the cluster topology or the distribution of space across fault domains. For example, consider a cluster with 100 GB free space available on the disks. However, 100 GB cannot be provisioned as a single 100 GB object due to the distribution of free space across fault domains. If there are three fault domains and each fault domain has 33 GB free space, then the largest object that you can create with FTT 1 is 33 GB.

Oversubscription reports the vSAN capacity required if all the thin provisioned VMs and user objects are used at full capacity. It shows a ratio of the required usage compared with the total available capacity. While calculating the oversubscription, vSAN includes all the available VMs, user objects, and storage policy overhead, and does not consider the vSAN namespace and swap objects.

Note Oversubscription is applicable only for vSAN hosts that are running 6.7 Update 1 or later.

Note PMem storage does not support What if analysis and Oversubscription.

- The Usage breakdown before deduplication and compression displays the amount of storage space used by VMs, user objects, and the system. You can view a pie chart that represents the different usage categories. Click the pie chart to view the details of the selected category.

Following are the different usage categories available:

Category	Description
VM (user objects) usage	<p>Displays the following:</p> <ul style="list-style-type: none"> ■ VM home objects - Usage of VM namespace object. ■ Swap objects - Usage of 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 - Usage of memory snapshot file for VMs. ■ Block container volumes (attached to a VM) - Capacity consumed by the container objects that are attached to a VM. ■ vSphere replication persistent state file - vSAN object used to store the persistent state file (PSF) at source site.
Non-VM (user objects) usage	<p>Displays iSCSI objects, block container volumes that are not attached to VM, user-created files, ISO files, VM templates, files shares, file container volumes, and vSAN objects used by the vSphere replication service at the target site.</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 - vSAN on-disk format overhead that may take up on the capacity drives. ■ ESA object overhead - vSAN ESA uses the capacity to store object metadata and to provide high performance. ■ 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. ■ Operations usage - Temporary space usage in a cluster. The temporary space usage includes temporary capacity used for rebalance operations or moving objects due to FTT changes.

Category	Description
	<ul style="list-style-type: none"> ■ Native trace objects - Capacity consumed by objects created for storing vSAN traces.
<p>Note PMem only supports VMDK, Non-Volatile Dual In-line Memory Module (NVDIMM), and file system overhead.</p>	

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*.

In vSAN ESA, Usage by Snapshots displays the snapshot usage by the vSAN datastore. You can delete one or more snapshots and free the used space, thus managing space consumption. To delete a snapshot, right-click the virtual machine > **Snapshots** > **Manage Snapshots**. Click **Delete** to delete a snapshot. Click **Delete All Snapshots** to delete all the snapshots. The following are the different usage snapshots available:

Snapshot	Description
Container volume snapshots	Displays the container volume snapshot usage in the vSAN datastore.
VMDK snapshots	Displays the VMDK snapshot usage in the vSAN datastore.
vSAN file share snapshots	Displays the file share snapshot usage in the vSAN datastore.
Current data	Displays the usage data that is not included in the snapshot usage data. You can calculate the current data by subtracting the total snapshot usage from the total used space.

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

The Capacity monitor displays two thresholds represented as vertical markers in the bar chart:

- Operations threshold - Displays the space vSAN requires to perform internal operations in the cluster. If the used space reaches beyond that threshold, vSAN might not be able to operate properly.
- Host rebuild threshold - Displays the space vSAN requires to tolerate one host failure. If the used space reaches beyond the host rebuild threshold and the host fails, vSAN might not successfully restore all data from the failed host.

If you enable reserved capacity, the Capacity monitor displays the following:

- Operations reserve - Reserved space in the cluster for internal operations.
- Host rebuild reserve - Reserved space for vSAN to be able to repair in case of single host failure. The Capacity monitor displays the host rebuild threshold only when the host rebuild reserve is enabled.

If the resynchronization of objects is in progress in a cluster, vSAN displays the capacity used in the capacity chart as operations usage. In case there is enough free space in the cluster, vSAN might use more space than the operations threshold for the resyncing operations to complete faster.

Click **Configure** tab to enable the capacity reserve. You can also click **Configure > vSAN > Services** to enable the capacity reserve. For more information on configuring the reserved capacity, see [Configure Reserved Capacity for vSAN Cluster](#).

In a cluster, if there is more utilization than the host rebuild threshold and the reserved capacity is not enabled, the capacity chart turns yellow as a warning. If the most consumed host fails, vSAN cannot recover the data. If you enable the host rebuild reserve, the capacity chart turns yellow at 80% of the host rebuild threshold. If the used space reaches beyond the operations threshold and the reserved capacity is not enabled, vSAN cannot perform or complete operations such as rebalance, resync object components due to policy changes, and so on. In that case, the capacity chart turns red to indicate that the disk usage exceeds the operations threshold. For more information about capacity reserve, see [About Reserved Capacity in vSAN Cluster](#).

Monitor Physical Devices in vSAN Cluster

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

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Click **Disk Management** to review all hosts, cache devices, and capacity devices in the cluster. The physical location is based on the hardware location of cache and capacity devices on vSAN hosts. You can see the virtual objects on any selected host, disk group, or disk and view the impact of the selected entity to the virtual objects in the cluster.

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 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.
- 4 Click  to filter the virtual objects based on name, type, storage policy, and UUID.
 - a Select the check box on one of the virtual objects and click **View Placement Details** to open the Physical Placement dialog box. You can view the device information, such as name, identifier or UUID, number of devices used for each virtual machine, and how they are mirrored across hosts.
 - b On the Physical Placement dialog box, select 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, expand the VM to which the container is attached.

- 5 Select the check box of the attached block type or file volumes and click **View Performance**. You can use the vSAN cluster performance charts to monitor the workload in your cluster. For more information on the vSAN cluster performance charts, see [View vSAN Cluster Performance](#).
- 6 Select the check box on one of the container volumes and click **View Container Volume**. For more information about monitoring container volumes, see [Monitor Container Volumes in vSAN Cluster](#).
- 7 Select the check box on one of the file volumes and click **View File Share**. For more information about file volume, see *Administering VMware vSAN*.

Monitor Container Volumes in vSAN Cluster

You can view the status of the container volumes in the vSAN cluster.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Monitor** tab.
- 3 Under **Cloud Native Storage**, select **Container Volumes** to view the container volumes in the vSAN cluster. You can view information about the volume name, label, datastore, compliance status, health status, and capacity quota.
- 4 Click  to view the following:
 - Click the **Basics** tab to view the volume details such as volume type, ID, datastore, storage policy, compliance, and health status.
 - Click the **Kubernetes objects** tab to view Kubernetes related data such as Kubernetes cluster, namespace, pod, persistent volume claim, labels, and so on.
 - Click the **Physical Placement** tab to view the type, host, cache, and capacity disk of the virtual object components.
 - Click the **Performance** tab to view the performance of the container volumes.
- 5 Select the check box for the volumes that have an out-of-date policy status. Click **Reapply Policy** to reapply the policy on the selected volumes.
- 6 Select the check box for the container volume you want to delete and click **Delete**.
- 7 Use the **Add Filter** option to add filters to the container volumes.

About Reserved Capacity in vSAN Cluster

vSAN requires capacity for its internal operations.

For a cluster to be able to tolerate a single host failure, vSAN requires free space to restore the data of the failed host. The capacity required to restore a host failure matches the total capacity of the largest host in the cluster. These values are represented as thresholds in the Capacity Monitor page:

- Operations threshold - Displays the space vSAN requires to run its internal operations in the cluster. If the used space exceeds the operations threshold, vSAN might not operate properly.
- Host rebuild threshold - Displays the space vSAN requires to tolerate one host failure. If the used space exceeds the host rebuild threshold and the host fails, vSAN might not successfully restore all data from the failed host.

For more information on the capacity thresholds, see [Monitor vSAN Capacity](#).

vSAN provides you the option to reserve the capacity in advance so that it has enough free space available to perform internal operations and to repair data back to compliance following a single host failure. By enabling reserve capacity in advance, vSAN prevents you from using the space to create workloads and intends to save the capacity available in a cluster. By default, the reserved capacity is not enabled.

If there is enough free space in the vSAN cluster, you can enable the operations reserve and/or the host rebuild reserve.

- Operations reserve - Reserved space in the cluster for vSAN internal operations.
- Host rebuild reserve - Reserved space for vSAN to be able to repair in case of a single host failure.

These soft reservations prevent the creation of new VMs or powering on VMs if such operations consume the reserved space. Once the reserved capacity is enabled, vSAN does not prevent powered on VM operations, such as I/O from the guest operating system or applications from consuming the space even after the threshold limits are reached. After you enable the reserved capacity, you must monitor the disk space health alerts and capacity usage in the cluster and take appropriate actions to keep the capacity usage below the threshold limits.

Note The reserved capacity is not supported on a vSAN stretched cluster, cluster with fault domains and nested fault domains, ROBO cluster, or if the number of hosts in the cluster is less than four.

To enable reserved capacity for the host rebuild, you must first enable the operations reserve. When you enable operations reserve, vSAN reserves 5% additional capacity in the operations reserve as a buffer to ensure you have time to react to the capacity fullness before the actual threshold is reached.

vSAN indicates when the capacity usage is high in a cluster. The indications can be in the form of health alerts, capacity chart turning yellow or red, and so on. Due to the reservation, vSAN might not have enough free space left. This results in the inability to create VMs or VM snapshots, creating or extending virtual disks, and so on.

Note You cannot enable reserved capacity, if the cluster is at a capacity higher than the specified threshold.

Capacity Reservation Considerations

Following are the considerations if you enable reserved capacity:

- When you enable reserved capacity with the host rebuild reserve and a host is put into maintenance mode, the host might not come back online. In this case, vSAN continues to reserve capacity for another host failure. This host failure is in addition to the host that is already in the maintenance mode. This might cause the failure of operations if the capacity usage is above the host rebuild threshold.

- When you enable reserved capacity with the host rebuild reserve and a host fails, vSAN might not start repairing the affected objects until the repair timer expires. During this time, vSAN continues to reserve capacity for another host failure. This can cause failure of operations if the capacity usage is above the current host rebuild threshold, after the first host failure. After the repairs are complete, you can deactivate the reserved capacity for the host rebuild reserve if the cluster does not have the capacity for another host failure.

Configure Reserved Capacity for vSAN Cluster

You can configure reserved capacity for a vSAN cluster to reserve capacity for internal operations.

Prerequisites

You can also configure reserve capacity to reserve capacity for data repair following a single host failure. Ensure that you have the following required privileges: **Host.Inventory.EditCluster** and **Host.Config.Storage**.

Verify that the vSAN cluster:

- Is not configured as a vSAN stretched cluster or ROBO cluster.
- Has no fault domains and nested fault domains created.
- Has a minimum of four hosts.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.

4 Click to edit the Reservations and Alerts.

Reservations and Alerts | Vsan2Cluster

Enabling operations reserve for vSAN helps ensure that there will be enough space in the cluster for internal operations to complete successfully. Enabling host rebuild reserve allows vSAN to tolerate one host failure. When reservation is enabled and capacity usage reaches the limit, new workloads fail to deploy.

[About Reserved Capacity](#)

The reserved capacity is displayed in the capacity overview:

Actually written 118.17 GB (19.72%)

☒ Operations reserve

☒ Host rebuild reserve

The default health alerts are system recommendations based on your reservation configuration.

☒ Customize alerts

⚠ Available capacity warning threshold % 70
Set available capacity threshold for receiving warning alert

❗ Available capacity error threshold % 90
Set available capacity threshold for receiving error alert

CANCEL **APPLY**

- 5 Click to enable or deactivate the operations reserve. On enabling the operations reserve, vSAN ensures that the cluster has enough space to complete the internal operations.
- 6 Click to enable or deactivate the host rebuild reserve. On enabling the host rebuild reserve, vSAN provides the reservation of space to repair data back to compliance following a single host failure. You can enable the host rebuild reserve only after you enable the operations reserve. After enabling, if you deactivate the operations reserve, the host rebuild reserve gets automatically deactivated.
- 7 Select **Customize alerts**. You can set a customized threshold to receive warning and error alerts. The threshold percentage is calculated based on the available capacity, which is the difference between the total capacity and the reserved capacity. If you do not set a customized value, vSAN uses the default thresholds to generate alerts.
- 8 Click **Apply**.

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 version 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.

Monitor the Resynchronization Tasks in 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 7.0 or later.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Select the **Monitor** tab.
- 3 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.
Total resyncing ETA	Estimated time left for the resynchronization to complete. 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.
Scheduled resyncing	Remaining number of objects to be resynchronized. 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. Resynchronization of objects starts once the timer expires. The pending category displays the objects with the expired delay timer that cannot be resynchronized. This can be due to insufficient resources in the current cluster or the vSAN FTT policy set on the cluster not being met.

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.

About vSAN Cluster Rebalancing

When any capacity device in your cluster reaches 80 percent full, vSAN automatically rebalances the cluster.

The vSAN cluster rebalancing continues until the space and components 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 FTT=0 reside on the host.

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

Configure Automatic Rebalance in vSAN Cluster

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

Your vSAN cluster can become unbalanced based on the space or component usage for many reasons such as when you create objects of different sizes, when you add new hosts or capacity devices, or when objects write different amounts of data to the disks. If the cluster becomes unbalanced, vSAN automatically rebalances the disks. Based on the space or component usage, this operation moves components from over-utilized disks to under-utilized disks.

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

Disk rebalancing can impact the I/O performance of your vSAN cluster. By default the rebalance threshold is set at 30 percentage and ensures that the cluster remains relatively balanced without significantly impacting the performance. If the cluster becomes severely imbalanced, such as after adding one or more hosts or disks, temporarily using a lower threshold of 10 or 20 percentage makes the cluster evenly balanced. This must be done during off-peak periods to minimize the performance impact during the rebalancing activity. Once the rebalancing is complete, you can change the threshold back to the default 30 percentage.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Configure** tab.
- 3 Under vSAN, select **Services**.

4 Click to edit **Advanced Options**.

Advanced Options | Vsan2Cluster

Object repair timer 60
The amount of minutes vSAN waits before repairing an object after a host is either in a failed state (absent failures) or in Maintenance Mode.

☐ **Site read locality**
When enabled, reads to vSAN objects occur locally. When disabled, reads occur across both sites for stretched cluster.

☒ **Thin swap**
When enabled, swap objects will not reserve 100% of their space on vSAN datastore; storage policy reservation will be respected.

☐ **Guest Trim/Unmap**
Guest Trim/Unmap cannot be disabled for cluster with vSAN ESA. When enabled, vSAN automatically reclaims blocks after Guest OS file deletions. VMs that are running need to be power cycled for the setting to take effect. Refer to the administrative guide for prerequisites.

☒ **Automatic rebalance**
When the cluster is unbalanced, rebalance starts automatically after enabling automatic rebalance. Rebalance can wait up to 30 minutes to start, giving time to high priority tasks like EMM, repair, etc. to use the resources before rebalancing.

Rebalancing threshold % 30
Determines when background rebalancing starts in the system. If any two disks in the cluster have this much variance then rebalancing begins. It will continue until it is turned off or the the variance between disks is less than 1/2 of the rebalancing threshold.

CANCEL **APPLY**

5 Click to enable or deactivate Automatic Rebalance.

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

What to do next

You can use the vSAN Skyline Health 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 Click  and type **vSAN** in the search box 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.

View vSAN Network Alarms


vSAN network diagnostics queries the latest network metrics and compares the metrics statistics with the defined threshold values.

If the value reaches above the threshold that you have set, vSAN network diagnostics raises an alarm. You must acknowledge and manually reset the triggered alarms to green after fixing the network issues.

Prerequisites

The vSAN performance service must be turned on.

Procedure

- 1 Navigate to the host in the vSAN cluster.
- 2 Click the **Monitor** tab.
- 3 Under vSAN, select **Performance**.
- 4 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.
- 5 Select . In the Threshold settings dialog box, enter a threshold value to receive warning and error alert.
- 6 Click **Save**.

Results

vSAN displays the performance statistics of all the network I/Os in use. vSAN network diagnostics result appears in the vCenter Server alerts. The redirection to the related performance charts is available in the vSAN network alerts generated by the network diagnostics service.

Using the VMkernel Observations for Creating vSAN Alarms

VMkernel Observations (VOBs) are system events that you can use to set up vSAN alarms.

vSAN alarms are used 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.componentthresh old</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.

Table 2-1. VOB IDs for vSAN (continued)

VOB ID	Description
esx.problem.vsan.clustering.disabled	vSAN clustering services are not enabled.
esx.problem.vsan.lsom.congestionthreshold	vSAN device memory or SSD congestion has been updated.
esx.problem.vsan.net.not.ready	A vmknic is added to vSAN network configuration without a valid IP address. This happens when the vSAN network is not ready.
esx.problem.vsan.net.redundancy.lost	The vSAN network configuration does not have the required redundancy.
esx.problem.vsan.no.network.connectivity	vSAN does not have existing networking configuration, which is in use.
esx.problem.vsan.vmknic.not.ready	A vmknic is added to the vSAN network configuration without a valid IP address.
esx.problem.vob.vsan.lsom.devicerepair	The vSAN device is offline and in a repaired state because of I/O failures.
esx.problem.vsan.health.ssd.endurance	One or more vSAN disks exceed the warning usage of estimated endurance threshold.
esx.problem.vsan.health.ssd.endurance.error	A vSAN disk exceeds the estimated endurance threshold.
esx.problem.vsan.health.ssd.endurance.warning	A vSAN disk exceeds 90% of its estimated endurance threshold.

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 Navigate to the vSAN cluster.
- 2 On the **Configure** tab, select **Alarm Definitions** and click **Add**.
- 3 In the Name and Targets page, enter a name and description for the new alarm.
- 4 From the **Target type** drop-down menu, select the type of inventory object that you want this alarm to monitor and click **Next**.

Depending on the type of target that you choose to monitor, the summary that follows the **Targets**, change.

- 5 In the Alarm Rule page, select a trigger from the drop-down menu.

The combined event triggers are displayed. You can set the rule for a single event only. You must create multiple rules for multiple events.

- 6 Click **Add Argument** to select an argument from the drop-down menu.
 - a Select an operator from the drop-down menu.
 - b Select an option from the drop-down menu to set the threshold for triggering an alarm.
 - c Select severity of the alarm from the drop-down menu. You can set the condition to either **Show as Warning** or **Show as Critical**, but not for both. You must create a separate alarm definition for warning and critical status.
- 7 Select **Send email notifications**, to send email notifications when alarms are triggered.
- 8 In the **Email to** text box, enter recipient addresses. Use commas to separate multiple addresses.
- 9 Select **Send SNMP traps** to send traps when alarms are triggered on a vCenter Server instance.
- 10 Select **Run script** to run scripts when alarms are triggered.
- 11 In the **Run this script** text box, enter the following script or command:

For this type of command...	Enter this...
EXE executable files	Full pathname of the command. For example, to run the cmd.exe command in the C:\tools directory, type: <pre>c:\tools\cmd.exe</pre>
BAT batch file	Full pathname of the command as an argument to the c:\windows\system32\cmd.exe command. For example, to run the cmd.bat command in the C:\tools directory, type: <pre>c:\windows\system32\cmd.exe /c c:\tools\cmd.bat</pre>

- 12 Select an advanced action from the drop-down menu. You can define the advanced actions for virtual machine and hosts. You can add multiple advanced actions for an alarm.
- 13 Click **Next** to set the Reset Rule.
- 14 Select **Reset the alarm** to green and click **Next** to review the alarm definition.
- 15 Select **Enable this alarm** to enable the alarm and click **Create**.

Results

The alarm is configured.

Monitoring vSAN Skyline Health

3

You can check the overall health of the vSAN cluster, including hardware compatibility and networking configuration and operations.

You can also check the advanced vSAN configuration options, storage device health, and virtual machine object health.


Read the following topics next:

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

About the vSAN Skyline Health

Use the vSAN Skyline health to monitor the health of your vSAN cluster.

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



You can use Overview to monitor the core health issues of your vSAN cluster. You can also view the following:

- Cluster health score based on the health findings
- View the health score trend for 24 hours
- View the health score trend for a particular period

Ensure that the **Historical Health Service** is enabled to view details of the Health score trend. Click **View Details** in the Health score trend chart to examine the health state of the cluster for a selected time point within 24 hours. Use **Custom** to customize the time range as per your requirement.

You can use the vSAN Health findings to diagnose issues, troubleshoot problems, and remediate the problems.

The health findings are classified as follows:

- **Unhealthy** – Critical or important issue(s) being detected that needs attention.
- **Healthy** – There are no issues found that needs attention.
- **Info** – Health findings which may not impact the cluster running state but important for awareness.
- **Silenced** – Health findings have been silenced without triggering vSAN health alarm by intention.

To troubleshoot an issue, you can sort the findings by root cause to resolve the primary issues initially and then verify if the impacted issues can also be resolved.

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

If you participate in the Customer Experience Improvement Program (CEIP), you can run health findings and send the data to VMware for advanced analysis. Click **Retest with Online health** and then click **OK**. Online notifications is enabled by default if the vCenter Server can connect to VMware Analytics Cloud without enrolling CEIP. If you do not want to participate in CEIP, you can still receive vSAN health notifications for software and hardware issues using Online notifications.

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 findings.

Viewing vSAN Health History

The vSAN health history helps you examine health issues by querying the historical health records. You can only view the historical health data of a cluster. By default, the health history is enabled. To deactivate the health history, select the cluster and navigate to the **Configure > vSAN > Services > Historical Health Service** and click **Disable**. If you deactivate the health history, all the health data collected on the vCenter Server database gets purged. The database stores the health data for up to 30 days depending on the available capacity.

Using the Skyline Health view, you can view the health history for a selected time range. The start date of the time range must not be earlier than 30 days from the current date. The end date must not be later than the current date. Based on your selection, you can view the historical health findings. Click **View History Details** to view the history of a health finding within a selected time period. The historical data is displayed as a graphical representation with green circles, yellow triangles, and red squares showing success, warning, and failure respectively. The detailed information about each health finding result is displayed in a table.

Using vSAN Support Insight

vSAN support insight is a platform that helps you maintain a reliable and consistent compute, storage, and network environment. VMware support uses the vSAN support insight to monitor the vSAN performance diagnostics and resolve performance issues. vSAN uses Customer Experience Improvement Program (CEIP) to send data to VMware for analysis on a regular basis. To deactivate CEIP, select **vSphere Client > Administration > Customer Experience Improvement Program > Leave Program**.

Check vSAN Skyline Health

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

Procedure

- 1 Navigate to the vSAN cluster.
- 2 Click the **Monitor** tab.
- 3 Under **vSAN**, select **Skyline Health** to review the vSAN health finding.
- 4 Under Health findings, perform the following:
 - Click **Unhealthy** to view the issues and the details. Click **Troubleshoot** to troubleshoot and fix an issue. You can sort the findings by root cause to resolve the primary issues and then verify if the impacted issues can be resolved.
 - Click **View History Details** to identify the status history of the health finding for a particular time period. The default time period is 24 hours. You can also customize the time period as per your requirement. The status of an unhealthy finding is displayed in yellow or red. Click the **Ask VMware** button to open a knowledge base article that describes the health finding and provides information about how to resolve the issue. You can also view the status history of the health finding for a given period using **History Details** tab.
 - You can click **Silence alert** on a health finding, so it does not display any warnings or failures.
 - Click **Healthy** to view health findings that are healthy. Click **View Current Result** to view the current status of the health finding. Click **View History Details** to identify the status history of the health finding for a particular time period. The status is displayed in green. You can also view the status history of the health finding for a given period using **History Details** tab.

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 findings.

- 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, and click **Save**.

Proactive Tests on vSAN Cluster

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

Note You must not conduct the proactive test in a production environment as it creates network traffic and impacts the vSAN workload.

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. You can run the proactive test at a maximum speed limit of 10Gbps. In vSAN ESA, the proactive test reports error when the result is zero bps and the Health Status displays the test results as info when the result is a non-zero number.

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

Managing Proactive Hardware

4

vSAN Proactive Hardware Management (PHM) informs you of any dying disks based on disk predictive failure events generated by the Original Equipment Manufacturer (OEM) vendor.

Based on this information provided, you can take the necessary remediation. PHM resides within the vSAN management service on the vCenter Server. The Hardware Support Manager (HSM) is registered with the vCenter Server. PHM collects vendor hardware information from HSM and sends it to vSAN.

Read the following topics next:

- [About Hardware Support Managers](#)
- [Deploying and Configuring Hardware Support Managers](#)
- [Registering Hardware Support Manager](#)
- [Associating and Dissociating Hosts](#)
- [Processing Hardware Failures](#)

About Hardware Support Managers

The deployment method and the management of a hardware support manager are determined by the respective OEM vendor.

Several of the major OEM vendors develop and supply hardware support managers. For example:

- Dell - The hardware support manager that Dell provides is part of their host management solution, OpenManage Integration for VMware vCenter (OMIVV), which you deploy as an appliance.
- HPE - The hardware support managers that HPE provides are part of their management tools, iLO Amplifier and OneView, which you deploy as appliances.
- Lenovo - The hardware support manager that Lenovo provides is part of their server management solution, Lenovo XClarity Integrator for VMware vCenter, which you deploy as an appliance.

You can find the full list of all VMware-certified hardware support managers in the VMware Compatibility Guide at <https://www.vmware.com/resources/compatibility/search.php?deviceCategory=hsm>.

Deploying and Configuring Hardware Support Managers

Regardless of the hardware vendor, you must deploy the hardware support manager appliance on a host with sufficient memory, storage, and processing resources.

Typically, hardware support manager appliances are distributed as OVF or OVA templates. You can deploy them on any host in any vCenter Server instance.

After you deploy the appliance, you must power on the appliance virtual machine and register the appliance as a vCenter Server extension. You might need to log in to the appliance as an administrator. Each hardware support manager might register with only one or multiple vCenter Server systems.

A vCenter Server plug-in user interface might become available in the vSphere Client after you deploy a hardware support manager appliance, but the hardware support manager might also have a separate user interface of its own. For example, OMIVV, iLO Amplifier, and Lenovo XClarity Integrator for VMware vCenter all have a vCenter Server plug-in user interface, which helps you configure and work with the respective hardware support manager.

For detailed information about deploying, configuring, and managing hardware support managers, refer to the respective OEM-provided documentation.

Registering Hardware Support Manager

You must register HSM with PHM that resides within the vSAN management service on the vCenter Server using the vendor management service.

For detailed information about registering hardware support managers, refer to the respective OEM-provided documentation.

Associating and Dissociating Hosts

After registering HSM with PHM, you need to associate appropriate hosts available in the vCenter Server with the HSM.

This enables PHM on each host. HSM informs PHM on any change in the managed host list. PHM associates the managed hosts available in a vSAN cluster. When a host is associated or dissociated with PHM, vCenter Server event gets generated. For detailed information about associating and dissociating hosts, refer to the respective OEM-provided documentation.

Processing Hardware Failures

PHM checks for HSM generated hardware failure events every 10 minutes.

You can customize the time interval using the vSAN configuration file.

Procedure

- 1 Log in to vCenter Server console as root.
- 2 Open the `/usr/lib/vsan-health/VsanVcMgmtConfig.xml` file.
- 3 Set the interval value using `healthUpdatePollIntervalInSeconds` xml tag.
- 4 Restart the vSAN Health service.

Results

PHM uses these events to generate alarms, which appears in the vSAN Skyline Health. For more information on the vSAN Skyline Health events, see the VMware knowledge base article at <https://knowledge.broadcom.com/external/article?articleNumber=367770>.

Monitoring vSAN Performance

5

You can monitor the performance of your vSAN cluster.

Performance charts are available for clusters, hosts, physical disks, virtual machines, and virtual disks.

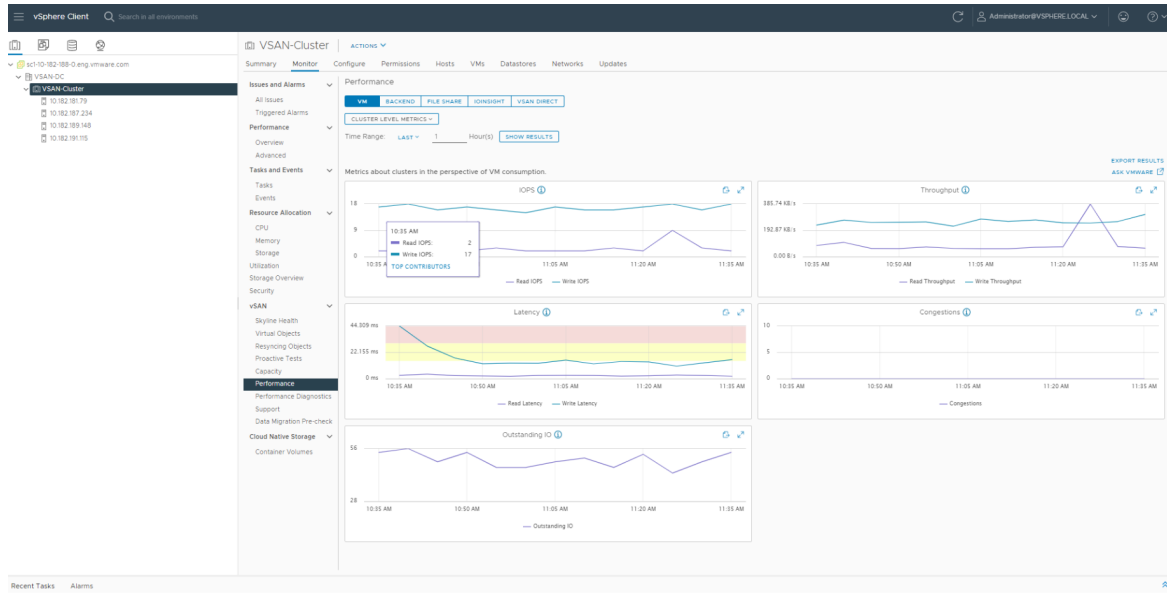
Read the following topics next:

- [About the vSAN Performance Service](#)
- [Configure vSAN Performance Service](#)
- [Use Saved Time Range in vSAN Cluster](#)
- [View vSAN Cluster Performance](#)
- [View vSAN Host Performance](#)
- [View vSAN VM Performance](#)
- [Use vSAN I/O Insight](#)
- [Use vSAN I/O Trip Analyzer](#)
- [View vSAN Performance Metrics for Support Cases](#)
- [Using vSAN Performance Diagnostics](#)
- [View vSAN Obfuscation Map](#)

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 vSAN OSA, you can optionally enable or deactivate the Performance Service. You can enable and configure the Performance Service. When you create vSAN ESA, the Performance Service is enabled by default. You can then 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 7.0 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.
- 3 Under vSAN, select **Services**.
- 4 (Optional for vSAN ESA cluster.) Click the Performance Service **Enable** button.

- 5 (Optional for vSAN ESA cluster.) In vSAN Performance Service Settings, select a storage policy for the stats database object.
- 6 (Optional for vSAN ESA cluster.) Click **Enable** to enable vSAN Performance Service.
- 7 Click **Edit** if you want to select a different storage policy in the vSAN Performance Service Settings.
- 8 (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.
- 9 (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.
- 10 Click **Apply**.

Use Saved Time Range in vSAN Cluster

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 7.0 or later.

Procedure

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

You can save the selected time range at the VM and the host level.

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 7.0 or later.
 - To view file service performance charts, you must enable vSAN File Service.
 - To view vSAN Direct performance charts, you must claim disks for vSAN Direct.
 - To view PMem performance charts, you must have PMem storage attached to the hosts in the cluster.
-

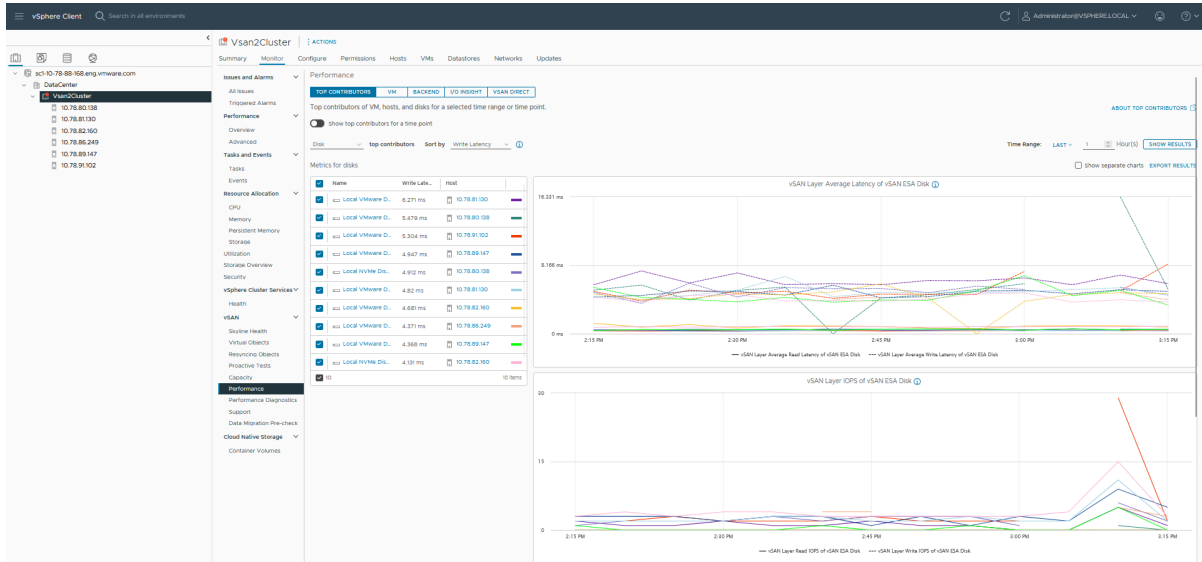
Prerequisites

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.
- 3 Under vSAN, select **Performance**.

4 Select Top Contributors.



Perform one of the following:

- Select a time range to view the hotspot entities in the charts. You can view the top 10 hotspot entities as aggregated metrics for the selected time range. You can view the hotspots of VMs, disk groups (vSAN OSA) or disks (vSAN ESA), host (backend), or host (frontend). You have the option to enable separate charts.
- Select a single timestamp to identify the VMs, disk groups (vSAN OSA) or disks (vSAN ESA), host (backend), or host (frontend) that consume the most IOPS, have the highest I/O throughput, or I/O latency. For example, based on the I/O latency graph of the cluster, you can select a timestamp and get the top contributors with latency statistics. You can also select a single contributor and view the latency graph. You have the option to switch between the combined view and table view.

5 Select VM.

Perform one of the following:

- Select **Cluster level metrics** to display the aggregated performance metrics for the cluster that you selected.
- Select **Show specific VMs** to display metrics for all the VMs selected. If you enable **Show separate chart by VMs**, vSAN displays separate metrics for all the VMs selected.

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. You can also select **Real-Time** as the time range that displays real-time data that is automatically refreshed every 30 seconds. The real-time statistics data is retained in the SQL database for seven days until it gets purged.

- 6 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.
- 7 Select **File Share** and choose a file. Select a time range for your query. Select **NFS performance** or **File system performance** based on the protocol layer performance or file system layer performance that you want to display. vSAN displays performance charts for vSAN file services, including IOPS, throughput, and latency.
- 8 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.
- 9 (Optional) Select **I/O Insight**. For more information on I/O Insight, see [Use vSAN I/O Insight](#).
- 10 Select **vSAN Direct** to display the performance data of the vSAN direct disks. Select a time range for your query. vSAN displays performance charts for vSAN direct, including IOPS, bandwidth, latency, and outstanding I/O.
- 11 Select **PMEM** to display the performance data of all VMs placed on the PMem storage. Select a time range for your query. You can also select **Real-time** as the time range that displays real time data that is automatically refreshed every 30 seconds. PMem displays performance charts including IOPS, bandwidth, and latency. For more information about PMem metrics collection settings, see <https://kb.vmware.com/s/article/89100>.
- 12 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 7.0 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.
- 3 Under vSAN, select **Performance**.
- 4 Select **VM**.
 - Select **Host level metrics** to display the aggregated performance metrics for the host that you selected.
 - Select **Show specific VMs** to display metrics for all the VMs selected on the host. If you enable **Show separate chart by VMs**, vSAN displays separate metrics for all the VMs selected on the host.

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. You can also select **Real-Time** as the time range that displays real-time data that is automatically refreshed every 30 seconds. The real-time statistics data is retained in the SQL database for seven days until it gets purged.

- 5 In vSAN ESA, select **Backend Cache**. Select a time range for your query. vSAN displays the performance charts for the backend cache operations of the host, including the overall backend cache statistics, the overall cache miss by the different types, cache miss by types for the different transactions, and the catch latency for the different transactions.
- 6 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.
- 7 Perform one of the following:
 - 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.
 - In vSAN ESA, select **Disks**, and then select a disk. Select a time range for your query. vSAN displays performance charts for the disk, including vSAN layer IOPS, throughput, and latency. It also displays the physical or firmware layer IOPS, throughput, and latency.
- 8 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.
- 9 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.

- 10 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.
- 11 (Optional) Select **I/O Insight**. For more information on I/O Insight, see [Use vSAN I/O Insight](#).
- 12 Select **vSAN Direct** to display the performance data of the vSAN direct disks. Select a time range for your query. vSAN displays performance charts for vSAN direct, including IOPS, bandwidth, latency, and outstanding I/O.
- 13 Select **PMEM** to display the performance data of all VMs placed on the PMem storage. Select a time range for your query. You can also select **Real-time** as the time range that displays real time data that is automatically refreshed every 30 seconds. PMem displays the performance charts including IOPS, bandwidth, and latency. For more information about PMem metrics collection settings, see <https://kb.vmware.com/s/article/89100>.
- 14 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.
- 3 Under vSAN, select **Performance**.
- 4 Select **VM**. Select a time range for your query. vSAN displays performance charts for the VM, including IOPS, throughput, and latency.
- 5 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. The virtual SCSI latency performance charts display a highlighted area due to the IOPS limit enforcement.
- 6 (Optional) In the Virtual Disk, click **New I/O Insight Instance**. For more information on I/O Insight, see [Use vSAN I/O Insight](#).

- 7 Click **Refresh** or **Show Results** to update the display.

Use vSAN I/O Insight

I/O Insight allows you to select and view I/O performance metrics of virtual machines in a vSAN cluster.

By understanding the I/O characteristics of VMs, you can ensure better capacity planning and performance tuning.

Procedure

- 1 Navigate to the vSAN cluster or host.

You can also access I/O Insight from the VM. Select the VM and navigate to **Monitor > vSAN > Performance > Virtual Disks**.

- 2 Click the **Monitor** tab.
- 3 Under **vSAN**, select **Performance**.
- 4 Select the **I/O Insight** tab and click **New Instance**.
- 5 Select the required hosts or VMs that you want to monitor. You can also search for VMs.
- 6 Click **Next**.
- 7 Enter a name and select a duration.
- 8 Click **Next** and review the instance information.
- 9 Click **Finish**.

I/O Insight instance monitors the selected VMs for the specified duration. However, you can stop an instance before completion of the duration that you specified.

Note VMs monitored by I/O Insight must not be vMotioned. vMotion stops the VMs from being monitored and will result in an unsuccessful trace.

Results

vSAN displays performance charts for the VMs in the cluster, including IOPS, throughput, I/O size distribution, I/O latency distribution, and so on.

What to do next

You can view metrics for the I/O Insight instance that you created.

View vSAN I/O Insight Metrics

I/O Insight performance metrics chart displays the metrics at the virtual disk level.

When I/O Insight is running, vSAN collects and displays the metrics for selected VMs, for a set duration. You can view the performance metrics for up to 90 days. The I/O Insight instances are automatically deleted after this period.

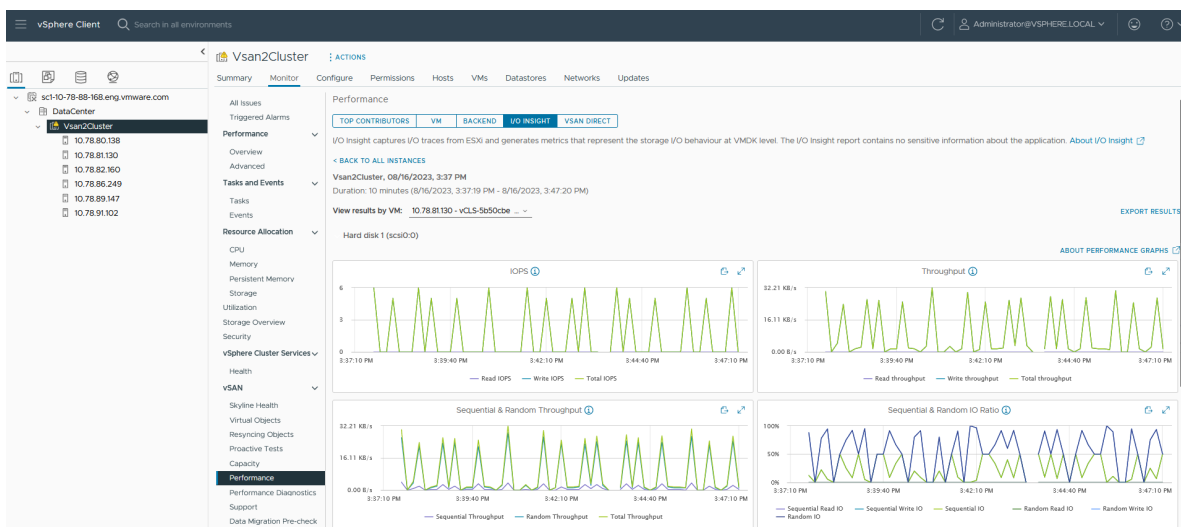
Procedure

- 1 Navigate to the vSAN cluster or host.

You can also access I/O Insight from the VM. Select the VM and navigate to **Monitor > vSAN > Performance > Virtual Disks**.

- 2 Click the **Monitor** tab.
- 3 Under **vSAN**, select **Performance**.
- 4 Select the **I/O Insight** tab. You can organize the instances based on time or hosts.

- 5 To view the metrics of an instance, click  and click **View Metrics**. You can optionally stop a running instance before completing the specified duration.

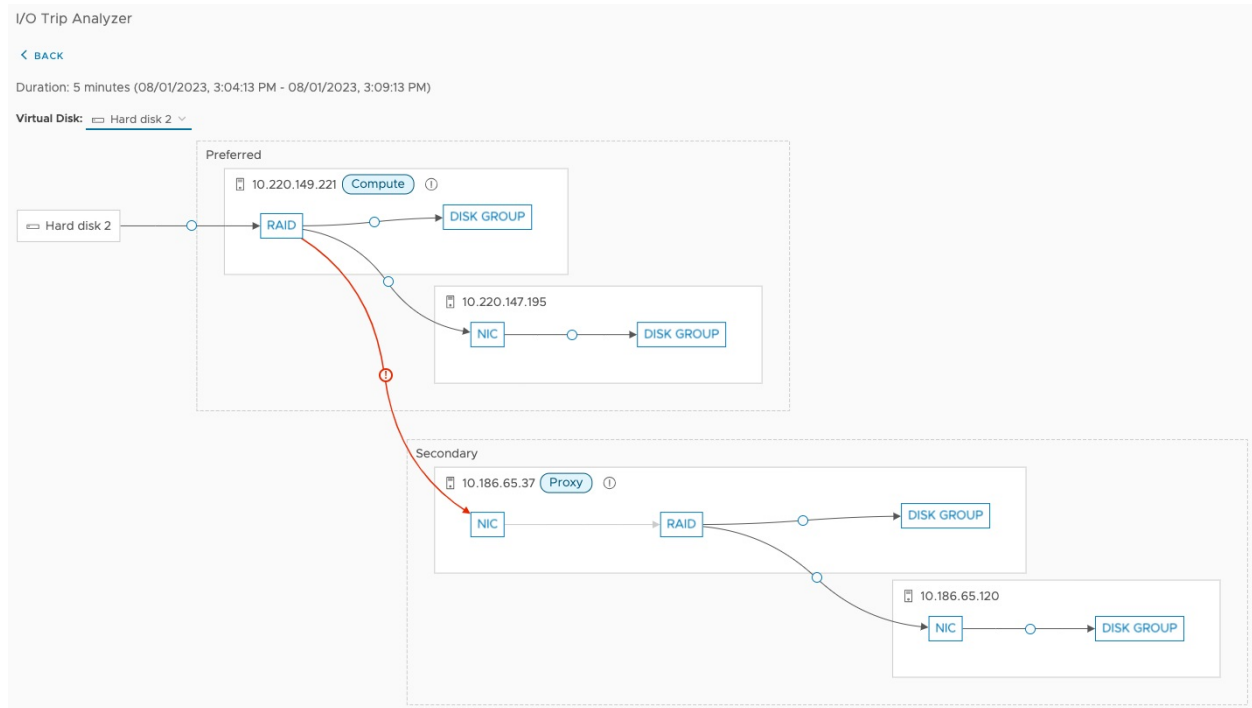


You can rerun an instance, and rename or delete the existing instances.

Use vSAN I/O Trip Analyzer

You can use vSAN I/O trip analyzer to diagnose the virtual machine I/O latency issues.

vSAN latency issues can be caused by outstanding I/Os, network hardware issues, network congestions, or disk slowness. The trip analyzer allows you to get the breakdown of the latencies at each layer of the vSAN stack. The topology diagram shows only the hosts with VM I/O traffic.



Note All the ESXi hosts and vCenter Server in the vSAN cluster must be running 7.0 Update 3 or later.

Using the I/O trip analyzer scheduler running 8.0 Update 1 or later, you can set the recurrence for I/O trip analyzer diagnostic operations. You can either set a one time occurrence or set the recurrence to later. On reaching the recurrence time, the scheduler automatically collects the results. You can view the results collected within 30 days.

Note The I/O trip analyzer supports stretched cluster and multiple VMs (maximum 8 VMs and 64 VMDKs) in one diagnostic run for a single cluster.

Prerequisites

The vSAN performance service must be enabled before you can run the I/O trip analyzer and view the test results.

Procedure

- 1 Navigate to the vSAN cluster, and select a VM.
- 2 Click the **Monitor** tab.
- 3 Under vSAN, select **I/O Trip Analyzer**.
- 4 Click **Run New Test**.
- 5 In the Run VM I/O Trip Analyzer Test, select the duration of the test.

- 6 (Optional) Select **Schedule for a future time** to schedule the test for a later time. You can either select **Start now** or enter a time based on your requirement in the **Custom time** field. Select the repeat options and click **Schedule**.

Note You can schedule only a single I/O trip analyzer per cluster. You can schedule another I/O trip analyzer after deleting the current scheduler. To delete a scheduler, click **Schedules > Delete**. You can also modify a schedule that you created. Click **Schedules > Edit**.

- 7 Click **RUN**. The trip analyzer test data is persisted and is available only for 30 days.

Note vSAN does not support I/O trip analyzer for virtual disks in a remote vSAN datastore.

- 8 Click **VIEW RESULT** to view the visualized I/O topology.
- 9 From the Virtual Disks drop-down, select the disk for which you want to view the I/O topology. You can also view the performance details of the network and the disk groups. Click the edge points of the topology to view the latency details.

Click the edge points of the topology to view the latency details. If there is a latency issue, click the red icon to focus on that area.

View vSAN Performance Metrics for Support Cases

Use the vSAN cluster performance metrics to monitor the performance of your cluster and determine the root cause of the performance issues.

You can use the vSAN Obfuscation Map to identify the obfuscated data sent to VMware. For more information on obfuscation map, see [View vSAN Obfuscation Map](#).

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.
- 3 Under vSAN, select **Support > Performance For Support**.
- 4 Select a performance dashboard from the drop-down menu.
- 5 Select hosts, disks, or NICs from the drop-down menu.
- 6 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 HClbench tool to run performance benchmark tests on the vSAN cluster, the time ranges of those tests appear in the drop-down menu.

7 Click **Show Results**.

vSAN displays performance charts for selected entities, such as IOPS, throughput, latency, congestions, and outstanding I/Os.

Using vSAN Performance Diagnostics

You can use vSAN performance diagnostics to improve the performance of your vSAN OSA 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 and to send data to VMware to analyze vSAN performance data.
- You must participate in the Customer Experience Improvement Program (CEIP).

Procedure

1 Navigate to the vSAN cluster.

2 Click the **Monitor** tab.

3 Under vSAN, select **Performance Diagnostics**.

4 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.

5 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.

6 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.

View vSAN Obfuscation Map

You can use the vSAN Obfuscation Map to identify the obfuscated data sent to VMware.

vSAN Obfuscation Map provides mapping of the obfuscated data sent to VMware as part of Customer Experience Improvement Program (CEIP) to facilitate communication during the Support Request process between vSAN user and VMware Global Support. Use notepad or any text editor to view the obfuscation map. For more information on obfuscation map, see the VMware knowledge base article at <https://kb.vmware.com/s/article/51120>.

Handling Failures and Troubleshooting vSAN

6

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.

Read the following topics next:

- [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](#)

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.

The log files, collected and packaged into a zip file, include the following:

- vCenter support bundle
- Host support bundle

The host support bundle in the cluster includes the following:

```
[ "Userworld:HostAgent", "Userworld:FDM",  
  "System:VMKernel", "System:ntp", "Storage:base", "Network:tcpip",  
  "Network:dvs", "Network:base", "Logs:System", "Storage:VSANMinimal",  
  "Storage:VSANHealth", "System:BaseMinimal", "Storage:VSANTraces"]
```


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. vSAN connects to the FTP port 21 or HTTPS port 443 of the target server with the domain name *vmware.com*, to automatically upload the support bundle.

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 OSA or vSAN ESA 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 storagepool list</code>	List storage pool claimed by vSAN ESA. This command is applicable only for vSAN ESA cluster.
<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.

Command	Description
<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 Cmmnds: 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.

Problem

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

Note Validation checks are also performed as a part of a host synchronization process.

Solution

If vCenter Server detects any configuration problems, it displays error messages. Use the following methods to fix vSAN configuration problems.

Table 6-1. 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 vSAN Alarms](#) and [Using the vSAN Default Alarms](#).

vSAN supports two types of failure states for components:

Table 6-2. 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 find how a failure in the cluster affects the virtual machine.

Table 6-3. 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 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 object state, storage policy, and vSAN UUID.
If the inventory object is Unhealthy, the vSphere Client indicates the reason for the unhealthy state in brackets.

Examine the Compliance of a Virtual Machine in vSAN

Use the vSphere 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 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 object state, 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 6-4. 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 <code>.vmtx</code> 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 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 finding	The Disk operation health finding 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 vSAN Cluster

When a magnetic disk or flash capacity device fails, vSAN evaluates the accessibility of the objects on the device.

vSAN 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 reProtection.</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 reprotect will occur.</p>

Storage Pool Device Is Not Accessible in vSAN ESA Cluster

When a storage pool device fails, vSAN evaluates the accessibility of the objects on the device.

vSAN 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

vSAN responds to the storage pool 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 reProtection.</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 reprotect 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.

vSAN 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 vSAN Cluster

If a host stops responding due to failure or reboot of the host, vSAN waits for the host to recover and rebuilds the components 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 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 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 vSAN Cluster

When the connectivity between the hosts in the cluster is lost, vSAN determines the active partition.

vSAN 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 vSAN Cluster

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

vSAN 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	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.

vSAN 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.

vSAN Stretched Cluster Failure Handling

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

Table 6-5. How vSAN 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 Skyline Health 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 findings verify your hardware against the *VMware Compatibility Guide*. For more information about using the vSAN Skyline Health, see [Chapter 3 Monitoring vSAN Skyline Health](#).

Examining Performance in 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 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 5 Monitoring vSAN Performance](#). For more information about using performance data in a vSAN cluster, see the *vSAN Troubleshooting Reference Manual*.

Network Misconfiguration Status in 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 the vSAN Cluster

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

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 `.vmtx` 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.

vSAN Stretched Cluster Configuration Error When Adding a Host

Before adding new hosts to a vSAN 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 vSAN 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.

vSAN Stretched Cluster Configuration Error When Using RVC to Add a Host

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

Problem

After you use the RVC tool to add a host to a vSAN 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 vSAN Stretched Cluster

Before adding or removing the witness host in a vSAN 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 vSAN 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 vSAN Cluster

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 Client. An Encryption health finding 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 in vSAN Cluster

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.

vSAN ESA contains a single storage pool of flash devices. Each flash device provides caching and capacity to the cluster. For more information on how the vSAN ESA is designed, see the *vSAN Planning and Deployment* guide.

Replace a Flash Caching Device on a Host in vSAN Cluster

You must replace a flash caching device if you detect a failure or when there is a disk group upgrade.

Removing the cache device removes the entire disk group from the vSAN cluster. 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 entire disk group that contains the flash caching device that you want to remove. vSAN does not allow you to remove the cache disk. To remove the cache disk, you must remove the entire disk group.
- 4 Click  and click **REMOVE**.
- 5 In the Remove Disk Group dialog box, select any of the following data migration mode to evacuate the data on the disks.
 - **Full data migration** - Transfers all the data available on the host to other hosts in the cluster.
 - **Ensure accessibility** - Transfers data available on the host to the other hosts in the cluster partially. During the data transfer, all virtual machines on the host remains accessible.
 - **No data migration** - There is no data transfer from the host. At this time, some objects might become inaccessible.

- 6 Click **GO TO PRE-CHECK** to find the impact on the cluster if the object is removed or placed in maintenance mode.
- 7 Click **REMOVE** to remove the disk group.

Results

vSAN removes the flash caching device along with the entire disk group from the cluster.

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.

For more information on creating a disk group, claiming storage devices, or adding devices to the disk group in the vSAN Cluster, see [Device Management in a vSAN Cluster](#).

Replace a Capacity Device in vSAN Cluster

You must 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 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.

You can also watch the video about how to replace a failed capacity device in vSAN.



(Replacing a Failed Drive in vSAN)

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 that the cluster contains enough space to migrate the data from the capacity device.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Configure** tab, click **Disk Management** under vSAN.
- 3 Select the flash capacity device or magnetic disk, and click **Remove Disk**.

Note You cannot remove a capacity device from the cluster with enabled deduplication and compression. You must remove the entire disk group. If you want to remove a disk group from a vSAN cluster with deduplication and compression enabled, see "Adding or Removing Disks with Deduplication and Compression Enabled" in *Administering VMware vSAN*.

- 4 In the Remove Disk dialog box, select **Full data migration** to transfer all the data available on the host to other hosts in the cluster.
- 5 Click **Go To Pre-Check** to find the impact on the cluster if the object is removed or placed in maintenance mode.
- 6 Click **Remove** to remove the capacity device.

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 Storage Pool Device in vSAN ESA Cluster

The storage pool represents the amount of capacity provided by the host to the vSAN datastore.

Each host's storage devices claimed by vSAN form a storage pool. All storage devices claimed by vSAN contribute to capacity and performance.

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 storage pool device, verify that the cluster contains enough space to migrate the data from the storage pool device.

Procedure

- 1 Navigate to the vSAN cluster.
- 2 On the **Configure** tab, click **Disk Management** under vSAN.
- 3 Select the storage pool device, and click **Remove Disk**.

- 4 In the Remove Disk dialog box, select **Full data migration** to transfer all the data available on the host to other hosts in the cluster.
- 5 Click **Go To Pre-Check** to find the impact on the cluster if the object is removed or placed in maintenance mode.
- 6 Click **Remove** to remove the storage pool device.

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.
- 3 Claim a disk using the vSAN cluster > **Configure** > **vSAN** > **Disk Management**.

Replace a Storage Controller in vSAN Cluster

You must replace a storage controller on a host if you detect a failure.

Procedure

- 1 Place the host into maintenance mode and power down the host.
- 2 Replace the failed card.
The replacement storage controller must have a supported firmware level listed in the *VMware Compatibility Guide*.
- 3 Power on the host.
- 4 Configure the card for passthrough mode. Refer to the vendor documentation for information about configuring the device.
- 5 Exit maintenance mode.

Remove a Device from a Host in vSAN Cluster 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
```

The following are the commands available for managing vSAN ESA cluster:

Table 6-6. vSAN ESA Commands

Command	Description
<code>esxcli vsan storagepool add</code>	Add physical disk for vSAN usage.
<code>esxcli vsan storagepool list</code>	List vSAN storage pool configuration.
<code>esxcli vsan storagepool mount</code>	Mount vSAN disk from storage pool.
<code>esxcli vsan storagepool rebuild</code>	Rebuild vSAN storage pool disks.
<code>esxcli vsan storagepool remove</code>	Remove physical disk from storage pool. Requires one --disk or --uuid param.
<code>esxcli vsan storagepool unmount</code>	Unmount vSAN disk from storage pool.

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.

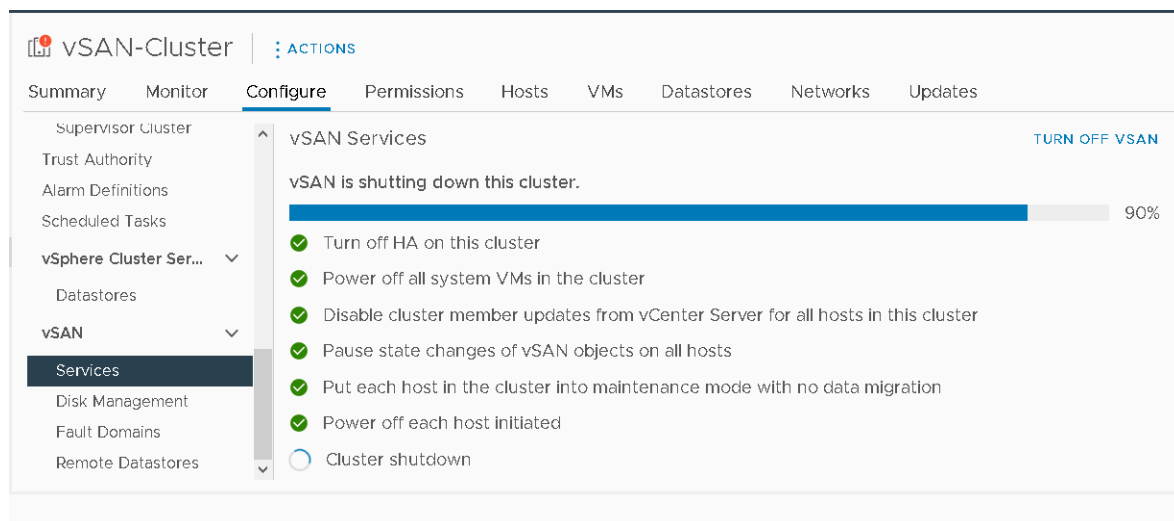
Shutting Down and Restarting the vSAN Cluster

7

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

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

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



Read the following topics next:

- [Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard](#)
- [Restart the vSAN Cluster](#)
- [Manually Shut Down and Restart the vSAN Cluster](#)

Shut Down the vSAN Cluster Using the Shutdown Cluster Wizard

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

The Shutdown Cluster Wizard is available with vSAN 7.0 Update 3 and later releases.

Note If you have a vSphere with Tanzu environment, you must follow the specified order when shutting down or starting up the components. For more information, see "Shutdown and Startup of VMware Cloud Foundation" in the *VMware Cloud Foundation Operations Guide*.

Procedure

1 Prepare the vSAN cluster for shutdown.

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

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

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

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

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

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

If vCenter Server appliance is deployed on the vSAN cluster, the Shutdown cluster wizard displays the vCenter Server notice. Note the IP address of the orchestration host, in case you need it during the cluster restart. If vCenter Server uses service VMs such as DNS or Active Directory, note them as exceptional VMs in the Shutdown cluster wizard.

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

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

5 Monitor the shutdown process.

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

What to do next

Restart the vSAN cluster. See [Restart the vSAN Cluster](#).

Restart the vSAN Cluster

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

Procedure

- 1 Power on the cluster hosts.

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

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

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

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

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

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

Manually Shut Down and Restart the vSAN Cluster

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

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

Note If you have a vSphere with Tanzu environment, you must follow the specified order when shutting down or starting up the components. For more information, see "Shutdown and Startup of VMware Cloud Foundation" in the *VMware Cloud Foundation Operations Guide*.

Procedure

- 1 Shut down the vSAN cluster.

- a Check the vSAN Skyline Health to confirm that the cluster is healthy.

- b Power off all virtual machines (VMs) running in the vSAN cluster, if vCenter Server is not hosted on the cluster. If vCenter Server is hosted in the vSAN cluster, do not power off the vCenter Server VM or service VMs (such as DNS, Active Directory) used by vCenter Server.

- c If vSAN file service is enabled in the vSAN cluster, you must deactivate the file service. Deactivating the vSAN file service removes the empty file service domain. If you want to retain the empty file service domain after restarting the vSAN cluster, you must create an NFS or SMB file share before deactivating the vSAN file service.

- d Click the **Configure** tab and turn off HA. As a result, the cluster does not register host shutdowns as failures.

For vSphere 7.0 U1 and later, enable vCLS retreat mode. For more information, see the VMware knowledge base article at <https://kb.vmware.com/s/article/80472>.

- e Verify that all resynchronization tasks are complete.

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

- f If vCenter Server is hosted on the vSAN cluster, power off the vCenter Server VM.

Make a note of the host that runs the vCenter Server VM. It is the host where you must restart the vCenter Server VM.

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

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

- h Log in to any host in the cluster other than the witness host.

- i Run the following command only on that host. If you run the command on multiple hosts concurrently, it may cause a race condition causing unexpected results.

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

The command returns and prints the following:

```
Cluster preparation is done.
```

Note

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

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

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

Perform this step on all the hosts.

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

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

2 Restart the vSAN cluster.

a Power on the ESXi hosts.

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

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

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

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

Perform this step on all the hosts.

c Log in to one of the hosts in the cluster other than the witness host.

d Run the following command only on that host. If you run the command on multiple hosts concurrently, it may cause a race condition causing unexpected results.

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

The command returns and prints the following:

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

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

```
esxcli vsan cluster get
```

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

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

g Restart the vCenter Server VM if it is powered off. Wait for the vCenter Server VM to be powered up and running. To deactivate vCLS retreat mode, see the VMware knowledge base article at <https://kb.vmware.com/s/article/80472>.

h Verify again that all the hosts are participating in the vSAN cluster by running the following command on each host.

```
esxcli vsan cluster get
```

i Restart the remaining VMs through vCenter Server.

j Check the vSAN Skyline Health and resolve any outstanding issues.

- k (Optional) Enable vSAN file service.
- l (Optional) If the vSAN cluster has vSphere Availability enabled, you must manually restart vSphere Availability to avoid the following error: `Cannot find vSphere HA master agent`.

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

- 1 **Configure > Services > vSphere Availability > EDIT > Disable vSphere HA**
- 2 **Configure > Services > vSphere Availability > EDIT > Enable vSphere HA**
- 3 If there are unhealthy or disconnected hosts in the cluster, recover or remove the hosts from the vSAN cluster. If vCenter Server uses service VMs such as DNS or Active Directory, note them as exceptional VMs in the Shutdown cluster wizard.

Retry the above commands only after the vSAN Skyline Health shows all available hosts in the green state.

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

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

```
reboot_helper.py recover
```

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

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

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

What to do next

Restart the vSAN cluster. See [Restart the vSAN Cluster](#).