vSAN Availability Technologies

Resilience Capabilities for vSAN 8 U3 and VMware Cloud Foundation 5.2

February 7, 2025

vmware by Broadcom

Table of Contents

Introduction

Reliable and resilient access to applications and their data is a top priority for nearly all IT organizations. Virtualization afforded every environment to provide high levels of availability for virtual machines (VMs) and the data they serve. When using a classic three-tier architecture consisting of compute resources (vSphere hosts), networking, and storage, this meant that data availability was the responsibility of the storage solution. This typically was in the form of a traditional modular storage array using its proprietary hardware and dedicated storage fabric.

Hyperconvergence changed this model drastically. vSAN HCI clusters aggregate storage resources into the same hosts that provided the compute resources for the cluster. The hypervisor provided the distributed storage services exclusively through its software stack using commodity hardware. vSAN also offers a disaggregated approach to storage. This deployment option gives you the power of a fully distributed storage solution built into the hypervisor, but disaggregated away from the compute resources, so it can provide a centralized shared storage solution for one or many vSphere clusters powering VMware Cloud Foundation (VCF).

Data Availability Concepts

Maintaining availability of data means that a system must understand what types of failures may occur in a storage system. Some failures such as a storage device sending a failure code are recognizable. Other failures, such as a host outage are not as easily recognizable because its absence is the only indicator of a known issue. Defining what a failure is can also be difficult, as a failure may not always occur in a clear and a distinct way. For example, an interruption in a network link between hosts may be considered tolerable if for a brief period but deemed a failure if it exceeds a certain time threshold. When a failure is determined, a system must then be able to repair the condition or ameliorate it to the best of its ability at that moment.

Any type of enterprise storage system must have all the proper intelligence and logic to handle these matters in a reliable and elegant manner. A distributed storage system like vSAN must factor in additional considerations due its distributed architecture. vSAN factors in these complex considerations into a robust solution that maintains availability under a wide range of conditions.

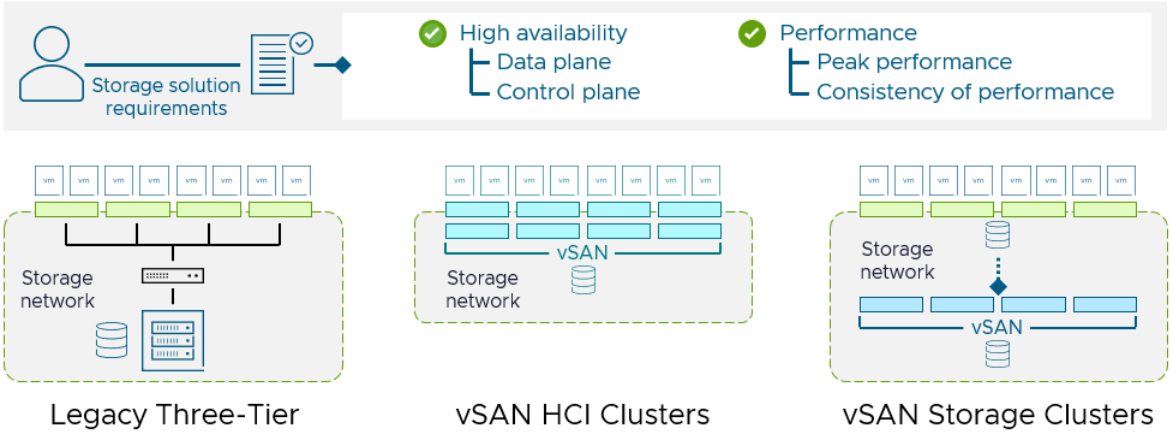

Figure. Core competencies of enterprise storage systems.

Scope of Topics

This document provides details on how vSAN provides high levels of data availability. The content describes general behavior of vSAN. Many of the topics apply to both the Original Storage Architecture (OSA) and the much more advanced Express Storage Architecture (ESA), with the latter being emphasized more in the configuration and failure scenarios. The information will also typically apply to aggregated vSAN HCI clusters as well as disaggregated vSAN storage clusters.

vSAN also has a highly durable control plane that maintains management of the system even under extreme circumstances. Availability of the control plane is not a focus of this document but may be discussed when appropriate. The concepts

discussed in this document generally apply to a stand-alone vSAN environment, as well as a VCF environment. **For VCF environments, please refer to the Administration Guide For VMware Cloud Foundation 5.2 for guidance and requirements specific to VCF.**

vSAN Architectural Basics

vSAN is a storage solution that provides shared storage for VMs running vSphere. It is a distributed storage solution, meaning that data is intelligently and automatically sprinkled across hosts that comprise the vSAN cluster to ensure data can remain available in the event of some known or unknown failure. It aggregates the storage devices across the hosts in a vSAN cluster and provides a single datastore that is easy to operate. This is quite different than shared storage using a traditional storage array, that provides resilience of storage through redundancy of discrete devices within the storage array, such as disks, storage controllers, and network uplinks. It is this distinction that helps make vSAN's distributed approach to storage much more scalable and resilient in the face of growing demand of data and workloads.

A standard vSAN cluster requires a minimum 3 hosts, although we recommend 4 as a minimum in most cases.. A single vSAN cluster can scale up to 64 hosts and be tailored to a size that best fits the needs of your organization. "vSAN Cluster Design – Large Clusters versus Small Clusters" details how you can use this flexibility to your advantage. Pair this capability with vCenter Server's ability to manage many clusters within a VCF workload domain, and many workload domains that can be a part of a single VCF instance, and you have a platform that makes compute and storage resources extremely scalable.

vSAN OSA versus vSAN ESA

The vSAN Express Storage Architecture (ESA) was announced in 2022, and offers a dramatic improvement in performance, efficiency, resilience, and ease of use. Some of the improvements in resilience comes from a change in how the two architectures use storage devices. For example, the Original Storage Architecture (OSA) used a construct known as a disk group. This was a unit of storage resources that allowed customers to pair 1-7 value-based capacity devices with a single higher performing caching device. This provided modest levels of performance using many lower performing storage devices. Unfortunately, this construct meant that a failure of a discrete storage device might impact the entire disk group. This larger blast zone meant that a single failure could impact more data, which would require more data would need to be resynchronized in order to regain the prescribed level of resilience.

vSAN ESA is different. Unlike the OSA which might be thought of as a two-tier architecture, vSAN ESA uses a single tier. This provides a smaller boundary of failure in the event of an unexpected issue with a storage device and yields much better efficiency and time needed to regain levels of compliance. For more information, see the post: "The Impact of a Storage Device Failure in vSAN ESA versus OSA."

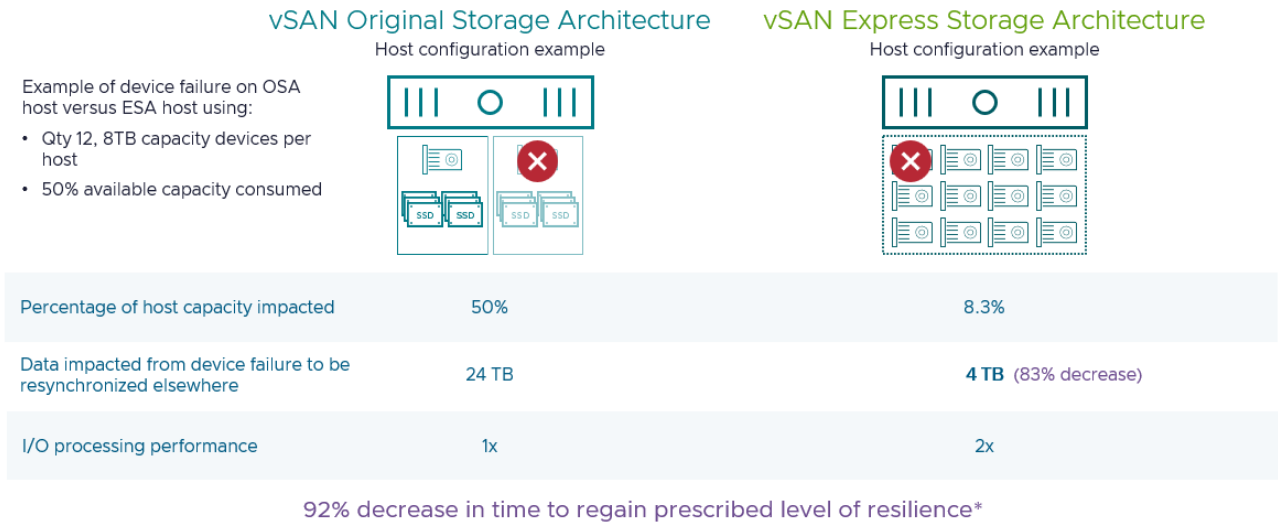

Figure. Comparing the impact of a storage device failure in vSAN OSA versus vSAN ESA.

vSAN Objects and Components

Unlike a clustered file system like VMFS used in a datastore residing on a storage array, vSAN uses a data structure that is analogous to an object store. It does this for several reasons, and are detailed in the post: "vSAN Objects and Components Revisited." Objects are the entities that make up a given VM, such as virtual disks (VMDKs) and configurations. They are the smallest unit of management within vSAN and give administrators capabilities not found with traditional storage. These objects can be as large as 62TB in size.

Components are smaller chunks of data that comprise an object. While components will be discussed in this document, they are simply an implementation detail of vSAN, and not a manageable entity. They are only discussed here and exposed in the UI because it can help an administrator better understand the concepts of object availability. They have their own limits in terms of the maximum size, and the maximum number of components that can reside on a host.

The components of an object in vSAN ESA are a concatenation of two types:

- **Performance leg components.** These are small components mirrored across 2 or more hosts and intended to temporarily store data and metadata in preparation to be written in a space efficient manner to the capacity leg.
- **Capacity leg components.** These are the components distributed across the number of hosts necessary for the assigned data placement type (such as FTT=2 using RAID-6)

All these components play a role in determining availability of data during failure conditions through a vote count of components, as described later in this document.

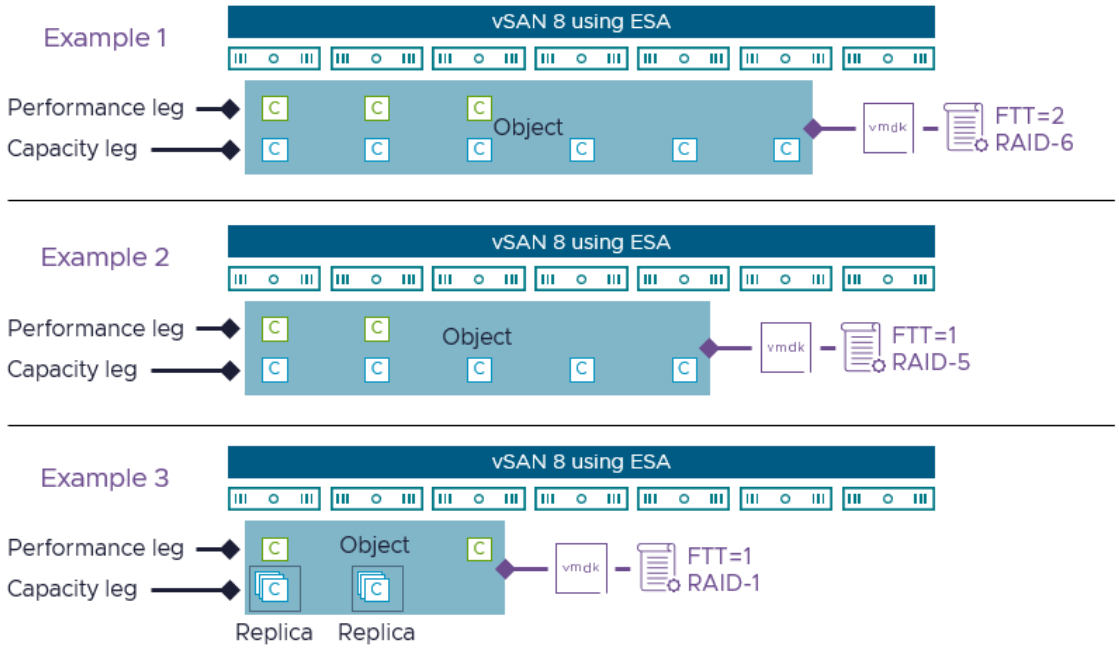

Figure. Concatenated components that make up an object in vSAN ESA.

vSAN's Cluster Level Object Manager (CLOM) automatically makes the placement decisions in order to satisfy the requirements of the assigned policy. This is what helps make the data resilient. The image below shows how a VM's virtual disk object is placed across a series of hosts. It is assigned a storage policy of FTT=2 using RAID-6 erasure coding. This results in a concatenation of 3 mirrored components in the performance leg that can tolerate up to 2 failures, and 6 components that make up a RAID-6 stripe with parity that can tolerate up to 2 failures.

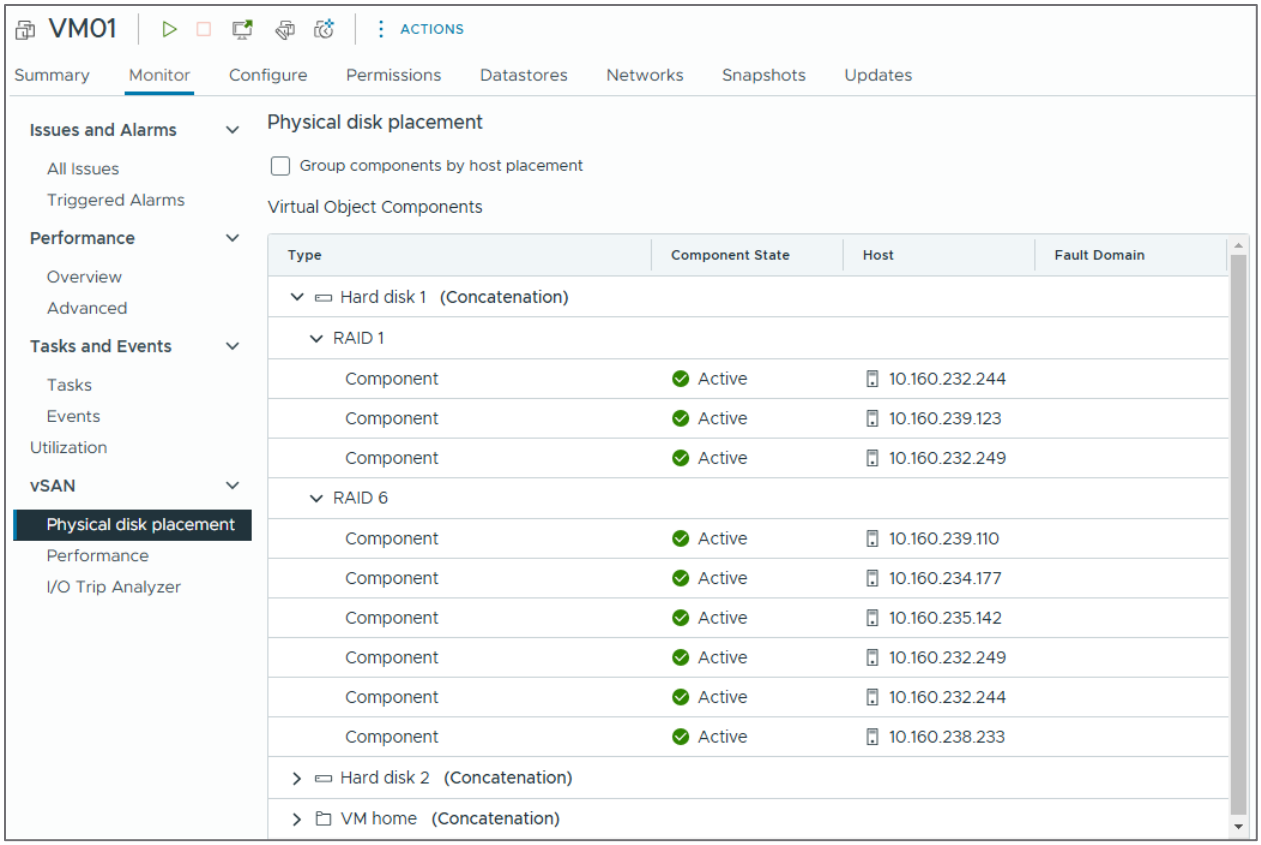

Figure. Virtual disk object for a VM in vSAN ESA, using a FTT=2 RAID-6 erasure code.

Components typically contain data payload and metadata of an object. In special circumstances, there may be a "witness" component. This is a very small component that simply contains metadata for the purpose of understanding the object availability. Witness components were common when using RAID-1 in vSAN OSA but are no longer used in ESA except for stretched clusters, and 2-Node clusters, where the witness components reside on a witness host appliance.

*Recommendation: There is very little need to spend time trying to fully understand components and their relation to availability. ESA uses components in much more sophisticated ways compared to the OSA and can be quite confusing to decipher in complex failure scenarios. However, **vSAN handles this logic so that you don't have to**. The information shared in this document is for informational purposes only.*

Storage Policies and Component Placement

Storage polices are an integral part of vSAN data management, as they define the desired state, or outcome the administrator wishes to have for one or more VMs. For example, one could set a higher level of resilience their business-critical VMs and a lower level of resilience for their test/lab VMs.

vSAN looks at the outcome defined in the storage policy, and not only applies it to the VMs associated with it but will constantly look to ensure that it is compliant with the prescribed policy. And if not, it will make the appropriate changes automatically to regain the prescribed level of resilience. For example, a VM using a level of failure to tolerate (FTT) of 2 using RAID-6 erasure code will always stripe the data with parity across at least 6 hosts so that if two of the hosts holding that object fail, the data will remain available. **vSAN has internal anti-affinity mechanisms in place so these components of an object will not reside on the same host.** If there is a failure, vSAN will repair the data structure to regain its prescribed level of resilience, assuming it has sufficient hosts to do so.

The vSAN Interactive Infographic is a helpful tool to understand how configuration changes and failure conditions will affect data resilience and availability in vSAN. It has been updated for vSAN ESA and includes many of the most common failure

scenarios. For simplicity, some of the illustrations and animations have been simplified for clarity. The number and type of components may deviate from what is illustrated.

Storage Policy Resilience Settings

As noted above, storage polices allow administrators to define the level of resilience desired for one or more objects. This is expressed as a storage policy rule by its level of failures to tolerate, or "FTT." A failure in this context would relate to what vSAN defines as a fault domain. In a standard vSAN cluster, this would be a host. In a stretched cluster, this would be a site, and in a standard vSAN cluster using the "Fault Domains" feature, this could be a collection of hosts such as a room or a rack.

The FTT setting defines the level of resilience desired for the object. An FTT of 1, or FTT=1 allows for one of the hosts holding that object to fail while remaining available. An FTT=2 allows for two of the hosts. **Note that the FTT setting only applies to failures for the hosts that the object resides on, NOT the total number of failures within a cluster**. This is why a vSAN cluster improves in its ability to maintain data availability as a cluster size grows, versus other distributed storage systems that tend to become more fragile as cluster sizes grow.

Figure. Failures to tolerate (FTT) options in vSAN.

vSAN ESA makes storage policies even easier. An optional "Auto-Policy Management" feature allows you to let vSAN determine the most appropriate storage policy based on the characteristics of the cluster. For more information, see the post: "Auto-Policy Management Capabilities with the ESA in vSAN 8 U1."

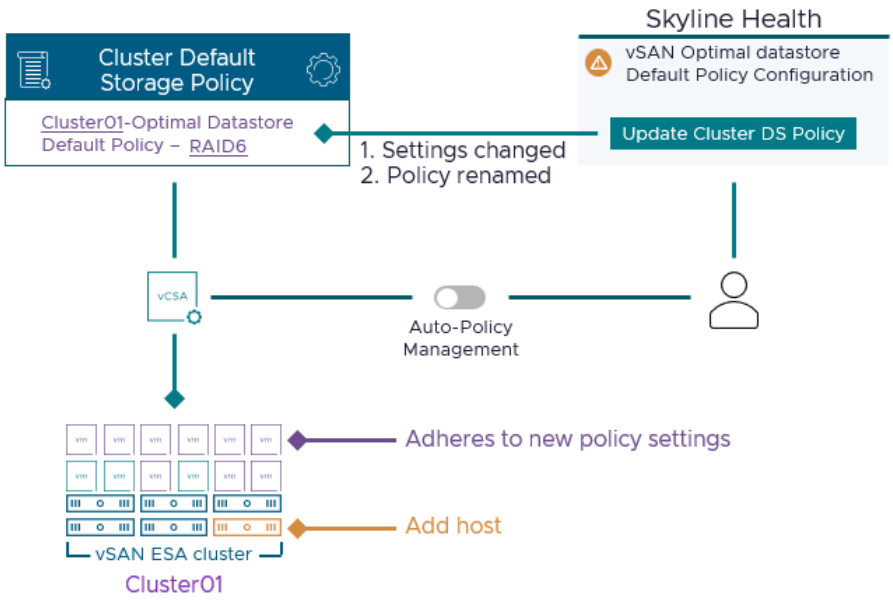

Figure. Auto-Policy Management in vSAN ESA.

Storage Policy Data Placement Scheme

While the storage policy rule "outcome" of FTT=1 or FTT=2 establishes the level of resilience desired, the description, **the data placement scheme defines how it is achieved.** For example, the "Failures to tolerate" rule may give options such as:

- 1 failure – RAID-1 (Mirroring)
- 1 failure – RAID-5 (Erasure Coding)
- 2 Failures – RAID-1 Mirroring
- 2 Failures – RAID-6 (Erasure Coding)

A RAID-1 mirror means that vSAN simply keeps a synchronous copy of the object data somewhere else on the cluster to provide the prescribed level of resilience, such as FTT=1. RAID-5 or RAID-6 6 indicates the use of an erasure code, which is the ability for vSAN to store data in components in an intelligent stripe with parity across a series of hosts to provide the prescribed level of resilience. Erasure codes are much more space efficient than simple mirroring but may require more hosts to achieve the desired level of resilience.

In the older vSAN OSA, data mirroring using RAID-1 was the common method for resilience. This offered the ability to maintain resilience without the cost in performance of the much more space efficient RAID-5/6 erasure coding. However, vSAN ESA can store data using RAID-5/6 erasure coding without any performance penalty. Thus, RAID-1 mirroring should not be used in anything other than site-level resilience for stretched clusters, and host level resilience in 2-Node clusters. For more information, see the post: "RAID-5/6 with the Performance of RAID-1 using the vSAN Express Storage Architecture."

When using RAID-5/6 erasure coding, each component in the performance leg contains a combination of data and parity data. There is no dedicated parity component. There is also not any dedicated witness component.

Each data placement scheme has a minimum number of hosts required in order store the data in a non-failure state. For example, FTT=1 using RAID-1 mirroring requires just three hosts, while FTT=2 using RAID-6 erasure coding requires a minimum of 6 hosts. The illustration below shows the data placement scheme options in vSAN ESA. The hosts in blue represent the minimum required for the storage policy, and the additional host in purple represents the recommended minimum so that data can regain its prescribed level of resilience in the event of a sustained host failure.

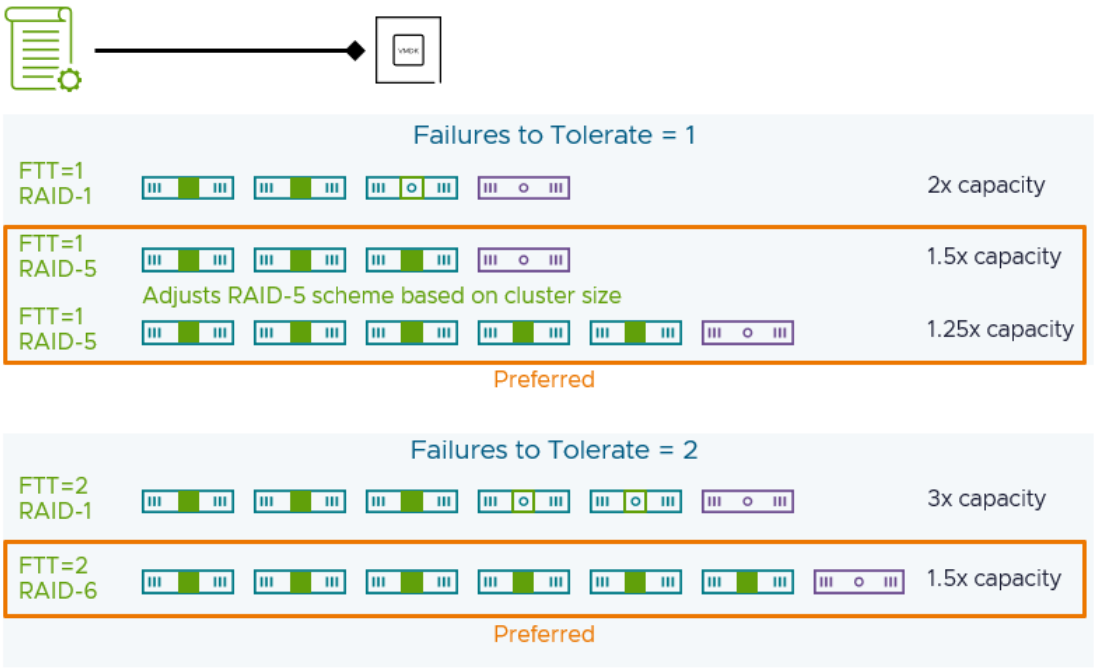

Figure. Data placement schemes and the minimum number of hosts required in vSAN ESA.

vSAN ESA features an adaptive RAID-5 erasure code that will change its data structure based on the number of hosts in the cluster. For clusters below 6 hosts, selecting a policy of FTT=1 using RAID-5 will stripe the data with parity across 3 hosts. For clusters with 6 or more hosts, selecting the same policy rule will stripe the data with parity across 5 hosts. For more information, see the post: "Adaptive RAID-5 Erasure Coding with the ESA in vSAN 8."

*Recommendation: Use the optional **Auto-Policy Management** feature in your ESA clusters. This will let the system decide the optimal storage policy based on the characteristics of your cluster. If you choose to do so, and it is on a cluster of 6 or fewer hosts, you may wish to **disable the "Host Rebuild Reserve"** option in the Reserved Capacity feature. This will help ensure that you are running at the highest levels of resilience. For more information, see the post "Auto-Policy Management Capabilities with the ESA in vSAN 8 U1."*

Fault Domains

vSAN uses a construct known as a "fault domain" to help it distribute data in a resilient way. By default, vSAN treats each host as a fault domain, as shown in the image below. In a standard, single site vSAN cluster, this helps keep data available in the event of a discrete host failure.

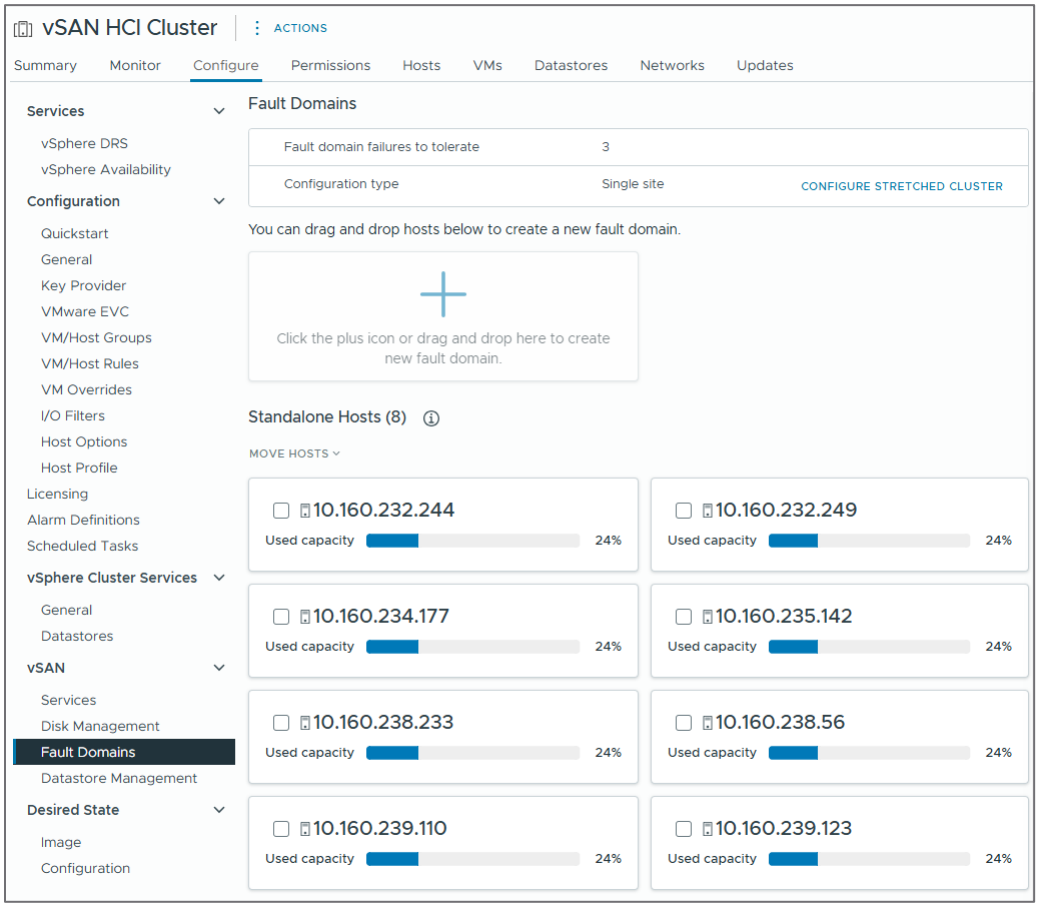

Figure. vSAN treating each host as its own fault domain.

This "Fault Domains" feature can be used to define a group of hosts to create a logical boundary of failure. It can use this information to write the data in such a way that if that defined fault domain is offline, the data will remain available. The manual creation of fault domains are often used to create groups of hosts in a cluster that represent their placement in a rack. This helps create rack-level resilience for a single site vSAN cluster. For more guidance on the use of vSAN Fault Domains, see the post: "Design and Operation Considerations when using vSAN Fault Domains." The post "Using Fault Domains in vSAN ESA" will provide additional guidance as it relates to vSAN ESA.

vSAN stretched clusters use a form of this construct, where the two data sites and the witness site are each defined as a fault domain. 2-Node vSAN clusters also use this method where the two data hosts and the witness site are each defined as a fault domain. The establishment of fault domains is automatically created for you when you select either one of those deployment types.

Unlike a traditional vSAN cluster, stretched clusters and 2-Node clusters allow you to establish a secondary level of resilience. In the case of a stretched cluster, this can help maintain data availability not only after a full site failure, but additional failures within the remaining site. For more information including an exhaustive list of failure handling scenarios with vSAN stretched clusters, see the "vSAN Stretched Cluster Guide."

Recommendation: Build a cluster with at least one more host than the minimum required by your storage policy. Upon a sustained outage or maintenance of a host, this will allow vSAN to automatically repair the data impacted by the failed host and reestablish the prescribed level of resilience assigned by the storage policy. If there is not an available host to rebuild the data to, the object will be available, but will be in a degraded state, and may not have a sufficient resilience setting to maintain availability of that object upon another failure.

vSAN Storage Clusters

vSAN can be deployed as an aggregated vSAN cluster, where compute and storage resources are provided by the same cluster. It can also be deployed as a vSAN storage cluster (previously known as "vSAN Max"). Much like a traditional storage array, a vSAN storage cluster can serve as a centralized shared storage solution for vSphere clusters. For more information, see the post: "vSAN Max and the Advantage of Scalability."

On the topic of availability, many of the same concepts and considerations for vSAN HCI clusters apply to vSAN storage clusters. See the "Disaggregated Storage using vSAN Storage Clusters" section of the vSAN FAQs, and the "vSAN Storage Clusters Design and Operational Guidance" document for guidance and recommendations.

Determining Availability of Data

Components of an object help determine the availability of an object. This is achieved through a vote count of components. This vote count of components helps determine if there is enough data to deem the object as available. This threshold is known as a "quorum" and is a safety mechanism to help vSAN determine under failure conditions when an object should be available or when it should be marked as unavailable to preserve its integrity.

An object must have more than 50% of its components accessible to meet quorum and deem the object as available. While most of the time each components have a value of "1" in vSAN's quorum voting mechanism, it may assign a value of two for some objects that have an even number of components. This is an implementation detail, and generally not necessary to know.

vSAN stretched clusters and vSAN 2-Node clusters have a unique ability to change the vote count of components in very specific situations, where a data site is offline due to a planned or unplanned event, followed by the witness site being offline for a planned or unplanned event. Assuming sufficient time between outages, vSAN will recalculate the votes to ensure that the one remaining site will remain available, even though a data site and a witness site is unavailable. For more information, see the "Adaptive Quorum Control" topic in the vSAN Stretched Cluster Guide.

vSAN is continuously monitoring the availability of all data and depending on events in a cluster (such as a transient network error, or a host failure), will make the determination on how to proceed with regaining its prescribed level of resilience.

Component States

vSAN components will typically show one of the following states described in more detail below:

Active

The most common state is "Active" which means the component is accessible and up to date.

Figure. Active components of an object in vSAN.

Reconfiguring

This state is usually found when a change to a storage policy is made, or a new storage policy is assigned to an object.

Figure. Component in a reconfiguration state for an object in vSAN.

Absent

This state represents components that can no longer be accessed. It does not define why it can't be accesses, but rather that they are simply not available. This is like an all paths down "APD" event with traditional storage. For example, in an 8-host cluster, imagine an object using a storage policy of FTT=2 using RAID-6. The data with parity would be spread across 6 of the hosts. If power was lost on one of the hosts storing a portion of that object, that specific component would be marked as "absent." Since vSAN is unaware of the reason for its absence, it will wait for a period for the component to be marked as active. This helps accommodate transient outages such as a temporary network error, or a host restart. Once it exceeds this time (the default is 60 minutes), it will begin to rebuild the component somewhere else in the cluster, assuming another previously unused fault domain (host) is available.

Figure. Absent component of an object in vSAN.

Degraded

A degraded component state is typically the result of the system receiving SCSI sense codes that indicate a device failure. This type of permanent device loss (PDL) will initiate rebuild or repair activities immediately, since the failure is known. Note that pulling a device to test availability would only simulate an APD event described above, and not a PDL described here, since the system never received a sense code.

Stale

Stale components, or "Active – Stale" as shown in the UI represent components that are available but are outdated. vSAN use sequence numbers to verify a component has the latest updates. This sequence number is normally kept consistent across components that make up an object. When a change to the object occurs, the data is written to both replicas, for example, and the sequence number is updated for the components. If a component is active, but its sequence number is different/older than the current sequence number for the object, the component is marked as Active - Stale. This usually occurs when the components of an object go offline and come back online concurrently at different times. Those stale components have an older sequence number and represent that they may not reflect the latest updated data. To preserve data integrity, vSAN marks these as "Active – Stale" until the components can be updated with the same sequence number.

Figure. Active-stale component of an object in vSAN.

Recommendation: Avoid changing the default rebuild timer setting of 60 minutes. A 1 hour timeout value provides a good balance of avoiding unnecessary rebuilds while minimizing the risk of downtime. If a situation occurs where a timelier rebuild is desired, it is possible to trigger a rebuild using the Repair Objects Immediately in the vSAN Health Check user interface.

Attempting to take a snapshot of a VM in vSAN OSA will fail if the storage policy for that VM is not fully compliant. This is not an issue in vSAN ESA.

Availability Handling

The storage of data across hosts in a resilient way is what allows vSAN to maintain availability under a wide range of conditions. For a more complete and interactive understanding how configuration changes and failure conditions will impact the availability of data in vSAN, see the vSAN Interactive Infographic. It has been updated for vSAN ESA and includes many of the most common failure scenarios. For simplicity, some of the illustrations and animations on the infographic have been simplified for clarity. The number and type of components may deviate from what is illustrated.

The event handling below shows outcomes of:

- Data available and resilient to its prescribe level of resilience
- Data available and degraded to a lesser level of resilience
- Data available but with no level of resilience
- Data unavailable

Unavailability of data does not mean "data loss." It simply refers to the state of availability for the VM to use or access the data. vSAN's data path always commits data to persistent media using advanced consensus algorithms to ensure data consistency. It does this prior to ever responding to the guest VMs with a write acknowledgement.

When discussing scenarios of object availability in vSAN, a failure only applies to the hosts that the object resides on, NOT the total number of failures wihin a cluster. An object, such as a VMDK is a relatively small boundary of data. vSAN will only place this on a sufficient number of hosts to comply with the storage policy resilience setting. Therefore, other failures that occur within a cluster will have no impact on the object if it does not use the other hosts that failed.

Availability Handling for an Object Assigned FTT=2 using RAID-6

This group of scenarios will describe the handling of an object on a vSAN ESA cluster with a storage policy that has FTT=2 using RAID-6. A RAID-6 object will have it's capacity leg components spread across 6 hosts, and its performance leg components spread across 3 hosts. FTT=2 using RAID-6 will be the most common storage policy in a standard, single site vSAN ESA cluster with 6 or more hosts.

Example 1: One Host Offline in a 7 Host Cluster

In this first example, a host is goes offline. The object remains available, but in a degraded state with a reduced effective resilience of FTT=1. After a 60 minute delay, that portion of the stripe with parity will be rebuilt on the host that does not contain the object. It will then regain its full prescribed level of resilience of FTT=2.

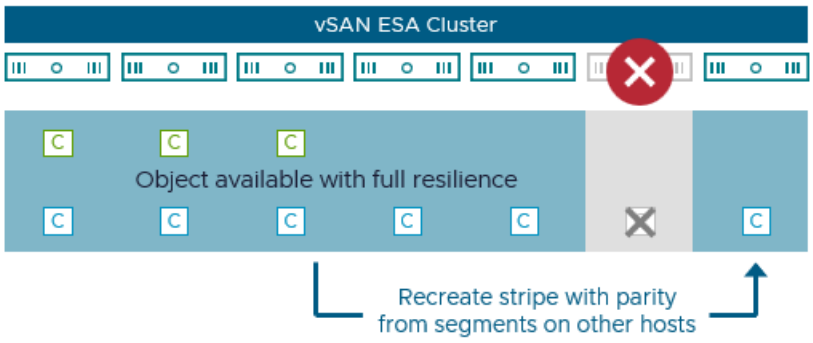

Figure. Regaining prescribed level of resilience after a host is offline.

Example 2: Two Hosts Offline in a 7 Host Cluster

In this second example, an additional host goes offline. The object remains available, but in a degraded state with a reduced effective resilience of FTT=1. The object cannot regain its prescribed level of resilience because there are not any hosts (fault domains) as a rebuild target.

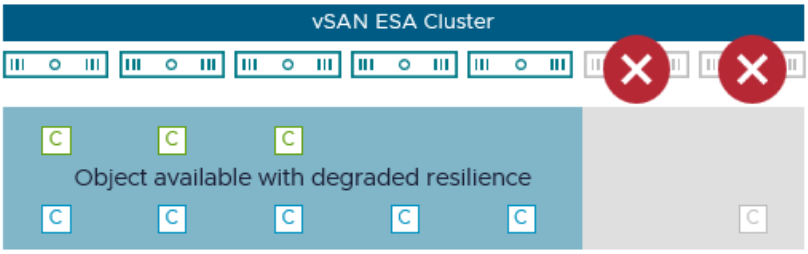

Figure. RAID-6 object available but with degraded level of resilience..

Example 3: Three Hosts Offline in a 7 host Cluster

In this third example, yet another host goes offline. The object remains available, but in a degraded state with no resilience - equivalent to an FTT=0. The object cannot regain its prescribed level of resilience because there are not any hosts (fault domains) as a rebuild target. It cannot remain any more host failures beyond what is shown.

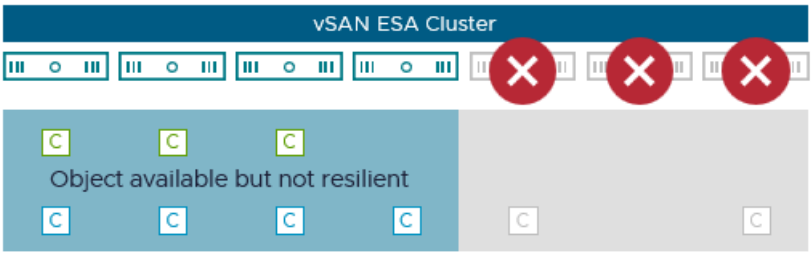

Figure. RAID-6 object available but with no resilience.

Example 4: Four Hosts Offline in a 7 host Cluster

In this fourth example, yet another host goes offline. The object is no longer available because a RAID-6 erasure code creates a stripe with double-parity across 6 hosts. It can tolerate 2 failures, but must have 4 of the 6 hosts available, otherwise the object will be unavailable.

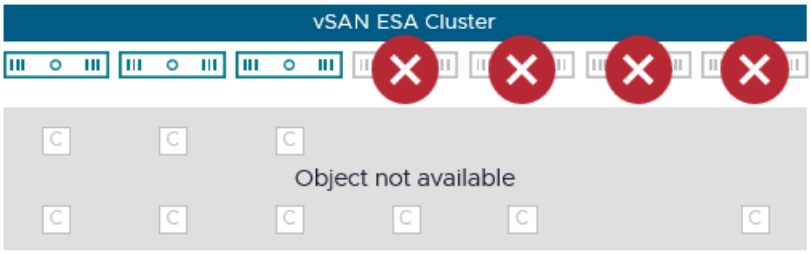

Figure. RAID-6 object unavailable.

The scenarios above demonstrate why we recommend having at least one more host than the minimum required by your desired storage policy resilience setting. It allows for for storage policies to regain their prescribed levels of resilience during any sustained outage of a host.

Durability of Data to Account for Multiple Events

Under times of planned or unplanned events, such as a host maintenance mode event or host failure, vSAN may create what are known as "durability components." These are special components that contain only updated blocks at the time a planned or unplanned outage occurred. For example, as illustrated in the figure below, imagine a VM object in a 4 host vSAN ESA cluster using a storage policy assigned with FTT=1 using RAID-5. In a normal operating state, the data with parity is striped across three hosts. Lets imagine the third host from the left is placed into maintenance mode using "Ensure accessibility." The object will remain available but no longer be resilient. At this time, vSAN creates a temporary durability component to capture all of the latest writes, in addition to the writes captures in the remaining stripe. If another host fails at this time, the object would no longer be available, but if the host originally placed into maintenance mode was brought back online, the delta changes in the durability component would be merged with the stripe, and the object would be available again.

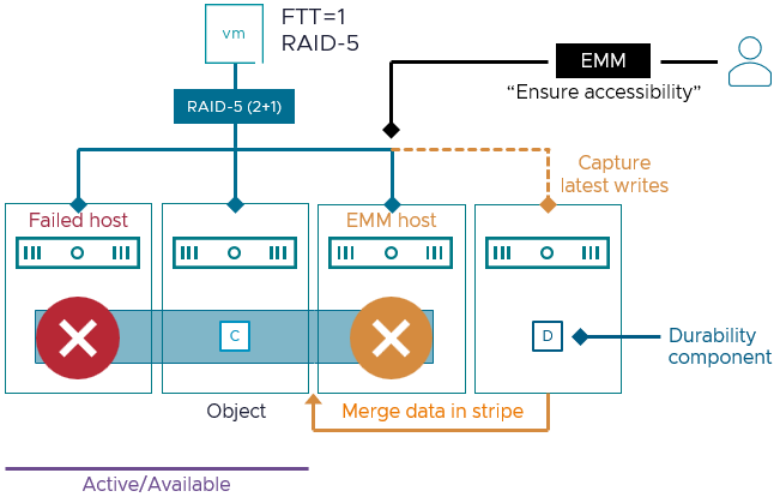

Figure. The use of durability components during planned or unplanned events.

We can see how this looks in practice. The UI shows that a "RAID_D" durability components were created at the time one of the hosts was entered into maintenance mode. One durability component represents missing component on the performance leg, while the other represents the missing component on the capacity leg.

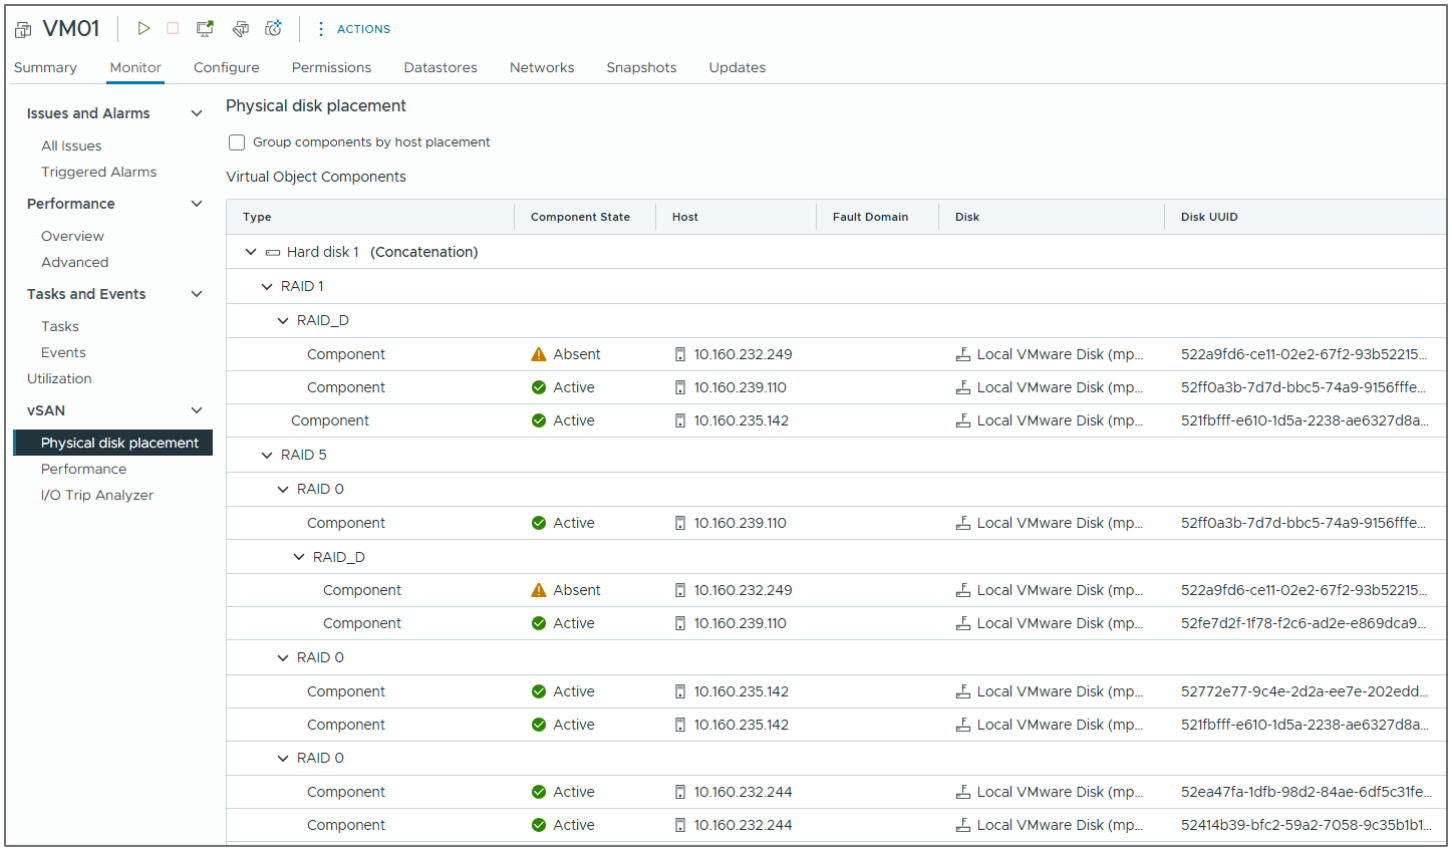

Figure. Durability components created during a maintenance mode event for an object using RAID-5.

Durability components are not used in any events other than maintenance or failure conditions. When they are no longer needed for an object, vSAN discards them.

The Role of the Network with vSAN

Since vSAN is a distributed storage system, it relies on the network to write the data in a resilient way and read the data in a reliable manner. This includes I/O from guest VMs, as well as "back-end" I/O activity that occurs to rebalance and repair data. Any transient or permanent interruptions in the network will impede vSAN's ability to write or read the data requested. The network that is used for vSAN is just as important as a dedicated storage fabric used with a traditional storage array.

Transport Method

vSAN uses TCP (unicast traffic) over Ethernet to transmit I/O payload across the hosts that comprise the cluster. On each host within a vSAN cluster, a VMkernel port is tagged for vSAN traffic. This VMkernel port is used to send traffic using vSAN's lightweight, purpose-built Reliable Datagram Transport (RDT) protocol. This protocol is used to help vSAN quickly build up and tear down sessions, and a myriad of other responsibilities.

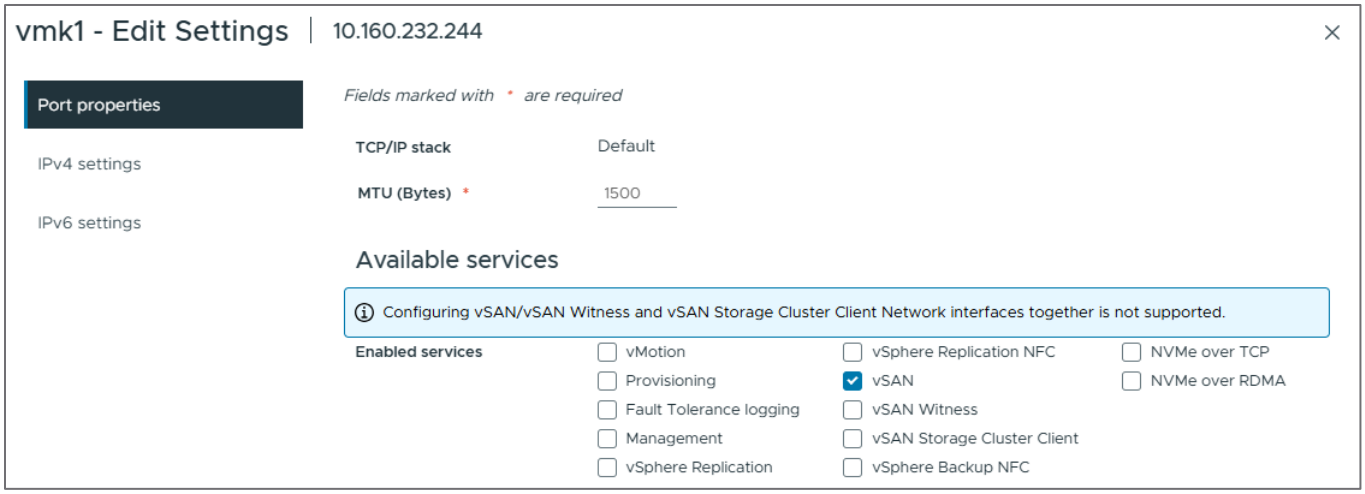

Figure. VMkernel port tagged for vSAN.

vSAN can optionally use RDMA over Converged Ethernet (RoCE v2). While TCP is highly durable, and can reliably deliver network traffic under poor conditions, RDMA can be ideal for high performance, low latency environments. RDMA does require special configurations on the switchgear and hosts and will also have other limitations.

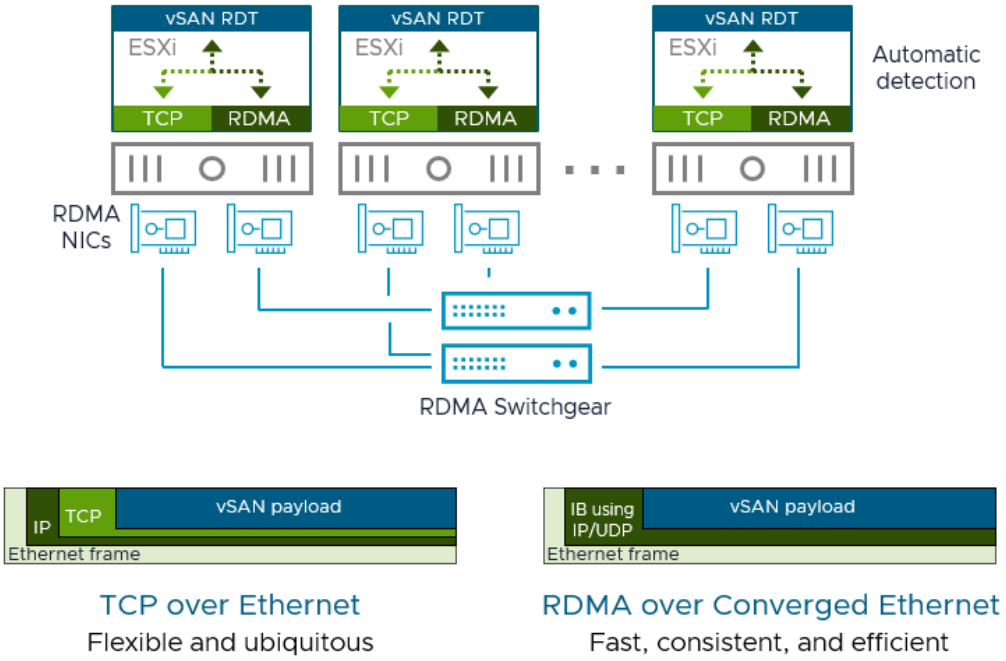

Figure. TCP over Ethernet versus RDMA over Ethernet for vSAN traffic.

Network Redundancy

Redundancy of a network ensures that any failure of a Network Interface Card (NIC) port, or a switch will not interfere with the connectivity to the rest of the cluster. For vSAN clusters, this redundant network communication is achieved in a very similar manner to that of traditional vSphere clusters.

- Hosts with two NICs consisting of one or more ports for each NIC.
- Hosts configured with a teaming policy for that vSAN VMkernel port that will allow it to fail over to another port in the event of a failure.

- A port from each NIC on the host connecting to one of two switches.

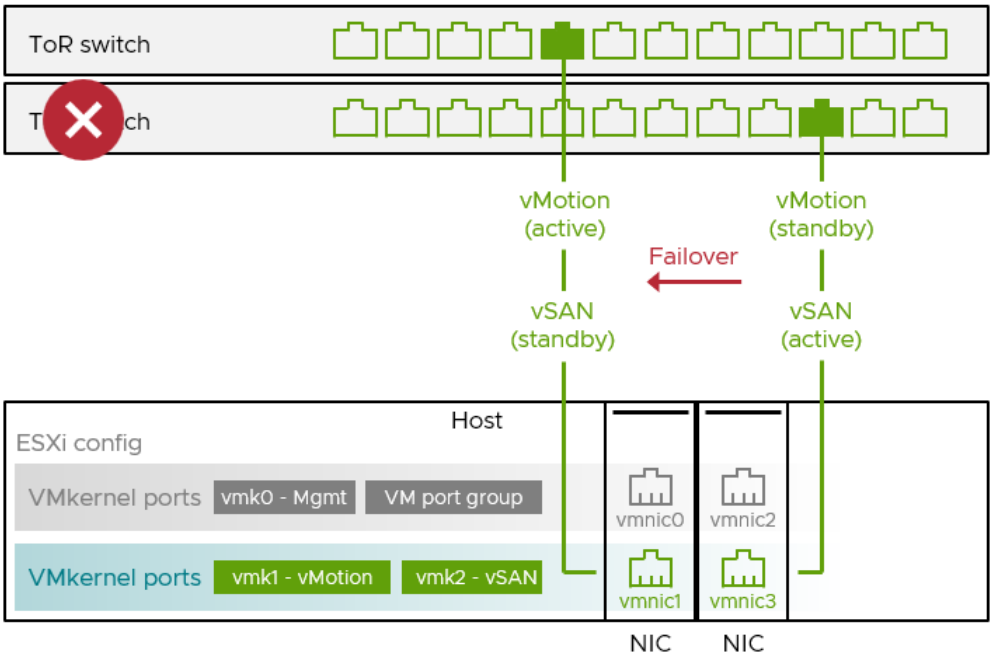

Figure. Redundancy of host NICs and switches in a vSAN environment.

The traits listed above would characterize a vSAN cluster that could continue to reliably transmit storage traffic in the event of communication failure due to discrete failures of a NIC and/or a switch failure.

Recommendation: For vSAN traffic, use a NIC teaming policy of "Active/Standby" with "Route Based on Originating Virtual Port ID." This will offer the most deterministic path for vSAN traffic, which is critical for storage performance and consistency. Note that ***VCF will by default use a teaming policy of "Active/Active" with "Route Based on Physical Load"*** *otherwise known as "Load-Based Teaming." While this approach teaming policy is adequate for VM port groups, it does not offer a deterministic path for storage traffic. You can easily override the VMkernel port settings for just the VMkernel port tagged for vSAN so that it is using "Active/Standby."* ***Overriding this vSAN VMkernel traffic to "Active/Standby" IS supported in VCF.***

Network Partitioning

In a distributed storage system like vSAN, a certain type of failure condition must be accounted for. All hosts in a vSAN cluster must be able to communicate with each other. But there may be a condition where one subset of hosts in a cluster can communicate to each other, but not to the other hosts in a vSAN cluster. This is known as a "network partition" and must be accounted for in the proper way.

An element of vSAN known as the Cluster Monitoring, Membership and Directory Services (CMMDS) is responsible for monitoring communication between the hosts in a vSAN cluster. When communication between hosts is sufficient, vSAN will indicate (in the UI, APIs, etc.) that **there is a single partition group**, reflecting a desired condition.

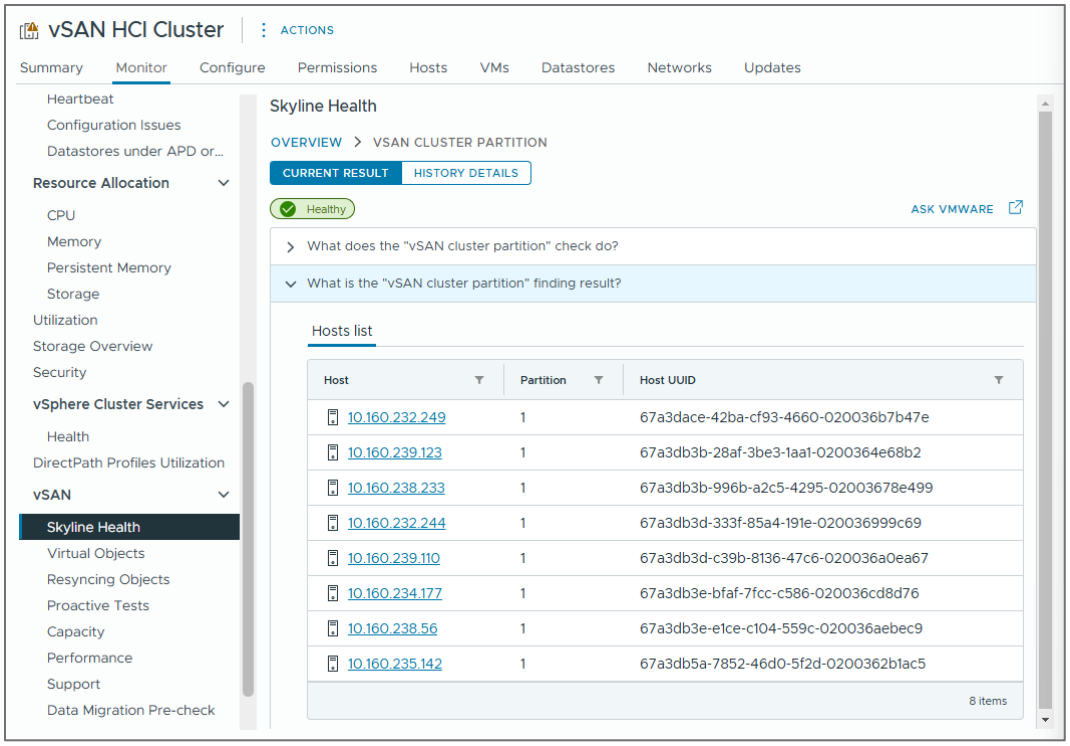

Figure. A vSAN cluster with a "healthy" partition state.

When communication between hosts is not sufficient, vSAN will indicate that there is more than one partition group. Each group will show the hosts that can communicate with each other, but not the other hosts. If this occurs, vSAN will use its quorum voting mechanism to determine:

- If any partition meets quorum for that VM or object
- If yes, ensure that the VM and data is available on that partition
- Mark any of the components in the other partition(s) that comprise that object as "absent."

This logic built into vSAN helps determine if and where the data should remain available, or unavailable. This is what prevents "split-brain" conditions where a VM and its data are updated independently in more than one location.

Since partitioning events can affect the availability of the data, the Skyline health cluster scoring, diagnostics and remediation dashboard will rank these issues with the highest level of importance. They will impact the cluster score the most, signifying an issue that should be dealt with immediately, and will stack rank this type of issue at the top of any other triggered issues. Paired with the helpful guidance to resolve the issue, this will help you identify and resolve the matter quickly.

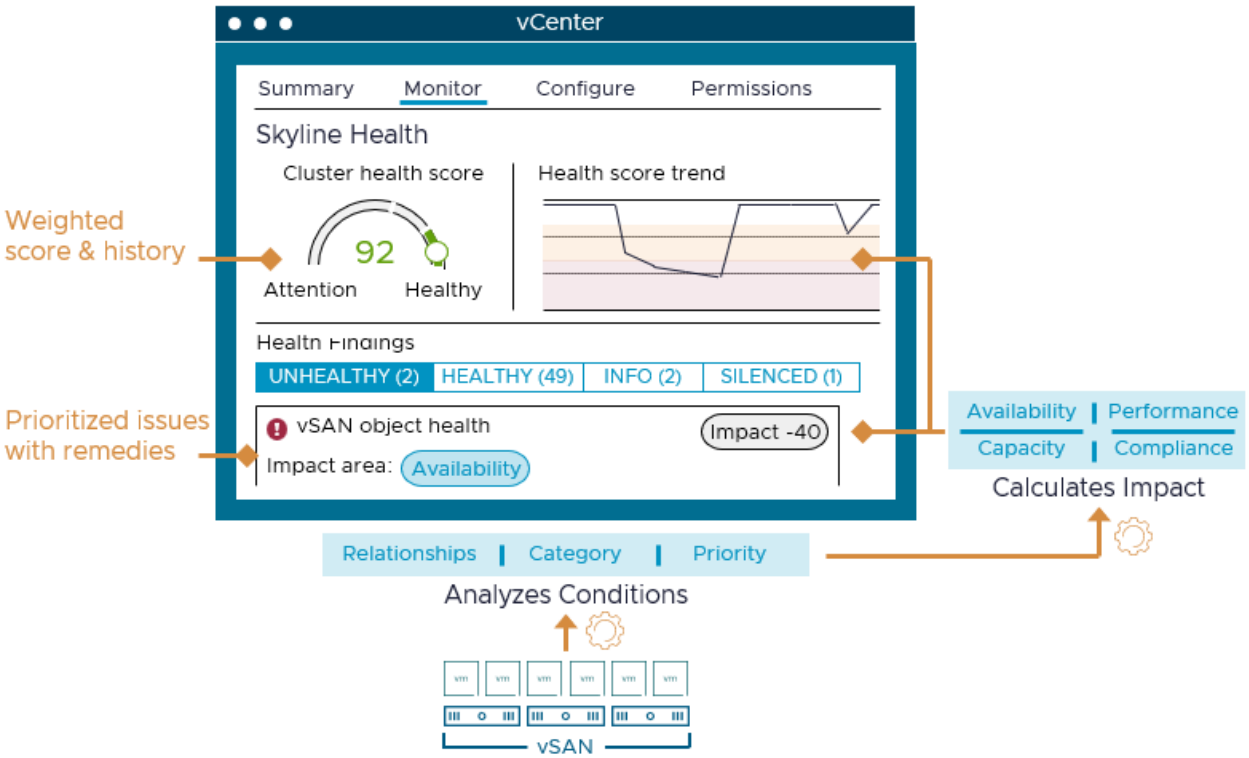

Figure. Skyline health cluster scoring, diagnostics and remediation dashboard.

Sklyline Health for vSAN has several health findings that will assist in detecting network related issues that may be the cause for partitions.

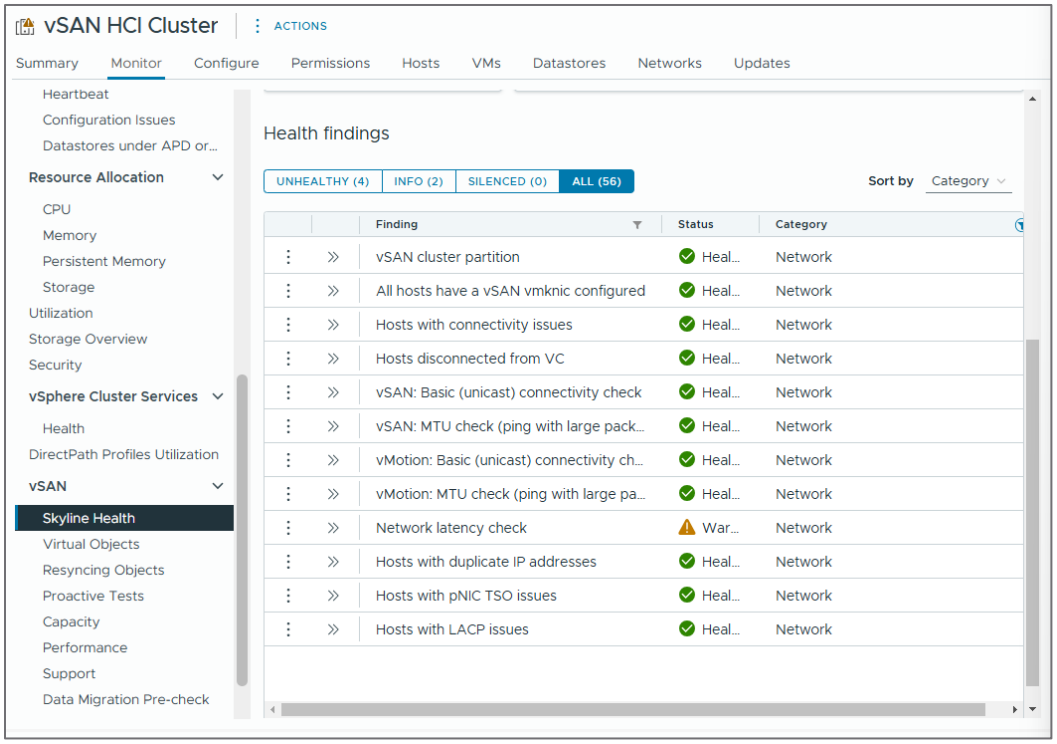

Figure. Network-related health checks in Skyline Health for vSAN.

Integration with vSphere High Availability (HA)

vSphere HA uses the same network connections as vSAN when vSAN is enabled. This helps ensure that vSphere HA and vSAN see the same network partition(s) in a cluster. Having a consistent view of network partitions across vSAN and HA enable accurate and reliable responses to host isolation and multiple network partitions.

vSphere HA is inherently reactive, meaning that it will only restart VMs if they are unavailable. vSAN also supports the use of vSphere Proactive HA in conjunction with a plug-in provided by the server OEM. This optional feature is capable of proactively removing VMs off of hosts suspected of immanent failure and may be helpful in maintaining the highest levels of uptime in your vSAN environment.

Resynchronization Activity

vSAN will ensure that data is stored per the desired resilience outcome defined in the object's storage policy. It achieves this desired outcome through "resynchronization" activity. This is I/O processing that is initiated by the system to ensure the data complies with the defined storage policy. Resynchronization may occur for several reasons, including but not limited to:

- Entering a host into maintenance mode.
- Rebalancing data across the cluster
- Storage policy changes
- Regaining prescribed level of resilience after a failure

Resync activity can be easily viewed in the UI, as shown below.

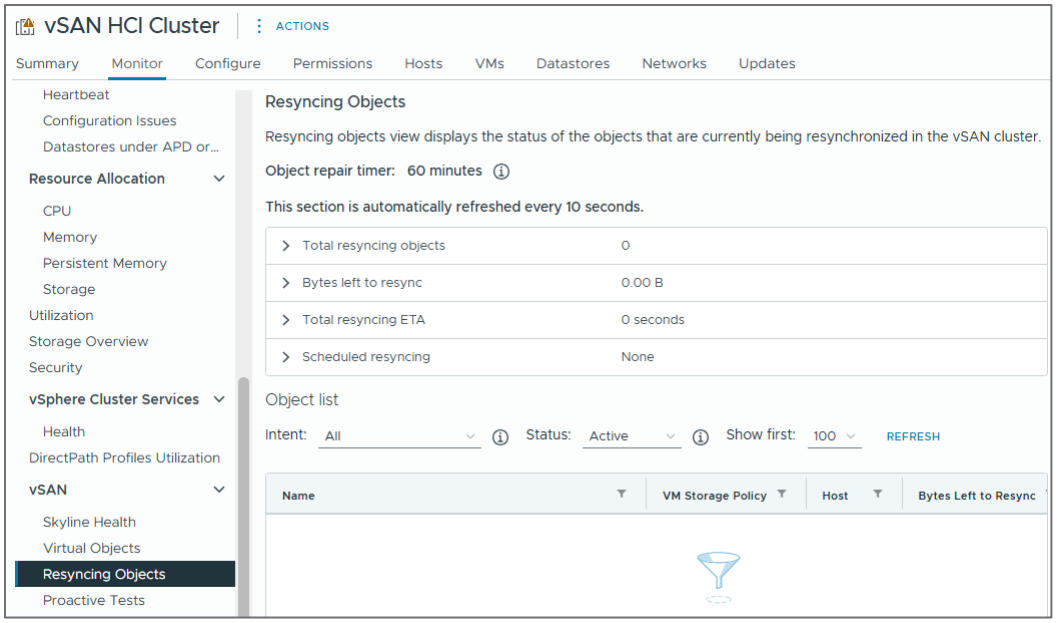

Figure. Resyncing Objects view

The guest VM is unaware of this resynchronization activity since it was not issued by the guest VM. There can be times where a large number of resynchronizations may have some impact on guest VM performance, but this was typically associated with earlier versions of vSAN, and is not a common issue anymore, especially with vSAN ESA.

Resynchronization Activity from Failures

As noted in the last bullet point above, resynchronization activity may occur after some type of failure has been recognized. For example, in an 8-host cluster, imagine an object using a storage policy of FTT=2 using RAID-6. The data with parity would be spread across 6 of the hosts. If power was lost on one of the hosts storing a portion of that object, that specific component would be marked as "absent." Since vSAN is unaware of the reason for its absence, it will wait for a period for the component to be marked as active. This helps accommodate transient outages such as a temporary network error, or a host restart. Once it exceeds this time (the default is 60 minutes), it will begin to rebuild the component somewhere else in the cluster, assuming another previously unused fault domain (host) is available.

How it rebuilds the data will depend a bit on the data placement scheme used. A repair of an object using erasure coding may involve recalculating the data and parity bits to reconstruct the remainder of the stripe. A repair of an object using mirroring may simply copy the data from the other replica. vSAN may also use the special durability components for repair processes, which can speed up repairs.

Minimizing Resynchronization Activity

Recent versions of vSAN are dramatically better at handling resynchronizations. These improvements come from a reduced boundary of failure, better processes to minimize the amount of data to resynchronize, and intelligent throttling to minimize the impact on VMs during resynchronization.

- Reduced boundary of failure after a storage device failure. (ESA)
- Much faster processing of resynchronization data. (ESA)
- Adaptive Network Traffic Shaping. (ESA)
- Adaptive Resync. (ESA and OSA)
- The use of Durability Components for faster repairs. (ESA and OSA)
- Smart Efficient repairs choose the best option to rebuild based on the conditions. (e.g. rebuild new replica versus resync existing components). (ESA and OSA)

- Partial Repairs achieves a best-effort repair to improve resilience. (ESA and OSA)
- Resumable resyncs can restart existing repair operations previously interrupted. (ESA and OSA)

Maintenance Mode

Entering a host into maintenance mode when configured as a vSAN cluster is a bit different than a traditional vSphere cluster. Since a vSAN cluster aggregates storage devices in the hosts that comprise the cluster, **placing a host into maintenance mode impacts the overall storage capacity**. When you place a host into maintenance mode, you may impact the level of failures that some objects can tolerate.

At the time that a host is entered into maintenance mode (EMM), a pre-check simulation is performed on the data that resides on the host so that vSAN can communicate to the user the type of impact the EMM will have, all without moving any data. If the pre-check results show that a host can be seamlessly placed in maintenance mode, decide on the type of data migration. Consider the storage policies that have been applied within the cluster. Some migration options might result in a reduced level of availability for some objects. The EMM option chosen will dictate the behavior and availability of data on a vSAN cluster.

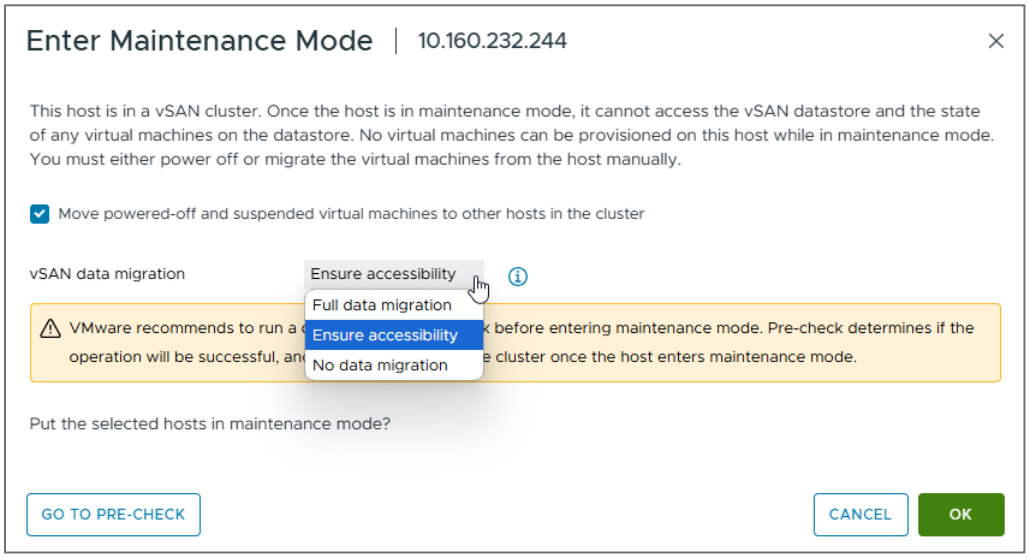

Figure. Host maintenance mode options in vSAN.

Full data migration—Evacuate all components to other hosts in the cluster.
This option maintains compliance with the FTT number but requires more time as all data is migrated from the host going into maintenance mode. It usually takes longer for a host to enter maintenance mode with Full data migration versus Ensure accessibility. Though this option assures the absolute availability of the objects within the cluster, it causes a heavy load of data transfer. This might cause additional latency if the environment is already busy. When it is recommended to use Full data migration:

- If maintenance is going to take longer than the rebuild timer value.
- If the host is going to be permanently decommissioned.
- If you want to maintain the FTT method during the maintenance.

Ensure accessibility (default option)—Instructs vSAN to migrate just enough data to ensure every object is accessible after the host goes into maintenance mode.
vSAN searches only for data with RAID-0 and move/regenerate them on a host different than the one entering in maintenance mode. All the other objects with RAID-1 and higher, should already have at least one copy residing on different host within the cluster. Once the host comes back to operational, the data components left on the host in maintenance mode update with

changes that have been applied on the components from the hosts that have been available. Keep in mind the level of availability might be reduced for objects that have components on the host in maintenance mode.

This maintenance mode is intended to be used for software upgrades or node reboots. Ensure accessibility gives the opportunity to avoid needless Full data migration, since the host will be back to operational in a short time frame. It is the most versatile of all EMM options.

The "Ensure accessibility" option will behave differently than in object creation or failure handling conditions. It puts availability of that data at the highest level of importance, with no regard to the chosen data structure or anti-affinity rules assigned by the storage policy. It will essentially defy any and all elements of the storage policy and how the object is built, which typically dictates availability, for the single purpose of maintaining availability of the object data.

No data migration—No data is migrated when this option is selected.
A host will typically enter maintenance mode quickly with this option, but there is a risk if any of the objects have a storage policy assigned with FTT=0. The "No data migration" option is best suited for full cluster shutdowns, or environments where network changes are being made.

.Recommendation. In most cases, using the "Ensure Accessibility" option will be your best choice when placing a host into maintenance mode. This will minimize the amount of data movement and will make your server maintenance operations much faster and more efficient.

For more guidance on the use of maintenance mode in a vSAN cluster, see the "vSAN Operations Guide."

Handling Failed Storage Devices

Storage devices are one of the more challenging entities to determine failure, as their operational state is not always a binary "working" versus "failed" matter. In many cases, a device may become highly unresponsive or perform poorly but still not be in a working state. As vSAN and the hardware available on the market has improved, several enhancements have been made to determine the operational state of storage devices. This, paired with the reduced impact of a storage device failure in vSAN ESA makes device handling much more sophisticated, and less impactful than in previous editions.

Degraded Device Handling (DDH)

Degraded Device Handling (DDH) is a monitoring technique within vSAN that will detect high device latency thresholds and patterns to these thresholds that suggest imminent failure of storage device in the capacity tier. This symptom-based approach that attempts to detect impending failures of storage devices. It was originally developed to help work around the shortcomings of the industry standard Self-Monitoring, Analysis, and Reporting Technology (S.M.A.R.T. or SMART). This standard never materialized as a reliable or consistent way to pole device telemetry data from SAS and SATA storage devices.

Iterative improvements have been made to DDH across multiple releases. Much of the effort went to reducing false positives by using a different algorithm to detect high latency. For example, more advanced logic will observe performance traits of a time period by breaking a configurable time period (4-hour default) into a series of configurable intervals (20-minute default). A device is considered unhealthy if performance traits exceed the configured average performance trait for that device. These are randomly selected non-contiguous intervals to better represent a pattern of degrading performance. If an issue is detected with a storage device, alerts will be triggered with those components being marked as absent. vSAN will strive to rebuild that data from other devices, and if not possible, will automatically migrate all active components from a degraded device. Other improvements include a less invasive approach when a degraded device is detected. Instead of the device being unmounted, it will be flagged and logged accordingly. As new devices have higher performance capabilities, additional logic is introduced to help account for these faster performing devices.

DDH is a symptom-based solution, and can be prone to false positives. For example, high latency can be a leading indicator of a failing device, but it can also be an indicator high sequential writes using large I/O sizes, or a period in which garbage collection activities occur. Therefore, DDH should not be viewed as a definitive technology to detect failing devices, but a complimentary mechanism that provides additional insight and protection for storage devices behaving erratically.

Recommendation. If you have experienced DDH triggering devices as failed, consider moving to vSAN ESA. vSAN ESA eliminates the use of dedicated storage controllers on hosts, as well as legacy storage devices using SAS and SATA protocols. These were much more prone to non-deterministic levels of latency.

NVMe Device Endurance Tracking (ESA only)

vSAN ESA can take advantage of a new SMART standard specification exclusively for NVMe-based storage devices. This allows vSAN to check for and alert against any NVMe storage device that is nearing its expected number of device writes. These alerts can be easily customized within vCenter Server. Unlike previous specifications, SMART data polling has little to no effect on the storage device itself, which ensures this as a reliable way to move forward on intelligent device telemetry. For more information, see the post: "Health and Performance Monitoring Enhancements in vSAN 8 U2" and "Enhanced Intelligence in vSAN 8 U3 and VMware Cloud Foundation 5.2."

Low-Level Metadata Resilience (ESA only)

vSAN ESA introduced a capability that helps address Unrecoverable Read Errors (UREs) in storage devices. UREs typically occur throughout the life span of a storage device. While vSAN employs all different types of resilience techniques of data and metadata in vSAN, we wanted to make the ESA more robust and efficient in the event of these UREs. ESA will mirror some low-level metadata on each discrete device to help ensure that the metadata created and managed by vSAN's local log structured object manager (LSOM), is resilient on the device, in the event of a discrete, 4KB URE. Even though the mirrored metadata resides on the same device, given the nature of UREs and their typical impact radius (typically at a 4KB granularity), this approach provides a statistically reasonable level of resilience to protect against UREs. Best of all, is that this new capability is completely transparent to the administrator and requires no interaction or management.

Proactive Hardware Management (ESA only)

vSAN 8 U3 (and VCF 5.2) introduced a new framework that provides a uniform, vendor neutral way to take device metadata which can serve as leading indicators of impending failure, to then predict and present its results in the form of health findings in Skyline Health so that appropriate action can be taken. This uses the vLCM framework already in place, so ReadyNode vendors can adopt it easily. It will surely prove to be a powerful way to expose key device telemetry details to address health-related issues. For more information, see the post: "Enhanced Intelligence in vSAN 8 U3 and VMware Cloud Foundation 5.2."

Availability versus Protection

Availability, resilience, durability, protection and disaster recovery are often conflated in ways that can be challenging for users to clearly understand the differences. The information below will shed some light on their similarities and differences as they relate to vSAN and its availability technologies.

Availability of Data

Storage systems typically focus on availability of data as it relates to the systems that are using that data. Availability is closely tied with "resilience" as this is the typical way that high availability is expressed and implemented in systems. vSAN has a plethora of techniques mentioned throughout this document to ensure that data remains resilient to failure, or highly available. This can include anything from vSAN's data placement schemes to stretched cluster topologies that maintain the availability requirements for customers. vSAN even has durability mechanisms in place (such as durability components) that help data to remain intact and accessible over multiple disruptions.

Protection of Data

Protection of data typically refers to the steps taken for a more definitive step, such as a recovery to a specific data and time. This is achieved through other mechanisms such as asynchronous replication such as vSphere Replication, or backups using third party tools that use special VMware APIs for Data Protection (VADP). Protection strategies are commonly built around a "3-2-1" rule. This refers to the notion of having three copies of data, with two backups using different media types or targets, and one copy living independently, outside the domain of failure. Protection strategies can even be augmented by tools that may not be a protection mechanism entirely on its own, such as snapshots.

vSAN 8 U3 and VCF 5.2 includes a new "vSAN Data Protection" feature. This uses vSAN ESAs snapshot engine to produce highly efficient snapshots local to the same datastore that can be used for quick recovery. It is fully integrated into vSAN, and can be a great way to augment existing protection strategies. Why is it called "vSAN Data Protection" when you may have always been told that snapshots are not backups? This is an excellent question and is answered in the "vSAN Data Protection" section of the vSAN FAQs.

Disaster Recovery

Disaster Recovery is an umbrella term that implies multiple elements in place to maintain full business continuity in the event of a major event or disaster. Disaster Recovery usually employs a combination of asynchronous replication to a remote location (vSphere Replication) and an orchestration solution such as VMware Live Recover (VLR), or more specifically, VMware Live Site Recovery (VLSR) previously known as Site Recovery Manager, or SRM.

Multiple Concepts Working Together

It is possible to use elements of each to provide a comprehensive approach for a highly resilient, protected environment. For example, one could power a VCF environment using vSAN stretched clusters to provide an active/active data center that will remain up in the event of an entire site outage. One can then protect these workloads using asynchronous replication to a target living at a third location using VLSR should there be a need to fail over to the DR site in the event of a dual site failure. One could pair this with a thorough 3-2-1 data protection strategy that perhaps includes the use of a third-party backup software, augmented by the unique power of vSAN Data Protection and its highly efficient and scalable snapshots.

Summary

vSAN can provide all the resilience capabilities you'd expect to maintain the highest levels of availability in your data center. Understanding the basic concepts around how vSAN stores data will help you design, operate, and optimize your environment to all new levels.

Additional Resources

The following are a collection of useful links that relate to bandwidth sizing for vSAN stretched clusters.

vSAN Interactive Infographic. This tool allows you to dynamically choose various configurations and failure scenarios to better understand how vSAN maintains supreme levels of availability.

vSAN Failure Handling Guide. Provides information on how vSAN handles specific failure scenarios.

Performance Recommendations for vSAN ESA. This is a collection of recommendations to help achieve the highest levels of performance in a vSAN ESA cluster. Many of these same recommendations apply to vSAN storage clusters.

vSAN Proof of Concept (PoC) Performance Testing. This is a collection of recommendations that will guide users to test the performance of a vSAN cluster. While it is currently written for the OSA, many of the testing methods used are also applicable to the ESA.

Design and Sizing for vSAN ESA clusters. This post offers some nice guidance on using the vSAN Sizer for the ESA that summarizes some key points that can be found in the VMware vSAN Design Guide.

vSAN Network Design Guide. This network design guide applies to environments running vSAN 8 and later.

vSAN technical blogs. Stay up to date on the most recently published technical information about vSAN. These posts are created by the vSAN Technical Marketing team.

VMware Resource Center. The location for design guides, operations guides and other technical white papers on vSAN. These assets are created by the vSAN Technical Marketing and Product Enablement teams.

Official vSAN documentation. The location for all "how to" documentation on vSAN.

About the Author

Pete Koehler is a Product Marketing Engineer in the VCF division at Broadcom. With a primary focus on vSAN, Pete covers topics such as design and sizing, operations, performance, troubleshooting, and integration with other products and platforms.