vSAN Frequently Asked Questions

Common questions and answers for vSAN

January 28, 2025

vmware by Broadcom

Table of Contents

General Information

What are the hardware requirements when running vSAN?

The hardware requirements will depend on the capacity and performance requirements for your environment. The newer, more powerful and efficient vSAN Express Storage Architecture (ESA) has different hardware requirements than the original storage architecture (OSA). For the latest minimum hardware requirements for vSAN ReadyNodes, see the "vSAN ESA ReadyNode Hardware Guidance" document.

For a standard, single site topology, vSAN typically requires a minimum of three hosts, and can support as many as 64 hosts in a single cluster. Although, it does have a "2-Node" deployment option that is more appropriate for remote/edge environments. This will consist of two hosts storing the data resiliently, and a third host at a central location that helps determine availability of the data.

What are the typical hardware deployment options available?

vSAN runs on commodity servers running x86 processors. Hardware componentry that has been certified to work with vSAN will be available in the form of vSAN ReadyNodes. These are preconfigured hosts from your favorite server OEM vendor and purchased with a single SKU. You can purchase this same hardware using discrete hardware components from these same OEM vendors. These are known as ReadyNode Emulated configurations.

Turn-key appliances such as Dell EMC VxRail, Hitachi UCP HC, and Lenovo ThinkAgile VX Series are also available.

How much memory is required for a vSAN host?

The requirements will vary depending on demands whether ESA or OSA is being used, as well as the demand and performance expectations of the workloads. vSAN OSA requires 32GB of RAM per host, while ESA requires 128GB. A given vSAN ReadyNode profile may have higher minimum requirements, however. For the latest minimum hardware requirements for vSAN ReadyNodes, see the "vSAN ESA ReadyNode Hardware Guidance" document. When a vSAN cluster is running, one can view the amount of memory used by vSAN by highlighting the cluster, clicking **Monitor > vSAN > Support > Performance for Support > Memory > vSAN Memory**.

What are the processor requirements for a vSAN host?

The requirements will vary depending on demands whether ESA or OSA is being used, as well as the demand and performance expectations of the workloads. For the latest minimum hardware requirements for vSAN ReadyNodes, see the "vSAN ESA ReadyNode Hardware Guidance" document.

In the vSAN Original Storage Architecture (OSA), what is a disk group?

The vSAN OSA uses a unit of storage resources known as a disk group. Each disk group consists of a flash device serving as a cache tier and one or more capacity devices providing capacity. Each host has a minimum of one and up to a maximum of five disk groups. Each disk group consists of precisely one cache device and a minimum of one up to a maximum of seven capacity devices. When using the vSAN OSA, most vSAN configurations have between 1 and 3 disk groups.

The vSAN Express Storage Architecture (ESA) does not use a concept of disk groups, but rather, a "storage pool." This simply represents all devices in a host that have been claimed for the purpose of being used by the vSAN ESA. In the ESA, all storage devices contribute to both capacity and performance, so unlike the OSA, there is no dedicated caching or capacity tier, nor is there a funneling of I/O through a disk group. The design of the ESA allows for it to no longer need the concept of a disk group, and as the result offers far better performance, and improved durability of data. See the post: "The Impact of a Storage Device Failure in vSAN ESA versus OSA" for more information.

When using the vSAN OSA, why are flash devices needed for the cache tier?

The vSAN OSA is built with two distinct tiers: A caching/buffering tier, and a capacity tier. The flash device in the cache tier of each disk group is used as a write buffer in all-flash vSAN configurations. These cache devices are typically higher-endurance, lower-capacity devices. Data is de-staged from the cache tier to the capacity tier. Capacity tier devices are more commonly lower-endurance, higher-capacity flash devices. A large portion of reads in an all-flash vSAN cluster are served directly from

the capacity tier. An all-flash configuration provides a good balance of high performance, low latency, endurance, and cost-effectiveness. In versions prior to vSAN 8, the write buffer device was limited to 600GB of logical capacity. This has been increased to 1.6TB in vSAN 8. For more information, see the post: Increased Write Buffer Capacity for the vSAN 8 Original Storage Architecture. The vSAN OSA does not support a write buffer device using partitioning through NVMe namespaces, or any other partitioning method.

The vSAN Express Storage Architecture (ESA) is designed differently, and does not have these distinct tiers.

How should I size my cache tier when using vSAN OSA?

For the OSA, given the typical device sizes on the market today, it is simply best, and easiest to plan for a 1.6TB cache device per disk group in the OSA. This will provide an optimal amount of write buffer and caching capacity to improve the performance of an OSA-based cluster. For the ESA, this is not an issue since it does not use a dedicated caching device.

Does vSAN use vSphere vVols?

No. vVols are designed for use with external storage arrays. vSAN uses local drives to create a shared datastore. vSphere Virtual Volumes and vSAN can be used in the same cluster and both provide the benefits of storage policy-based management.

Can I use existing storage arrays (block or file) in the same cluster as a vSAN cluster?

Yes, vSphere can access and use traditional VMFS and NFS datastores alongside vSAN and vSphere Virtual Volumes—all in the same cluster. If your SAN or NAS solution is compatible with vSphere Virtual Volumes, management is easier and more precise as storage policies can be used to manage all of your storage on a per-VM basis.

In most cases, vSphere Storage vMotion can be used to migrate VMs between these various datastore types. This feature makes it easy to migrate existing workloads when there is a need to perform maintenance or retire and older storage solution.

How can I size a vSAN cluster so that it meets my capacity and performance requirements?

The vSAN ReadyNode Sizer can help you determine suitable vSAN ReadyNode server configurations, and will work for sizing a cluster using the vSAN Original Storage Architecture (OSA) and the vSAN Express Storage Architecture (ESA). Tools such as Live Optics and RVTools can help you asses a current environment to begin the sizing exercise.

What can I change in a vSAN ReadyNode?

vSAN ReadyNodes are extremely flexible in their ability to be tailored to meet the requirements of your environment. For more information, see "What you can (and cannot) change in a vSAN ReadyNode."

Can I add a host that does not have local storage devices to a vSAN cluster?

A host with no local storage can be added to a vSAN cluster, but is not recommended.

Recommendation: Use uniformly configured hosts for vSAN deployments. While compute only hosts can exist in a vSAN environment, and consume storage from other hosts in the cluster, VMware supports vSAN clusters with asymmetrical host configurations, VMware does not recommend having unbalanced cluster configurations, and thus minimize significant levels of asymmetry. This will help prevent potential availability and capacity issues during failure conditions, and minimize deviation in performance. See the post: "Asymmetrical vSAN Clusters - What is Allowed, and What is Smart" for more information.

Can a vCenter Server Appliance (VCSA) be installed on a single host on a new cluster?

Yes. VCSA deployment wizard includes the ability to claim disks and turn on vSAN on a single host. This enables administrators to deploy vCenter Server to a new environment where vSAN will be the only datastore.

Can I mix different hardware componentry in a vSAN cluster?

Clusters have hosts that use modest levels of dissimilar hardware. It is highly advised to keep the hardware within a cluster as similar in hardware componentry and symmetrical in hardware resources as possible. This principle not only applies to vSAN clusters, but traditional vSphere clusters as well. For more information, see the post: "Asymmetrical vSAN Clusters – What is Allowed, and What is Smart."

Are there any vSphere features that are not supported with vSAN?

Nearly all vSphere features such as VMware vSphere vMotion™, VMware vSphere Distributed Resource Scheduler™, VMware vSphere High Availability, VMware vSphere Network I/O Control, and VMware vSphere Replication™ are compatible and supported with vSAN. VMware vSphere Fault Tolerance is supported for VMs running on vSAN except for stretched clusters.

The following vSphere features are not supported with vSAN:

- VMware vSphere Distributed Power Management
- VMware vSphere Storage DRS™ VMware vSphere Storage I/O Control

Can I share a single vSAN datastore across multiple vSAN and vSphere clusters?

Yes. By default, a vSAN datastore is directly accessible only by the hosts and VMs in the vSAN cluster. When using vSAN's disaggregated offering, the datastore from one vSAN HCI cluster can be mounted to another vSAN HCI cluster. This is known as "vSAN HCI with Datastore Sharing." vSAN storage clusters is a vSAN deployment option that provides a centralized shared storage solution for your vSphere clusters and augment storage for your vSAN HCI clusters.

How does vSAN store objects such as VM configuration files and virtual disks?

vSAN stores data in a way that is very analogous to an object store. Items such as a VM's configuration (VMX) and virtual disks (VMDKs) are stored as objects. An object consists of one or more components. The size and number of components depend on several factors such as the size of the object and the storage policy assigned. The following figure shows common virtual machine objects.

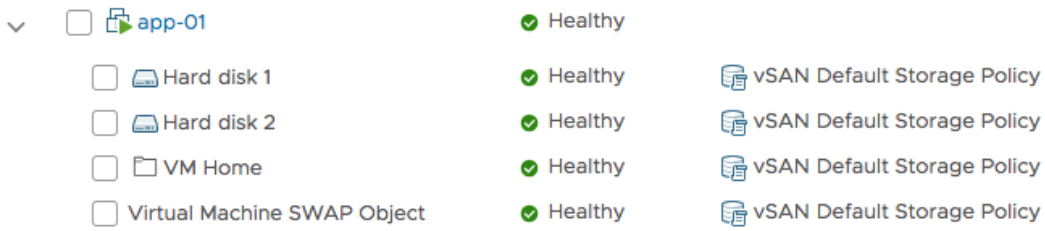

Each object commonly consists of multiple components. Components are just an implementation detail, and not a manageable entity in vSAN. For more detail, see the blog post: [vSAN Objects and Components Revisited](). The image below provides details on the number of components and the hosts where they are located for a vSAN OSA cluster. The vSAN Default Storage Policy is assigned which contains the rules Failures to Tolerate (FTT) = 1 and RAID-1 (mirroring). As a result, two components are created - two copies of the virtual disk - and there is a witness component that is used to achieve quorum if one of the other components is offline or there is a split-brain scenario. The vSAN ESA stores objects slightly different, but the principles are the same.

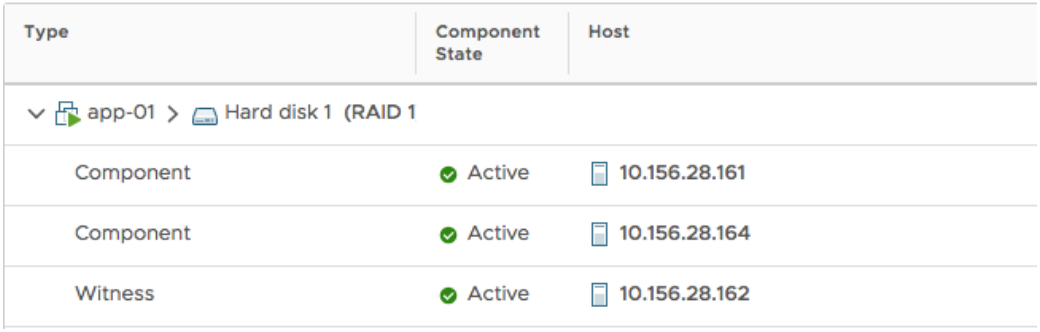

Note: The witness component should not be confused with the witness host virtual appliance discussed earlier in this document as they are two different items.

When a VM is migrated to another host, are the VM's objects migrated with the VM?

The concept implied in the question is often referred to as "data locality". vSAN does not require host-based data locality to achieve excellent performance. vSAN does not tax the vSAN backend network with moving multiple gigabytes of data every time a VM is migrated to another host. There is no need to do this considering modern 10Gb and 25/100Gb networking paired with high performing storage devices.

Where can I find guidance on vSphere boot devices for hosts in a vSAN cluster?

The vSAN Design guide will provide more information on recommended boot devices. The trend of using SD cards and USB devices as a boot device for the hypervisor is a trend that is falling out of favor, and will have limited support in future editions of vSphere. Using a device (such as an SSD, M.2 or BOSS card, etc.) that offers true persistent storage will allow crash dumps and logging to occur on the device, and is a much more useful and reliable configuration.

Does vSAN support 3D XPoint or Intel Optane storage devices?

The vSAN OSA supports the use of Intel Optane storage in the caching tier. Note that Intel announced the end of development of their Optane storage devices.

How is cost modeling different with vSAN versus traditional three-tier architectures?

One of the most compelling traits of vSAN is the incremental and linear scaling it possesses, and the smooth and predictable cost model. These two traits have historically been a challenge for data centers, because incremental growth can mean large capital expenditures when an environment reaches various resource thresholds. Traditional monolithic arrays, and the storage fabric that they run on, become much more costly as the demand of that shared storage unit increases. vSAN's architecture is different. A cluster serving your workloads can easily be scaled out one host at a time, or even scaled up using more or higher density storage devices that can be purchased as a direct expense, as opposed to a large capital expenditure. All of this can be done easily in a non-disruptive manner that is prescriptive, and predictable.

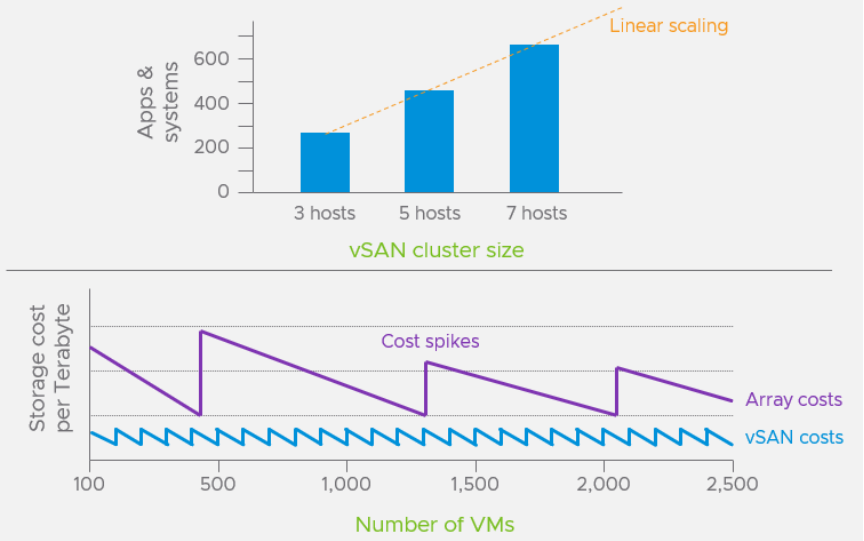

Express Storage Architecture (ESA)

What is the vSAN Express Storage Architecture?

The vSAN Express Storage Architecture (ESA) is an optional, alternative architecture in vSAN that is designed to process and store data with all new levels of efficiency, scalability, and performance. This optional architecture is optimized to exploit the full capabilities of the very latest in hardware. It was introduced in vSAN 8, and enhanced in Update 1, 2 and 3. It can be selected at the time of creating a cluster. The ESA in vSAN is an alternative to the Original Storage Architecture (OSA) found in all previous editions of vSAN, as well as an optional architecture in the very latest version.

Does vSAN 8 U3 include the Original Storage Architecture found in past editions of vSAN?

Yes! vSAN includes the new Express Storage Architecture (ESA) and the Original Storage Architecture (OSA). The OSA is the architecture that will be used for all in-place cluster upgrades, and new cluster installations using hardware that is not qualified for use with the ESA. The ESA can be used for new cluster installations using qualified hardware. The ESA has a much lower TCO, especially given recent announcements of a new entry-level ReadyNode classification, and the support of Read-Intensive storage devices.

Won't the use of NVMe-based TLC storage devices make ESA more expensive?

No. The architecture and abilities of the vSAN ESA will often make all NVMe-based configurations more affordable than clusters running the vSAN OSA using SAS devices. The comparatively better TCO comes in three forms: 1.) Elimination of purchasing caching/buffering devices and storage controllers. 2.) Running space efficient erasure codes without any performance compromise and do so on clusters as small as three hosts. 3.) Improvements in vSAN's compression architecture that will likely bring more capacity savings to existing environments. Items #1 & #2 are often enough to make running the ESA more cost efficient per Terabyte than running a similar cluster configured with SAS devices and running the vSAN OSA. With the recent introduction of a new entry-level ReadyNode classification, and support of Read-Intensive storage devices, the TCO will favor the ESA even more.

My vSAN cluster already runs all NVMe devices. Why should I consider running ESA?

Past editions of vSAN have supported the use of all NVMe-based storage devices. NVMe-based storage devices are certainly much higher performing than their SAS and SATA counterparts. While the original storage architecture (OSA) in vSAN could provide a fast storage platform, the ESA was built with these next generation devices in mind. When using vSAN ReadyNodes approved for the ESA, using ESA will be able to exploit the full potential of these storage devices, offering near device-level performance and consistency while improving operational simplicity, and driving down TCO.

There is a lot of talk about "efficiency" with ESA. What does this mean, and why is it so important?

The vSAN ESA is designed for the capabilities of storage devices of today and tomorrow. All new technologies are emerging that may dramatically increase the capacity capabilities of a storage device. Even if a storage system has an extremely fast data path, a system must use a minimal amount of CPU resources per I/O, otherwise it could run out of CPU resources during high loads. Increased storage densities and new storage techniques can also place a burden in scaling metadata - the data about the data. The vSAN ESA was carefully designed to address these challenges, and as a result, "efficiency" may be one of the most compelling aspects of the vSAN ESA, and what it brings for our customers. An efficient system allows you to do more with what you already have.

Is the Original Storage Architecture (OSA) going away?

No, not in the foreseeable future. We recognize many of our customers have invested heavily in a wide variety of hardware, and customers can continue to use the vSAN OSA with confidence for those configurations. The OSA will generally receive new management improvements that are applicable to both architectures. Continuing to upgrade your clusters to the latest version of vSAN using the OSA is a great way to use your existing hardware most effectively. New clusters can be configured with the ESA in mind, which will drive down TCO while benefiting from all of the capabilities that the ESA provides. The **ESA can deliver new capabilities and levels of performance and efficiency that are simply not possible in with the OSA**.

Will vSAN ESA support the use of spinning media in a hybrid configuration?

No. In vSAN 8, the vSAN ESA supports high-performing, NVMe-based TLC flash storage devices in vSAN ReadyNodes approved for the ESA. Spinning media is simply not able to support the capabilities that make the ESA so special. Spinning media devices will continue to be supported when using the OSA in a vSAN cluster.

What do I need to run the vSAN Express Storage Architecture in my environment?

The Express Storage Architecture (ESA) can be used with vSAN ReadyNodes or emulated ReadyNodes approved for use with the ESA. Our partners participating in the ReadyNode program have ReadyNodes approved for use with the ESA.

For vSAN 8 and the ESA, the servers must be vSAN ReadyNodes approved for use with the ESA. The ReadyNodes can also be "**emulated**" which means they are not purchased as an official ReadyNode using a single SKU, but are built using the same OEM server hardware and meets the minimum ReadyNode requirements. For more information, see the post "Support for ReadyNode Emulated Configurations in vSAN ESA."

The ESA does not support the old "Build Your Own" (BYO) configuration in the same way the OSA provided, but our ReadyNode program for vSAN ESA offers tremendous flexibility in configuration of hosts that fall within a ReadyNode profile, which essentially provides a guided "BYO" option.

Can the ESA support different storage device sizes in the same host?

Yes. While customers should strive for a relatively symmetrical cluster, we understand that market conditions over time can make purchasing storage devices of the same capacity challenging. The recommendations for cluster symmetry are very similar to using the original storage architecture. For more information, see the post: Asymmetrical vSAN Clusters - What is Allowed, and What is Smart. Initially, vSAN 8 can only be run on vSAN ReadyNodes approved for use with the ESA, which will naturally invite consistency and symmetry across a cluster.

While vSAN ESA supports both Read-Intensive (RI) and Mixed-Use (MU) devices, mixing devices with different endurance levels within a host, or across hosts within a cluster is not supported. For more information, see the post: "Expanded Hardware Compatibility for vSAN Express Storage Architecture."

Is there a migration path to vSAN 8 and/or the ESA in vSAN 8?

For customers who have existing clusters running previous versions of vSAN, and wish to upgrade to vSAN 8, they can perform an in-place upgrade of the cluster as usual. **In this scenario, the cluster will then be upgraded to vSAN 8 and use the original storage architecture (OSA).** For customers with ReadyNodes approved for use with the ESA, a cluster can be created, and during the installation process, The "vSAN ESA" can be selected and it will proceed to install and configure the cluster as such. When the configuration is complete, customers can migrate the VMs using vMotion and Storage vMotion.

For more information on transiting to the ESA, see the post: Migrating to the Express Storage Architecture in vSAN 8.

Will the vSAN ESA look and operate in the same way as past vSAN versions?

Yes. Many of the aspects of the Express Storage Architecture are "under the hood" architectural changes. vSAN will continue to operate in the same way as past editions. **It truly remains the software you already know.** In some ways, operational aspects become much easier. The vSAN ESA does not use disk groups, so attending to a maintenance issue with a storage device becomes easier with a smaller area of impact. Storage policy decisions become easier because RAID-5/6 can deliver the performance of RAID-1. In fact, when it comes to storage policies, the ESA makes this process even easier with the new Auto-Policy Management capability found exclusively in the ESA.

For more information, see the blog post: Availability and Serviceability Improvements in the vSAN Express Storage Architecture.

How do I size a new cluster running the ESA in vSAN 8?

The vSAN Sizer (at https://vcf.broadcom.com/tools/vsansizer/login.html) is a sophisticated sizing tool that will step you through the process of sizing your **performance and capacity** needs accurately. The vSAN Sizer will be updated to accommodate running vSAN using the OSA, or ESA. The desired architecture can be selected during the sizing process.

Note that **when comparing potential space and resource savings for the ESA versus the OSA, use the ReadyNode Sizer instead of the ReadyNode Quick Sizer.** The Quick Sizer will always show a cluster size based on the user input, not the number of hosts that the ReadyNode Sizer will calculate. The ESA results in a lower TCO due to the following:

- One can always use space efficient RAID-5/6 erasure coding without any compromise in performance.
- Reduced costs per host with the removal of cache devices, and no controllers.
- Potential capacity improvements through improved compression.
- Potential to use fewer hosts to serve the same number of workloads as a similar OSA cluster.

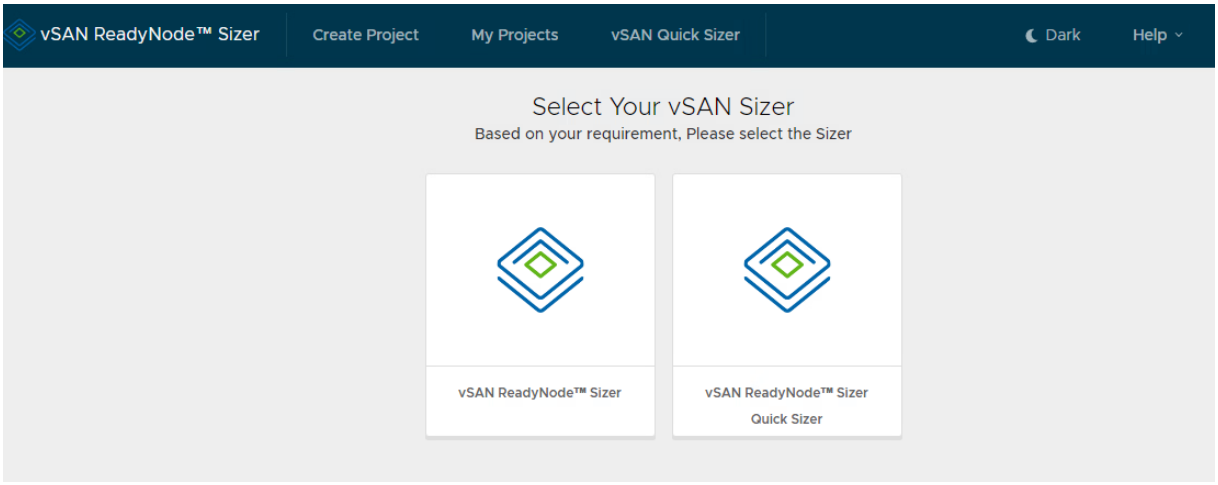

The results of the vSAN Sizer can provide a good representation of what will be needed in an environment. See the post: "Improved Capacity Reporting in VMware Cloud Foundation 5.1 and vSAN 8 U2" to learn how entering in correct inputs into the sizer can yeild an accurate, real world result.

Sizing vSAN Max clusters will share many similarities to a vSAN HCI cluster, but differs since it generally doesn't need to account for resources used to run VM instances on the cluster, since the majority of those will be running on vSphere clusters. More information on sizing vSAN Max will be forthcoming.

I see the minimum hardware requirements, and it indicates that faster networking is required on the ESA for the majority of ReadyNode profiles. Does the ESA use more network resources to process data?

No. While many ReadyNode profiles will define 25Gb or 100Gb as the minimum requirement, this requirement reflects the ESA's ability to deliver near device-level performance of the high-performing NVMe-based storage devices approved for use. Ensuring sufficient network resources allows for vSAN to exploit the full performance capabilities of the devices under maximum load. **If you are migrating production workloads from a vSAN OSA cluster to a vSAN ESA cluster, on average you will see fewer CPU and network resources used for those same workloads**. This is because the vSAN ESA uses fewer CPU cycles and fewer network resources to process and store I/O when compared to the vSAN OSA. **To account for environments that may be powering modest workloads, we've introduced a new ESA-AF-0 ReadyNode profile that requires just 10Gb networking.** Please use the vSAN ReadyNode sizer to determine what is the best option for you. While vSAN ESA is very efficient in transmitting data across the network, it is an extremely fast storage stack, and under heavy load may hit the physical limits of your cluster. Therefore, we recommend 25Gb or higher networking when using vSAN ESA.

I'm looking for my favorite ReadyNode on the compatibility list for ESA, but do not see it. What am I doing wrong?

This may occur for a few reasons.

- **Completion of certification.** Server manufacturers are responsible for the certification of systems, and the certification status will be dependent on the desires of the server manufacturer.
- **Generation of ReadyNode.** Among other hardware requirements, the Express Storage Architecture requires ReadyNodes that use chipsets no older than Intel's Ice Lake family of processors. ReadyNodes using processors older than Intel Ice Lake (Debuted in April, 2021) use an older generation of PCIe, and also do not have sufficient number of PCIe lanes to support the quantity of NVMe storage devices supported by the ESA. These older servers will also be incapable of providing line-rate speeds of modern, 100Gb NICs.

For the latest list of ReadyNodes compatible with ESA, visit the Broadcom Compatibility Guide (BCG) for vSAN.

Does the Express Storage Architecture in vSAN 8 use a caching device?

The ESA in vSAN 8 up to and including vSAN 8 U3 does not use a dedicated caching device. The Express Storage Architecture is a flexible architecture using a single tier, so all performance and capacity needs are being performed by the same storage devices.

What is a storage pool in the vSAN ESA?

The vSAN ESA removes the concept of disk groups and discrete caching and capacity tiers, and replaces it with a "storage pool." This storage pool is comprised of all storage devices selected on the host to provide storage resources to vSAN. Since the ESA in vSAN 8 does not have a dedicated caching tier to deliver performance, all selected storage devices in a storage pool on a host will contribute to capacity and performance. This improves the TCO as all devices contribute to capacity, and it also dramatically simplifies the provisioning process, and reduces the impact of a storage device failure. There can only be one storage pool per host. And these contribute to a single vSAN datastore per cluster.

Was vSAN rewritten?

No. Many aspects of vSAN Original Storage Architecture (OSA) are used throughout the Express Storage Architecture (ESA). vSAN has solved many of the great challenges of distributed storage systems, and we wanted to build off of these capabilities already built into vSAN. The ESA simply helps customers capitalize on the capabilities of this latest generation (and beyond) of high-performance hardware. Its architecture allows for the most efficient use of resources, so that vSAN is best positioned to power all application types.

How does the vSAN ESA deliver the performance of RAID-1 mirroring while using RAID-5/6 erasure coding?

vSAN introduces a new, patented log-structured file system (LFS) and an optimized log-structured object manager to deliver significant efficiencies throughout the stack. The vSAN LFS allows us to ingest writes quickly, efficiently and in a durable manner, while preparing the data and metadata for an efficient, full stripe write. The new LFS in the ESA takes advantage of our approach to writing data resiliently by first quickly writing using a redundant mirror, and packages it in a way that allows vSAN to write the data to a stripe with parity, all while maintaining the metadata in a very efficient and fast manner.

For more information, see the post: RAID-5/6 with the Performance of RAID-1 using the vSAN Express Storage Architecture. Stretched clusters will also see a performance and efficiency improvement as well. For more information, see the post: Using the vSAN ESA in a Stretched Cluster Topology.

Does the ESA in vSAN 8 U1 through U3 offer better performance than the ESA found in vSAN 8?

Yes. The ESA in vSAN 8 U1 and U2 introduces several enhancements that drive better performance than the initial release of vSAN ESA. The new adaptive write path allows vSAN to determine the type and size of incoming I/Os, and if it meets certain criteria (such as the size of I/O), will bypass committing the data to our durable log, and simply perform a full-stripe write while writing the data to the durable log. We've also increased the parallelism of certain processes in the data path that will drive better performance with single VMDKs. These improvements can provide a significant performance improvement in the most demanding workloads.

See the post "vSAN 8 U1 Express Storage Architecture - Faster than Ever" for more information on how the performance in the ESA compares to the OSA in vSAN 8 U1.

Which storage policy data placement scheme (RAID-1, RAID-5, or RAID-6) should I use for VMs powered by a cluster running the ESA?

Since the vSAN ESA eliminates the trade-off of performance versus deterministic space efficiency, the data placement scheme recommended is largely dependent on the size and type of vSAN cluster, and the level of resilience desired by the customer. While customers can prescribe storage policies per VM to whatever they desire, the guidance below will help you determine which data placement scheme makes the most sense for your organization, for vSAN ESA powered clusters.

- **For clusters with 7 or more hosts**. Select FTT=2 using RAID-6. This spreads the object data (and parity) across 6 hosts. It offers a very high level of resilience while being able to store the data in a space-efficient manner. With 7

hosts, one spare fault domain (host) will be available to regain the prescribed level of resilience in a host failure or maintenance condition.

- **For clusters with 6 hosts.** Select FTT=1 using RAID-5. This spreads the object data (and parity) across 5 hosts. It offers the resilience of data while providing supreme levels of space efficiency. With 6 hosts, one spare fault domain (host) will be available to regain the prescribed level of resilience in a host failure or maintenance condition.
- **For clusters with 3-5 hosts.** Select FTT=1 using RAID-5. This spreads the object data (and parity) across 3 hosts. It offers the resilience of data while providing great levels of space efficiency. With 4-5 hosts, one spare fault domain (host) will be available to regain the prescribed level of resilience in a host failure or maintenance condition.
- **For 2-Node clusters**. Select FTT=1 using RAID-1. This mirrors the data across two hosts and uses the virtual witness host appliance to help determine quorum.
- **Stretched clusters.** Select FTT=1 using RAID-1. This mirrors the data across two sites and uses the virtual witness host appliance to help determine quorum.

Beginning with vSAN 8 U2, the ESA offers a new Auto-Policy Management feature that helps determine and configure the ideal default storage policy for a cluster based on the 1.) type of cluster, the number of hosts in a cluster, and whether or not the "Host Rebuild Reserve" is enabled or not. While it has slightly different recommendations than the guidance provided above, it can offer a simple way for customers to ensure that their default storage policy for a given cluster is set to the ideal configuration. For more information, see the post: Auto-Policy Management Capabilities with the ESA in vSAN 8 U1.

Will the Auto-Policy Management feature in ESA change all of my storage policies?

No. The new Auto-Policy Management feature is an optional feature that is disabled by default, and designed to be non-invasive to your existing environment. Once enabled, it creates a cluster-specific default storage policy so that it does not impart sub optimal policy settings for other clusters. VM's already using the default policy will continue to do so. It will only use the default policy for newly created VMs, and trigger a health alarm for VMs, and describe to the user the ideal default storage policy to be used. It is up to the user to make the change.

vSAN 8 U2 improves on this functionality by providing a single one-click remediation step in the triggered health finding to help customers change the cluster specific default storage policy to the optimal settings. For more information, see the post: Auto-Policy Remediation Enhancements for the ESA in vSAN 8 U2

I read that the vSAN ESA can support RAID-5 on 3 hosts? How does this work? I thought vSAN required 4 hosts at a minimum for RAID-5?

The vSAN OSA uses a 3+1 stripe with parity data placement scheme, and as a result, requires 4 hosts minimum to run in a non-error state. The vSAN ESA uses two new RAID-5 erasure codes, that are automatically determined by vSAN based on the size of the cluster, and will adapt to the size of the cluster automatically.

- **6 or more hosts.** When RAID-5 is applied to objects in a cluster of this size, the vSAN ESA will use a 4+1 data placement scheme. This results in the capacity for data to be stored resiliently will only consume 1.25x the size of the original object. This is much better space efficiency for storing data resiliently than the 2x required by a RAID-1 mirror.
- **3-5 hosts.** When RAID-5 is applied to objects in a cluster of this size, the vSAN ESA will use a 2+1 data placement scheme. This results in the capacity for data to be stored resiliently will only consume 1.5x the size of the original object. This is much better space efficiency for storing data resiliently than the 2x required by a RAID-1 mirror, and you get to enjoy this guaranteed level of space savings with just a 3-host cluster.*

** While the new RAID-5 erasure code in the ESA allows it to be used on a cluster as small as 3 hosts (similar to the minimum for a RAID-1 mirror), VMware still recommends sizing a cluster host count to one additional host beyond the requirements of the storage policies used (e.g. 4 host cluster for FTT=1 using RAID-1 or RAID-5)*

For more information on the new RAID-5 in ESA, see the post: Adaptive RAID-5 Erasure Coding with the Express Storage Architecture in vSAN 8.

I see vSAN objects have more components now. Should I be concerned with that?

The vSAN ESA uses a modified data structure found in the original storage architecture of vSAN. As a result, objects in a vSAN ESA-powered cluster will typically have more components than the same objects in a vSAN OSA-powered cluster. To help accommodate for this, the vSAN ESA increases the component count limit from 9,000 components to 27,000 components. In vSAN 8 u2, the limit of VMs per host has increased from 200 VMs to 500 VMs per host when using the ESA. The limit of 200 VMs per host in the OSA remains the same.

What is a "capacity leg" and "performance leg" in the vSAN ESA, and what do I need to know about them?

The ESA uses a modified object format to store data in an object in two legs: The performance leg, and the capacity leg. These are contained all within the same object, and interact with the vSAN log-structured filesystem (LFS). This uses the same underlying distributed object manager to store the payload of data and metadata but does so in a manner that delivers performance and capacity all from the same tier. **It is simply an implementation detail and does not impact the design or operations of an environment.** It is a great example of how VMware integrated a new, optional architecture within an existing code base.

I'm looking at a RAID-1 object in my vSAN ESA cluster, and I don't see any witness components. Where did they go?

When using the vSAN ESA, Storage policies that use RAID-1 mirroring will not have a dedicated witness component for an object, except for in 2-node and stretched cluster topologies, where a virtual witness host appliance is used. Quorum for these RAID-1 objects is determined by the voting of components that comprise both the capacity leg and the performance leg of an object, thus a witness component is no longer necessary.

I see compression capabilities in the ESA is provided as a storage policy rule, as opposed to a cluster-based service in the OSA. Can I turn it on and off at will?

Yes, but you may not see all of the benefits of space efficiency if doing so. Compression is enabled by default, and can be turned off through a storage policy rule. Changing the policy rule will only affect new I/Os written. It will not retroactively inflate or compress existing blocks.

For more information on compression in the ESA, so the post: vSAN 8 Compression - Express Storage Architecture

When should I not use compression in the vSAN ESA?

Most workloads will benefit from the use of vSAN's compression. However, some corner-case workloads, such as video, PostgreSQL databases, and other applications may use their own compression capabilities. In these cases, using a storage policy with the compression capability turned off will save CPU cycles.

It is said that the ESA in vSAN 8 can have up to 4x better compression than the OSA. Will I see this benefit in my own environment?

Compression is an opportunistic space efficiency feature, so the benefit seen on real workloads will vary. The improvement refers to the theoretical maximum that vSAN can compress a 4KB data block. The vSAN OSA could theoretically compress a data block 2:1, or 2x. The vSAN ESA can theoretically compress a data block as little as 8:1, or 8x, with more granularity (e.g. 7:1, 6:1, 5:1, etc.). This is on a per data block basis. The benefits that real workloads will see will depend on the type of data stored. Unlike the OSA, the vSAN ESA's compression will reduce network bandwidth as well, since it operates higher up in the stack.

For more information on compression in the ESA, see the blog post: vSAN 8 Compression - Express Storage Architecture

The ESA in vSAN supports encryption. Does this mean that it supports data-at-rest or data-in-transit encryption, or both?

The vSAN ESA encrypts data high in the storage stack, prior to writes being sent to other hosts. This means that when vSAN encryption is used with a cluster running the ESA, it will be encrypted in-flight, and at rest. We still provide the option to enable Data-in-Transit encryption for clusters using the ESA to ensure that all encrypted packets remain unique. The ESA also eliminates previous decrypt, and re-encrypt processes, which reduces overhead.

For more information, see the post: Cluster Level Encryption with the vSAN Express Storage Architecture

How much faster will the vSAN ESA be than the OSA?

This is a constantly evolving answer, but a good comparison of performance of the vSAN ESA and OSA in vSAN 8 U1 is provided in the blog post: vSAN 8 U1 Express Storage Architecture - Faster than Ever.

The answer will depend on hardware configurations and workloads, and if synthetic tests are used, versus real-world workloads. For the latter, the measure of success will be monitoring the level and consistency of the latency, as provided by the vSAN performance service. Synthetic testing can be a useful exercise for stress tests but provide limited functional benefit for translating how real workloads will behave. An illustration of how synthetic generators compare to real workloads can be found in the blog post: Performance when using vSAN Encryption Services.

Achieving optimal performance will also be easier to do with the ESA. For example, one will not need to adjust the stripe width storage policy rule to achieve better performance. For more information, see the post: Stripe Width Storage Policy Rule in the vSAN ESA. For general recommendations to optimize performance with vSAN ESA, see the post: Performance Recommendations for vSAN ESA.

Why does the ESA include the ability to automatically manage vSAN-related network traffic?

With the tremendous efficiency of the vSAN ESA, storage traffic can be processed through the stack at nearly device-level rates. This higher rate of processing I/Os in a server means higher rates of I/O traversing the network, which can potentially lead to the network being the bottleneck when running highly demanding workloads. To help accommodate for this, the ESA in vSAN 8 includes an adaptive traffic shaping capability for vSAN I/O traversing a network. This helps ensure that when network contention occurs, that vSAN will properly prioritize VM I/O over resynchronization activity. This can help deliver more consistent performance for these demanding workloads that may otherwise be saturating a network link.

For more information, see the post: Adaptive Network Traffic Shaping in the vSAN Express Storage Architecture.

How do I see if my new cluster running the vSAN ESA is performing better than my other OSA based vSAN clusters?

Performance can be checked by using the metrics provided by the vSAN performance service. **Monitoring latency as seen by the guest VM is a good way to determine if the storage can meet the demands of the workloads.** VMs with higher latencies on other clusters will have the potential to perform better when running on a cluster powered by the vSAN ESA. A VM demonstrating low, consistent latency is the desired outcome for all workloads. Sometimes if you have applications performing batch processes, the time taken to complete the batch processing can be compared from one cluster to another. A shorter time to complete a batch process is a good indicator of improved performance.

How do I see if my new cluster running the vSAN ESA is more efficient than my other OSA based vSAN clusters?

The "efficiency" of a cluster can be measured two different ways. 1.) The amount of computational resources (CPU) across the cluster for a given set of workloads. 2.) The amount of storage capacity across the cluster to protect the data resiliently. Prior to migrating workloads running on an existing vSAN cluster running the OSA, look for the average CPU consumption across the cluster, then compare it to the CPU consumption on the new cluster. The same can be performed for capacity. Although, note that performing a capacity comparison will only be accurate if you maintain the same FTT level for all VMs.

Why did VMware introduce a new snapshotting capability within the vSAN ESA?

The introduction of the Express Storage Architecture in vSAN 8 means that vSAN can now manage the data in new and interesting ways. Building a new, native snapshotting capability was a great way to exploit the capabilities of the ESA, and help support the needs and use-cases of our customers. Customers will be able to create point-in-time states of data with minimal degradation in the performance of a VM, no matter how many snapshots are taken. The new native snapshot capability is integrated directly in vSphere and fully supports our broad backup partner community with the continued support of VADP backup integration. For more information, see the post: Scalable, High-Performance Native Snapshots in the vSAN Express Storage Architecture.

vSAN 8 U3 introduces more capabilities with ESA snapshots through its vSAN Data Protection capability. More information on vSAN Data Protection can be found in the aforementioned section in this FAQ.

I've always heard vSAN described as analogous to an object store, but I hear the vSAN 8 ESA uses a new file system. I'm confused. What is used in the vSAN ESA?

vSAN (The OSA and ESA) uses a data structure that most analogous to an object store, which is an ideal approach for a distributed storage solution like vSAN. The ESA in vSAN 8 introduces a new "log-structured file system" known as the "vSAN LFS." This does not refer in any way to a traditional file system such as NTFS, ext4, or a cluster file system like VMFS. It is a common industry reference to a method of how data and metadata are written and appended to a circular log buffer and persisted to the storage subsystems. For the vSAN ESA, its LFS helps vSAN ingest data quickly and efficiently and allows data and metadata to be prepared and stored with high levels of efficiency, scalability, and performance.

I see disaggregation is supported when using the ESA in vSAN 8 U1 or later. Are there any features within HCI disaggregation that are unavailable in the ESA?

When using the ESA in vSAN 8 U1, disaggregation has similar capabilities and limits as found when using the OSA in vSAN 8 U1, with limits in support of disaggregation over stretched cluster topologies and connectivity across clusters when using multiple vCenter Server instances. This limitation was eliminated for the ESA in vSAN 8 U2 and newer.

How much capacity overhead is consumed when using the vSAN ESA?

The amount of overhead used is about the same as the vSAN OSA. But may be effectively less as data can be stored in more efficient ways. For a more detailed answer, see the post: Capacity Overheads for the ESA in vSAN 8 as well as Improved Capacity Reporting in VMware Cloud Foundation 5.1 and vSAN 8 U2.

Are there any features or capabilities in vSAN that are not available when using the ESA in vSAN 8 U3?

Yes. The ESA in vSAN 8 U2 has almost eliminated any disparity in discrete features. Deduplication in the ESA is not supported at this time.

Were there any updates made to the OSA in vSAN 8 and newer?

Yes! We have introduced several improvements that apply to both the ESA, and OSA, as well as an improvement that apply to the OSA exclusively at this time.

Express Storage Architecture (ESA) and Original Storage Architecture (OSA):

- vSAN 8 U3. VM I/O Trip Analyzer Cluster Level View.
- vSAN 8 U3. Enhanced network health findings for RDMA-based networks
- vSAN 8 U2. Support of KMS servers using the key expiration attribute.
- vSAN 8 U2. Skyline Health remediation enhancements
- vSAN 8 U2. Top Contributors enhancements
- vSAN 8 U2. I/O Trip Analyzer support of 2-node and stretched cluster topologies
- vSAN 8 U2. Witness traffic separation configuration in the UI.
- vSAN 8 U1. Customized sizing of namespace objects.
- vSAN 8 U1. Skyline Health intelligent cluster health scoring.
- vSAN 8 U1. High resolution performance statistics.
- vSAN 8 U1. VM I/O Trip Analyzer task scheduling.
- vSAN 8 U1. Diagnostics support improvements.
- vSAN 8. Removal of the 32-node flag (and required reboot) on clusters growing beyond 32 hosts.
- vSAN 8. Network uplink metric enhancements.
- vSAN 8. vSAN cluster shutdown enhancements.
- vSAN 8. vSAN Proactive Insights - allowing for cloud-connected health checks without enabling CEIP.

Original Storage Architecture (OSA) exclusively:

- vSAN 8 U1. Support of vSAN Stretched clusters for disaggregation.
- vSAN 8 U1. Support of disaggregated clusters using multiple vCenter Servers.
- vSAN 8. For All-Flash configurations, the OSA in vSAN increases the logical size of buffer devices from 600GB to 1.6TB. For many of our customers who already have larger devices, this can mean the potential for improved levels of performance consistency in their vSAN clusters running vSAN 8 with the OSA. For more information, see the post: Increased Write Buffer Capacity for the vSAN 8 Original Storage Architecture.
- vSAN 8. When running disaggregation using vSAN 8 with the OSA, we've increased the maximum number of client clusters that can be connected to a server cluster from 5 to 10. For more information, see the post: HCI Mesh Scalability Improvements in vSAN 8.
- vSAN 8. When running File Services using vSAN 8 with the OSA, we've introduced several enhancements that improve day-to-day operations and usability.

I want to make sure my new ESA cluster is running as fast as possible. What steps should I take to ensure that I'm getting optimal performance from my ESA cluster?

We have provided a collection of Performance Recommendations for vSAN ESA to help users ensure that they are getting the most out of their environment. Be sure to refer to this resource often, as it will continue to be updated based on the capabilities of the ESA.

Note that with vSAN, **performance of a VM is derived from the host hardware and the network** used to interconnect the hosts in the vSAN cluster, **not the cluster host count**. While increasing the host count of a cluster will increase the aggregate IOPS and bandwidth achieved by the cluster, in most cases it will not improve the discrete performance capabilities observed by the VM. VM performance will be a function of the host hardware and network connectivity. See the post: "Performance Capabilities in Relation to vSAN Cluster Size" for more information.

Does vSAN ESA support cluster updates using vSphere Update Manager, otherwise known as VUM?

No. The **vSphere Lifecycle Manager (vLCM) is the only supported method for cluster lifecycle management when using vSAN ESA.** Skyline Health for vSAN will trigger a health finding/alert when the cluster is not configured for vLCM. For more information, see the KB: vSAN Health Service - vSphere Lifecycle Manager (vLCM) configuration

Availability

What happens if a host fails in a vSAN cluster?

vSAN will wait for 60 minutes by default and then rebuild the affected data on other hosts in the cluster. The 60-minute timer is in place to avoid unnecessary movement of large amounts of data due to a temporary issue. As an example, a reboot takes the host offline for approximately 10 minutes. It would be inefficient and resource-intensive to begin rebuilding several gigabytes or terabytes of data when the host is offline briefly. vSAN will also write all subsequent updated blocks of data to an additional host, in addition to the writes that are being committed to the other replica object. These are known as "durability components" and helps improve durability of data in both planned, and unplanned events. See the vSAN Interactive Infographic for a better understanding on how vSAN handles various conditions in a cluster.

How does vSAN handle a dividing or isolation of parts of a network, known as a network partition?

vSAN uses a quorum voting algorithm to help protect against "split-brain" scenarios and ensure data integrity. An object is available for reads and writes as long as greater than 50% of its components are accessible.

As an example, a VM has a virtual disk with a data component on Host1, a second mirrored data component on Host2, and a witness component on Host 3. Host1 is isolated from Host2 and Host3. Host2 and Host3 are still connected over the network. Since Host2 and Host3 have greater than 50% of the components (a data component and a witness component), the VM's virtual disk is accessible.

However, if all three hosts in our example above are isolated from each other, none of the hosts have access to greater than 50% of the components. vSAN makes the object inaccessible until the hosts are able to communicate over the network. This helps ensure data integrity. See the vSAN Interactive Infographic for a better understanding on how vSAN handles various conditions in a cluster.

What happens if a storage device fails in a vSAN host?

vSAN can not only handle host failures with ease, but also storage device failures. When a device is degraded and error codes are sensed by vSAN, all of the vSAN components on the affected drive are marked degraded and the rebuilding process starts immediately to restore redundancy. If the device fails without warning (no error codes received from the device), vSAN will wait for 60 minutes by default and then rebuild the affected data on other disks in the cluster. The boundary of failure of a single storage device failure will depend largely on whether one is using vSAN OSA or vSAN ESA. For more information, see the post: "The Impact of a Storage Device Failure in vSAN ESA versus OSA."

What happens if there is not enough free capacity to perform all the component rebuilds after one or more host failures?

In cases where there are not enough resources online to comply with all storage policies, vSAN will repair as many objects as possible. This helps ensure the highest possible levels of redundancy in environments affected by the unplanned downtime. When additional resources come back online, vSAN will continue the repair process to comply with storage policies. We recommend maintaining enough reserved capacity for rebuild operations and other activities such as policy changes, etc.

What happens if there are multiple failures (loss of hosts, etc.) that exceed the configured threshold of failures?

Some vSAN objects will become inaccessible if the number of failures in a cluster exceeds the failures to tolerate (FTT) setting in the storage policy assigned to the given object. If a VM object is using a storage policy that uses FTT=2, and three of the hosts that contain this object fail simultaneously, the object will be offline (but not lost) to preserve the integrity of the data. The FTT level does not specify the total failures that the cluster can tolerate, but rather, specifies the failure that the assigned VM object can tolerate. For example, a large cluster can have multiple failures, but if only one of the hosts for a given object is impacted, the data remains available.

When using the vSAN Express Storage Architecture (ESA), use the more resilient FTT=2 using RAID-6 and enjoy the benefits of space efficient storage and high levels of resilience without any performance cost.

How do I protect VMs residing on vSAN?

Many third-party data protection products use VMware vSphere Storage APIs for Data Protection (VADP) to provide efficient, reliable backup and recovery for virtualized environments. These APIs are compatible with vSAN just the same as other datastore types such and VMFS and NFS. Nearly all of these solutions should work with vSAN. Other methods of protecting the data in your vSAN environment can be achieve through solutions like VMware Live Recovery, or VMware Live Site Recovery more specifically.

Protection mechanisms can often be layered on top of each other to address specific use cases. For example. One could have a traditional VADP-based backups to provide backups that follow the typical 3-2-1 rule of protecting data (three copies of data across two types of media, and one copy offsite) paired with something like vSAN Data Protection - introduced in vSAN 8 U3 - that would augment this effort by having readily available, highly efficient snapshots that can be reverted when needed.

Does vSAN work with VMware Live Recovery, specifically VMware Live Site Recovery (VLSR) previously known as SRM?

Yes.

Is there a way to stop vSAN resynchronizations?

vSAN provides ways to gracefully stop resynchronizations, which means that in some cases, existing resynchronizations will finish to completion while others are halted prior to starting. vSAN also has protections in place that will pause resync

operations if disk space usage meets or exceeds critical thresholds. Subsequently, the operations are resumed when sufficient capacity is made available.

How is vSAN impacted if vCenter Server is offline?

vCenter operates as the management plane for vSAN and is the primary interface to manage and monitor vSAN. However, vCenter does not affect the data plane i.e., the VM I/O path. When vCenter Server is offline, vSAN continues to function normally. VMs continue to run, and application availability is not impacted. Management features such as changing a storage policy, monitoring performance, and adding a disk group are not available.

vSAN has a highly available control plane for health checks using the VMware Host Client—even if vCenter Server is offline. Hosts in a vSAN cluster cooperate in a distributed fashion to check the health of the entire cluster. Any host in the cluster can be used to view vSAN Health. This provides redundancy for the vSAN Health data to help ensure administrators always have this information available. One may wish to deploy a management cluster for additional durability of their data center. For more information, see the post: "Using vSAN as a Management Cluster."

Does the vSAN iSCSI Target Service support Windows Server Failover Cluster (WSFC) configurations?

Yes. Additional details are considerations are outlined in the following Knowledge Base article: "Using SQL Server Failover Clustering with vSAN iSCSI Target Service."

What happens if a vSAN cluster loses power?

The architecture of vSAN ensures that writes are always written to qualified persistent media in a redundant way - ensuring the availability and integrity of data. While vSAN strives for durability under the harshest of environments, a data center design that uses redundant power or standby temporary power is always encouraged.

Once power is restored to the entire cluster, the hosts will begin to power on. The initial booting of ESXi may take a little longer than usual to ameliorate service states and data. This time difference will be relatively insignificant if the cluster is using vSAN ESA. When the vSAN hosts complete the boot process, HA will initiate the power-on of VMs previously turned on, and vSAN will begin any resynchronizations to ensure full data resilience and honor the storage policies prescribed to the data.

For environments that only have temporary power during times of a sustained power loss, vSAN now has shutdown workflows in the UI and programmatically to help provide for a speedy shutdown. See the post: "Automation Improvements using PowerCLI 13.1 with vSAN 8 U1" for more information.

Does vSAN store data in a crash consistent manner?

Yes. vSAN ensures that data is persisted to disk before write acknowledgements are ever sent back to the guest VM. It is also written in a resilient way defined by the resilience setting in the prescribe storage policy.

Upon an unexpected failure, such as a power loss of a host where some of the VM data resides, the VM will continue to write data without issue, but in a less resilient way until the data is automatically reconstructed elsewhere. If a power loss occurred on a host where the VM instance was running, then vSphere HA will restart that VM elsewhere, and work with the data persisted to disk.

When a VM is protected (backed up) using our APIs for data protection (3rd party solutions using VADP, or vSAN Data Protection built into vSAN 8 U3), the data and its data structure is persisted to disk in a state that is preserved at the time of the crash. This may or may not include preserving the memory state of the VM. This is what is referred to as "crash consistency." When referring to three-tier architectures using traditional storage arrays, the additional step of a "stun" may need to occur on all VMs residing in a LUN. This would pause the I/O so that the storage array could create a LUN-based snapshot in a consistent manner. For crash-consistent backups using vSAN Data Protection, **a VM stun is unnecessary for vSAN** as we are fully aware of the I/Os and if they are persisted to disk. This makes snapshots in vSAN Data Protection fast, efficient, and scalable. For more information, see the post: "Superior Snapshots using VMware vSAN Data Protection."

Crash consistency is what most backup vendors provide, because it is purely working with the data that has been fully committed to disk. It is sufficient in most cases. "Application consistency" is a more sophisticated step that involves quiescing the data and application state perform a coordinated flush to disk to ensure that changes aren't currently being held in

memory. Some Operating Systems have mechanisms in place that allow for application or file based consistency (such as Microsoft's Volume Shadow Copy Service, or VSS), while other operating systems must use pre and post backup scripts to quiesce an application. Backup applications often need to coordinate with these mechanisms to provide application consistent backups.

The ability for vSAN to provide crash consistency applies to both the VM object data, as well as any associated snapshots using VADP, or vSAN Data Protection introduced in vSAN 8 U3.

Cloud-Native Storage

What is Cloud-Native Storage, or CNS?

Cloud-Native Storage(CNS) is a term used to describe the storage that can be provisioned to Cloud-Native Applications (CNAs). These CNAs are typically containerized, deployed and managed by a Container Orchestrator like Kubernetes, Mesos, Docker Swarm, etc. The storage consumed by such apps could be ephemeral or persistent, but in most cases, it is required to be persistent. CNS is supported when using the vSAN OSA, and when using the ESA in vSAN 8 U1 (block based RWO volumes only) and both block and file when using the ESA in vSAN 8 U2.

What is a Container Storage Interface, or CSI?

Container Storage Interface (CSI) is a standardized API developed for container orchestration platforms to interface with storage plugins. This API framework enables vSAN & vVOLS to be able to provision persistent volumes to Kubernetes based containers running on vSphere.

Can a vSAN datastore be used to provision persistent storage for a Kubernetes cluster?

Yes, vSAN supports provisioning persistent volumes to Kubernetes based workloads. A brief walkthrough is available here - Cloud-Native Storage on vSAN.

vSAN 7 supports block-based (RWO) or file-based (RWM) persistent volumes when using a standard vSAN cluster. In vSAN 7 U3, block-based persistent volumes are supported when using a vSAN stretched cluster.

How can Kubernetes administrators provision appropriate storage intended for respective containers on vSAN?

Kubernetes administrators can simply associate "storage classes" of the respective containers to Storage Policies. vSAN uses the standard Storage Based Policies Management(SPBM) to provision persistent volumes on the vSAN datastore.

Data Persistence Platform (DPp)

What is the vSAN Data Persistence platform (DPp)?

The vSAN DPp – introduced in vSAN 7 U1 and enhanced in later editions - is a framework within vSphere and vSAN that simplifies and deployment, operation, and optimizations for modern stateful applications and services built for cloud native architectures, running on vSAN. It is supported when using the vSAN OSA, and the ESA when using vSAN 8 U1 and newer.

Why was vSAN DPp created?

Traditional applications rely on the feature-set of an enterprise storage system and hypervisor to provide data and application availability and resilience. Stateful applications are built differently. They often use a shared-nothing architecture (SNA), are typically responsible for their own levels of data and application availability and resilience and have the mechanisms in place to do so. The vSAN DPp provides an ability for these apps to maintain these responsibilities while ensuring that it can coordinate with the underlying hypervisor and distributed storage system (vSAN) for events such as maintenance, failures, and expansion.

What applications can use the vSAN DPp?

The only applications or services capable of using the vSAN DPp will be those Independent Software Vendors (ISV) who have built and submitted an operator (and an optional vCenter plugin) that is certified by VMware. The certified operators available

for installation on a given version of vSphere will be listed in vCenter by highlighting the cluster and selecting "**Configure > Supervisor services**."

What is an operator, and who installs and updates the operators?

In this context, an operator is a small piece of software that provides application control. This provides and interface for bidirectional communication with the underlying platform: VMware vSAN. The 3rd party operators used by DPp will be enabled and updated by the vSphere Administrator, and done so on an as-needed basis.

Where can the 3rd party operators be downloaded?

In versions up to and including vSAN 7 U2, the operators are integrated into the installation, eliminating the need to download operators. With vSAN 7 U3, the installation process is asynchronous, meaning that the operators are not bundled with the installation of vCenter, and can be installed and updated to the latest edition at any time by the administrator.

Where do the operators run?

All 3rd party operators run in the supervisor cluster.

vSAN Direct Configuration

What is vSAN Direct Configuration, and how does it relate to the vSAN DPp?

The vSAN DPp provides two ways for eligible applications to consume storage provided by a vSAN cluster. Both types provide affinity or "colocation" of the application state and the storage it consumes (through persistent volumes) to the host powering the workload. It is available in the vSAN OSA.

- **vSAN SNA.** Storage consumption is through a single shared vSAN datastore, using a traditional vSAN data path, and providing the typical suite of storage services provided by vSAN such as availability, resilience, space efficiency, and data services.
- **vSAN Direct Configuration.** Storage consumption through a collection of discrete devices on the vSAN hosts. Direct access is provided to those ESXi storage devices which provides optimized capacity consumption and improved TCO.

Note that vSAN Direct Configuration predates the technologies introduced in vSAN ESA. It is not supported in vSAN ESA.

For eligible applications, what are the key considerations when determining whether to use a cluster providing storage using "vSAN SNA" versus "vSAN Direct Configuration?"

Clusters providing storage courtesy of the vSAN SNA can be used for modern applications using the vSAN DPp, and traditional workloads powered by VMs. They can also provide the features associated with a traditional vSAN datastore if those capabilities are not provided by the shared nothing application. Clusters providing storage courtesy of the vSAN Direct Configuration can only be used for applications with approved operators and are subject to additional restrictions on hardware devices used. Clusters using vSAN Direct Configuration can offer supreme levels of efficiency for shared-nothing applications and can offer anti-affinity of data across devices on a given host.

Will all DPp applications and services need to rely on application-level replication for resilience, or can vSAN's resilience through storage policies be used?

All applications and services running through vSAN DPp and residing on a cluster running vSAN Direct Configuration must rely on application-level replication to provide resilience. All applications running through vSAN DPp and residing on a cluster running vSAN SNA can rely on application-level replication for resilience (preferred), or use data-level resilience through vSAN storage policies.

If I have a traditional application running in a VM, can I use a cluster providing storage courtesy of vSAN Direct Configuration?

No. Only those applications with their own 3rd party operator are eligible to use storage courtesy of vSAN Direct Configuration.

vSAN File Services

How is vSAN File Services integrated into vSAN?

vSAN File Services is powered and managed by the vSphere platform that deploys a set of containers on each of the hosts. These containers act as the primary delivery vehicle to provision file services and are tightly integrated with the hypervisor. vSAN file services support up to 100 shares per cluster up to vSAN 8 U2. In vSAN 8 U3, using the ESA, vSAN file services supports up to 250 shares per cluster.

Can I run VMs on top of a file share provided by vSAN File Services?

No, it is not supported to mount NFS to ESXi for the purpose of running virtual machines. The NFS shares may be used to mount NFS directly to virtual machines running on the vSAN cluster, but may not be used to store VMDKs for running virtual machines.

In vSAN 8 U1, vSAN offers a new namespace object that can have a customizable size. This can be used as a way to store vSphere content libraries and ISOs.

What is the minimum number of hosts required in a cluster to deploy vSAN File Services?

For a standard vSAN cluster, a minimum of 3 hosts is required to configure vSAN File Services. It will run with as many as 2 remaining hosts. vSAN File Services will auto-scale 1 container per host up to 64 per cluster. vSAN File Services will work on a 2-Node cluster however, which only has two data nodes.

Is vSAN File Services supported on a stretched cluster and 2-Node cluster?

Yes, as of vSAN 7 U2, both topologies are supported.

What is the estimated resource overhead of each host when running vSAN File Services?

One protocol services VM instance runs on each host in the cluster. Each VM instance is configured with 4GB of RAM and 4vCPU. By default, there are no reservations applied to the resource pool associated with the entities required for vSAN File Service.

How is vSAN File Services monitored?

vSAN File Service can be monitored with vSAN Skyline Health Services. A new health check called "File Service - Infrastructure health" monitors several parameters and includes an automated remediation option.

What protocols and authentication methods are supported?

NFSv3, NFSv4.1, SMB v2.1 and SMBv3 are supported. Both NFS and SMB file shares are now able to use Kerberos based authentication when using Microsoft Active Directory.

Can a single share provide access using NFS and SMB at the same time?

Simultaneous access of a single share using both NFS and SMB is not supported at this time.

How can snapshot functionality be used in vSAN File Services?

The snapshot mechanism introduced in vSAN 7 U2 allows for our Backup Partners and Independent Software Vendors (ISVs) to provide point-in-time backup and recovery capabilities of files in vSAN file service shares into their backup solutions. The full capabilities of the snapshot mechanism are accessible via API. Using file share snapshots as a backup source requires backup products supporting the functionality. Backup vendors are currently working to provide support for vSAN file shares. Until then, organizations can use PowerCLI and backup vendor PowerShell modules to add newly created snapshots as a backup source.

Do I need to migrate or manage the file services VMs?

No. vSAN automatically manages the File Server VM. The containers are automatically shut down and removed. It will be recreated once an available host is no longer in maintenance mode.

Do I need to create or add vmdks or objects to expand storage to vSAN File Services?

vSAN File service uses elastic scalability and will create additional components as needed without any manual intervention.

How can I limit the consumption of shares provided by vSAN File Services?

Soft and hard share quotas can help manage capacity consumption.

- **Hard quotas** prevent users from writing data to disk. Hard quotas automatically limit the user's disk space, and no users are granted exceptions. Once users are about to reach their quota, they must request help.
- **Soft quotas** send alerts when users are about to exceed disk space. Unlike hard quotas, there is no physical restriction to prevent users from saving their data. However, you do get alerts and can create a corporate policy to help manage data.

Can I provision file shares to Cloud-Native workloads?

Yes, vSAN File Services can be used to provision file shares to container workloads as well as traditional workloads.

How do NFS shares recover from host failure or migrate during upgrades?

vMotion and vSphere HA are not used as part of migration, or failure recovery. Services within vSphere monitor for failure or maintenance activities and drive the relocation of services. The containers powering vSAN file services, automatically restart on other hosts, independent of vSphere HA.

While by default you will have one container per host, additional containers will run on a host in a case where a host (or hosts) have failed. When a host enters maintenance mode the container powering a given share or group of shares is recovered on a different host.

How is vSAN File Services updated?

An updated OVF can be automatically downloaded or manually updated to the vCenter managing the cluster. A non-disruptive rolling upgrade will proceed across the cluster replacing the old containers with the new version.

Disaggregated Storage using vSAN Storage Clusters

What is disaggregated storage in vSAN?

Disaggregated storage in vSAN uses a unique software-based approach to disaggregate, or decouple compute and storage resources. **vSAN HCI aggregates storage and compute resources into a single cluster, while disaggregation decouples storage and compute resources.** First introduced in vSAN 7 U1, it has since been enhanced up to and including vSAN 8 U3. In the latest version of vSAN, datastore sharing using traditional vSAN HCI clusters allows the datastore of a vSAN HCI cluster (a "server cluster") to be mounted by other vSAN HCI clusters (considered a "client cluster") for the purposes of consuming storage resources. Disaggregation with vSAN storage clusters (previously known as vSAN Max) is a topology where a dedicated cluster is providing storage resources only to one or more vSphere clusters. This provides full independence of scaling between your compute resources and your storage resource, but still using all of the technology of vSAN ESA with consistent unified management across all configuration types.

As of vSAN 8 U3 and VMware Cloud Foundation 5.2, vSAN storage clusters can be used as principal storage. It supports all ReadyNodes certified for vSAN storage clusters, as well as all cluster sizes recommended. At this time, it can be used for workload domains. The use of vSAN Max as principal storage in a management domain is not supported.

Note that when the server cluster is powered by vSAN ESA, only client clusters that are running vSAN ESA can mount the datastore of a server cluster. **At this time, client clusters running vSAN OSA cannot mount the datastore of a server cluster powered by vSAN ESA.** vSphere clusters running vSphere 8 or later can mount the datastore of a server cluster that is running vSAN OSA, or vSAN ESA, but not both technologies at the same time.

What capabilities does disaggregation in vSAN add to data center environments?

Full disaggregation of compute and storage resources using vSAN Max storage clusters provide dedicated storage clusters for vSphere clusters, or even vSAN HCI clusters. A vSphere cluster participating in this topology would be referred to as a "compute cluster" and would mount a remote vSAN datastore to consume the storage resources of that remote vSAN datastore. It uses native vSAN protocols for supreme levels of efficiency and functionality.

Another deployment option is "vSAN HCI with datastore sharing." It allows a vSAN HCI cluster to consume resources from other vSAN HCI clusters. This helps use stranded capacity between clusters and use unique data services or hardware capabilities provided by a given cluster (such as Data-at-Rest Encryption). It also allows administrators and architects the ability to scale compute and storage independently, easing design and operational complexities. For more information on which deployment option may be best for your environment, see: "vSAN HCI or vSAN Max - Which Deployment Option is Right for You?"

What are vSAN storage clusters? (aka vSAN Max)

vSAN Max is a new deployment option in vSAN that provides highly flexible disaggregated storage for vSphere clusters. It is powered by the vSAN Express Storage Architecture, or ESA. It provides our customers the ability to deploy a highly scalable storage cluster to be used as primary storage for vSphere clusters, or augment storage for traditional vSAN HCI clusters. This makes vSAN the premier storage platform for powering VMware Cloud Foundation. See the post "Introducing vSAN Max" for more information.

What would be some common use cases for vSAN storage clusters?

Some common uses cases would include, but are certainly not limited to the following:

- **Cost optimization for infrastructures and applications.** It is not unusual for customers to want to configure clusters to help minimize application costs. vSAN Max will allow customers to right size their compute resources to minimize these licensing costs. One can tailor the vSAN Max cluster to whatever the business needs in terms of performance, capacity, and data services provided.
- **Operational simplicity with unified storage.** One can use vSAN Max to extend the life of blade servers and other older hardware that was not ideal for HCI. If you've wanted to take advantage of vSAN, but wanted to retain compute and storage resources independent from each other, vSAN storage clusters is for you. It can easily serve as the shared storage platform for your data center, and be scaled easily and incrementally as your needs grow. It is a true, distributed scale-out storage solution that can be used as centralized shared storage by any of your vSphere clusters. assets.
- **Rapid Scaling for Cloud Native applications.** Cloud native applications need to be able to scale, which means the underlying infrastructure that supports it also must scale easily, incrementally, and economically. vSAN Max can be an ideal storage platform for cloud native applications.

How are vSAN storage clusters licensed?

Licensing for storage in VMware Cloud Foundation is based on two elements, 1.) The number of CPU cores licensed, and 2.) the amount of additional capacity (per TiB) purchased. For every core licensed in VMware Cloud Foundation, a complimentary 1 TiB of vSAN capacity is provided. Any additional capacity beyond the sum capacity provided by the core count would be through a capacity add-on license.

Which deployment option you choose is entirely up to you. While one may be preferable over the other based on your environment, the new licensing model for VMware Cloud Foundation largely removes licensing as a factor for considering one deployment option for another. As shown below, when a vSAN HCI cluster is provisioned, all cores and capacity licensing would reside on that cluster. But if a disaggregated approach was chosen, some licensing cores will reside on the vSphere clusters, while other licensing cores and capacity reside on the vSAN Max cluster. When using a disaggregated deployment option, **fewer cores are needed for each respective vSphere host, since vSphere hosts are only responsible for running VM instances, and vSAN Max hosts are only responsible for processing storage.**

Isn't a vSAN storage cluster jus a vSAN HCI cluster without any running VM instances?

Even though vSAN storage clusters share many characteristics to a vSAN HCI cluster, they are not the same. When a vSAN storage cluster deployment option is chosen, special tunings are automatically made that change the storage stack to accommodate for the fact that the vSAN storage cluster is going to be predominantly processing I/O and storing data, not running a substantial amount of VM instances. The tuning helps maintain more metadata in memory, which helps both read and write operations.

Are ReadyNodes certified for vSAN storage clusters the same as ReadyNodes certified for vSAN HCI?

No. vSAN Max will use special vSAN ReadyNode profiles that shares similarities to ReadyNodes for vSAN HCI clusters, but with much higher capacities. For more information, see the post: "ReadyNode Profiles Certified for vSAN Max." Additional flexibility in supported server configurations and host counts in a vSAN Max were also introduced recently. See the post: "Greater Flexibility with vSAN Max through Lower Hardware and Cluster Requirements."

Are there general design, sizing, and other recommendations for vSAN storage clusters?

Yes, the "vSAN Storage Clusters Design and Operational Guidance" document covers planning and sizing requirements and considerations, Day-0 deployment and configuration guidance, and Day-2 Operational guidance to help guide you through the process of providing a high performing, robust storage environment using vSAN storage clusters.

What versions of vSphere can be used to connect to a vSAN storage cluster datastore?

The hosts in a vSphere cluster attempting to mount a vSAN storage cluster datastore must be running vSphere 8 or later. vSAN clusters connecting to a vSAN storage cluster datastore must be running vSAN 8 or later, and MUST be using vSAN ESA.

vSAN storage clusters are built using vSAN ESA. When the server cluster (either vSAN Max deployment option, or a traditional vSAN HCI deployment option) is powered by vSAN ESA, only client clusters that are running vSAN ESA can mount the datastore of a server cluster. At this time, client clusters running vSAN OSA cannot mount the datastore of a server cluster powered by vSAN ESA.

How much additional CPU and memory is required for hosts in a vSphere cluster to communicate with a vSAN storage cluster?

The vSAN client service activated on a vSphere cluster is purpose-built for the role of communicating with a server cluster. The resources required on the hosts in the client cluster are relatively small, but will vary based on a number of factors, including the number of server clusters mounted, features used, workload activity, and whether the client service is for use with an OSA cluster, or ESA cluster. One may see up to about 17GB of memory used per host in the client cluster connecting to a vSAN Max cluster, which is powered by the ESA. Typically, the CPU demands on the hosts in the client cluster will be less than one core.

What type of cluster types and connectivity are supported with vSAN storage clusters?

The vSAN client service activated on a vSphere cluster is purpose-built for the role of communicating with a server cluster. The resources required on the hosts in the client cluster are relatively small, but will vary based on a number of factors, including the number of server clusters mounted, features used, workload activity, and whether the client service is for use with an OSA cluster, or ESA cluster. One may see up to about 17GB of memory used per host in the client cluster connecting to a vSAN storage cluster, which is powered by the ESA. Typically the CPU demands on the hosts in the client cluster will be less than one core.

What cluster types and connectivity are supported with vSAN storage clusters?

vSAN Max clusters can be configured either as a single site cluster, or stretched across two sites to provide site-level resilience. Client cluster types can include vSphere clusters (also known as vSAN compute clusters), vSAN HCI clusters, and vSAN HCI clusters that are stretched or in 2-node topologies. These can be managed by the same vCenter Server instance, or by different vCenter Server instances. For more information, see the post: "Flexible Topologies with vSAN Max."

There are certain limitations with client cluster connection types when mounting a datastore from a vSAN Max cluster or vSAN HCI cluster that is in a stretched cluster configuration. In general, **vSphere clusters stretched across two geographic sites is currently not supported when using a vSAN storage cluster or vSAN HCI cluster in a stretched cluster configuration**. For a detailed list of supported client configurations of vSAN Max in a stretched topology, see the table below.

Client Cluster Type	Server Cluster Type	Supported	Notes
vSAN HCI clusters (ESA) in a stretched cluster configuration.	vSAN storage cluster or vSAN HCI cluster (ESA) in a stretched cluster configuration	Yes	Provides resilience of data and high availability of running VM instances.
vSAN HCI clusters (ESA) when it resides in one of the data sites that the vSAN Max cluster resides.	vSAN storage cluster or vSAN HCI cluster (ESA) in a stretched cluster configuration	Yes	Provides resilience of data but no high availability of running VM instances.
vSphere clusters stretched across two sites using asymmetrical* network connectivity.	vSAN storage cluster or vSAN HCI cluster (ESA) in a stretched cluster configuration	No	Not supported.
vSphere clusters stretched across two sites using symmetrical* network connectivity.	vSAN storage cluster or vSAN HCI cluster (ESA) in a stretched cluster configuration	Yes	Supported, but less common, as it would require the same network capabilities (bandwidth and latency) between fault domains defining each site.
vSphere clusters when it resides in one of the data sites that the vSAN Max cluster resides.	vSAN storage cluster or vSAN HCI cluster (ESA) in a stretched cluster configuration	Yes	Provides resilience of data but no high availability of running VM instances.
Any type of client cluster running **vSAN OSA**	vSAN storage cluster or vSAN HCI cluster (ESA) in a single site or stretched cluster configuration	No	**Not supported.**

* Asymmetrical network connectivity would describe a topology where the network capabilities (latency & bandwidth) connecting the two sites (fault domains) would be less than the network capabilities between the client cluster and the server cluster within each site. This is most common with stretched cluster configurations using and inter-site link (ISL) between sites. Symmetrical network connectivity would describe a topology where the network capabilities connecting the two sites would be the same as the network capabilities between the client cluster and server cluster within each site. This configuration is less common, but might be found in environments where the two fault domains defining the sites in the stretched topology are simply server racks or rooms sitting adjacent to each other using the same network spine.

Aren't traditional storage arrays "disaggregated?" And if so, how is this any different?

Yes, a traditional three-tier architecture using a storage array provides storage resources that are "disaggregated" from compute resources. But that approach has inherent limitations. The design of a monolithic storage array connected to a dedicated storage fabric using a pair of redundant controllers must funnel all I/O through those controllers to a clustered filesystem like VMFS, then rely on locking mechanisms to prevent simultaneous access. vSAN Max, built off of vSAN ESA is

quite different. It is a fully distributed storage system that because of its design, offers incremental scalability of both capacity and performance, a common management plane, all running on commodity servers. This helps overcome most of the challenges with a traditional three-tier architecture using a storage array.

See the post: "vSAN Max and the Advantage of Scalability" for more information.

Can disaggregation be used to maintain cluster homogeneity of server vendors?

Absolutely! Disaggregation will allow you to keep servers from the same manufacturer in the same cluster, while consuming that storage from other vSphere or vSAN clusters. This can be ideal for organizations using multiple server vendors.

How is disaggregation with vSAN different than composable/modular infrastructures?

vSAN's disaggregation capabilities uses a software-based approach to disaggregation that can be implemented on any certified hardware. Composable infrastructure is based-off a hardware-centric approach that requires specialized hardware. Unlike other solutions, VMware treats the disaggregation at the cluster level. This helps avoid the challenges associated with stop-gap approaches such as storage-only nodes found with other solutions.

Recently, VMware partnered with Samsung to develop a proof of concept using vSAN Max that connected to dedicated NVMe-based JBOF enclosures to extend the concept of disaggregation. This is still a proof of concept at this time, and not commercially available.

Which protocol and data path does disaggregation in vSAN use?

It uses vSAN's native protocol and data path for cross-cluster connections, which preserves the vSAN management experience and provides the most efficient and optimized network transport. For example, storage policy-based management (SPBM) and the vSAN management interfaces are available on the client cluster, where the remote vSAN datastore essentially resembles a local vSAN datastore.

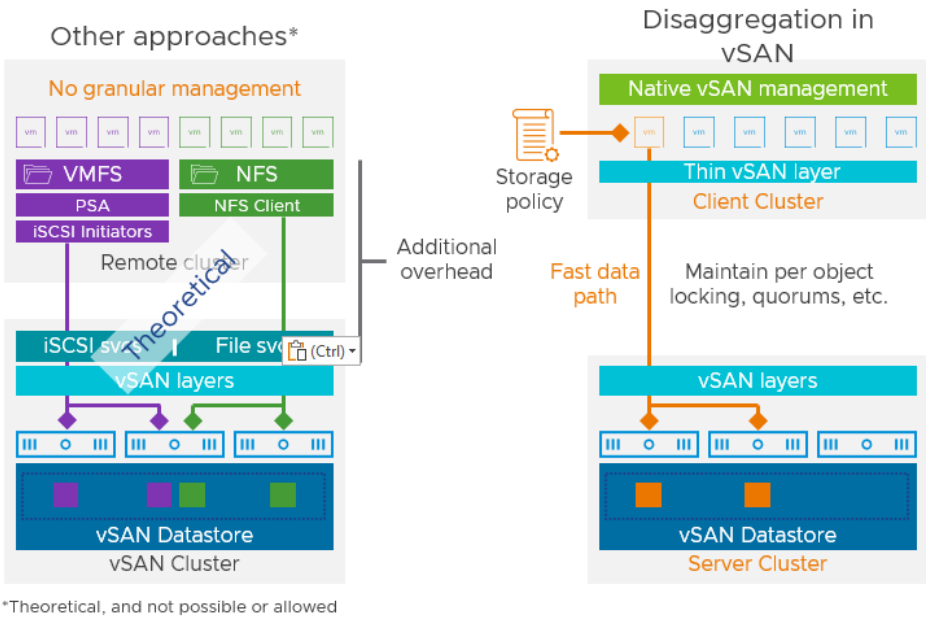

Using an approach as illustrated on the left side of the image above would have a complex data path that would inhibit performance, efficiency, scalability, and compatibility objectives. vSAN's stack was developed for the specific needs of a distributed system and has been adapted to a disaggregated architecture.

Can a vSAN storage cluster be mounted to vSphere clusters?

Yes, this is the primary use case, where vSAN storage clusters provides the centralized shared storage resources for vSphere clusters. When configuring a vSphere cluster for connection to a vSAN Max cluster, a thin layer of vSAN is activate on the

hosts participating in the client cluster and is used to communicate natively with the server cluster that is providing the storage capacity and services. For communication, disaggregation in vSAN uses native vSAN protocols for supreme levels of efficiency and functionality.

Do hosts in client clusters and the vSAN storage cluster need to be using the same CPU manufacturer?

No. The CPU manufacturer (e.g. Intel, AMD) used in hosts in a client cluster can be different than the CPU manufacturer of the hosts that comprise a vSAN storage cluster. For example, a vSphere client cluster using AMD CPUs can mount the datastore of a vSAN storage cluster comprised of hosts using Intel CPUs. The vSAN storage cluster must use hardware certified for use with vSAN storage clusters, while a client vSphere cluster only needs to adhere to the basic hardware compatibility list for vSphere.

Do hosts in client clusters (vSphere clusters) need to be certified vSAN ReadyNodes?

No. Hosts in client clusters that are mounting an external vSAN **datastore only need to comply with the VMware Compatibility Guide (VCG) for vSphere, not the VCG for vSAN.**

We do recommend that network connectivity between a client cluster and server cluster be as fast as possible, since this represents the fabric that storage traffic is transmitted. See the document "vSAN Storage Cluster Design and Operational Guidance" for further information on connectivity recommendations from client clusters to server clusters powered by vSAN Max.

What are some of the scaling capabilities with vSAN storage clusters?

Just like vSAN HCI clusters, vSAN storage clusters can be scaled incrementally to add more capacity and performance. Scale up by adding more or higher density storage devices in the hosts or scale out by adding more storage nodes. Cluster capacities are based on what is available in ReadyNodes certified for vSAN storage clusters. A client cluster can mount up to a maximum of 5 remote vSAN datastores, and a server cluster can export it's datastore up to a maximum of 10 client clusters. Up to 128 hosts can connect to a remote vSAN datastore - when counting the hosts from the client cluster(s) and the server cluster.

VMware has guidance on various design maximums. See the document "vSAN Storage Cluster Design and Operational Guidance" for more information.

Can VMs be provisioned to span across multiple remote datastores?

All objects related to VM (VMDKs, VM Home, etc.) **must reside on a single datastore** - which can be either local or remote. For example, a VM cannot have one of its objects residing on a local datastore, while another one of its objects resides on a remote datastore.

Does disaggregation integrate with other vSAN features?

Yes. For example, a vSAN storage cluster can provide resources to one or more client clusters, and be configured as a single-site vSAN storage cluster, or a highly resilient stretched cluster. The cluster can provide common data services such as encryption and data-at-rest encryption.

What are the network recommendations to implement disaggregation with vSAN?

Network requirements and best practices are very similar to vSAN Fault Domain configurations where data traffic travels east-west across multiple racks in the data center. In general, low-latency and high bandwidth network topologies are recommended for optimal performance. Sub-millisecond latency between two clusters is recommended to provide the most optimal workload performance, however higher latencies will be supported for workloads that are not latency sensitive. L2 or L3 topologies are supported, similar to other vSAN configurations. See the document "vSAN Storage Cluster Design and Operational Guidance" for more information specific to connectivity to a vSAN storage cluster.

More generalized guidance for vSAN HCI to vSAN HCI cluster disaggregation includes the following:

- Design for redundancy everywhere for highest availability (multiple links, NICs, TOR/spine switches, etc.)

- For all but the very smallest of vSAN storage cluster ReadyNode profiles, use NICs and storage-class switches capable of at least 25Gbps for connectivity between the hosts that make up the vSAN storage cluster. The **vSphere clusters do not need to meet this minimum**, but faster networking will help ensure that it is not the performance bottleneck.
- vSphere clusters ("compute clusters") acting as the client cluster can use 10GbE connectivity to the server cluster. This is to accommodate older server types such as blades where customers want to extend the life of the servers in use. However, we still highly recommend 25GbE or 100GbE for client cluster to server cluster communications.
- Use NIOC with vSphere Distributed Switches
- LACP offers benefits but is operationally complex; LBT or Active/standby failover is a simpler and more appropriate alternative in a large majority of cases.
- The mounting of a vSAN storage cluster or vSAN HCI datastore will run a precheck and look for network connectivity between the client clusters and the server cluster of 5ms or less. Since this connection represents storage traffic, it is best to strive for 1ms or less for network connectivity between client clusters and a server cluster.

Note: Currently, the use of multiple vSAN VMkernel interfaces i.e. Air-gapped network topology is not supported when using disaggregation with vSAN.

Are there any availability considerations with disaggregation in vSAN?

vSAN's disaggregation offering uses existing vSphere HA and vSAN availability concepts to provide both compute and storage high availability. vSphere HA will provide compute availability on the client cluster, and vSAN storage policies with FTT=N configured will provide storage availability on the server cluster. If the network connection between a client and server cluster is severed, the remote vSAN datastore on a client host will enter APD 60 seconds after the host becomes isolated from the server cluster. After that, vSphere will follow the current HA mechanisms for APD events and attempt to restart VMs after a 180 second delay.

Are storage policies integrated with disaggregation?

Yes, all VMs consuming storage on vSAN on a vSAN HCI cluster, or a vSAN storage cluster, are controlled by storage policy based management (SPBM). Simply apply the desired policy to the VM, and vSAN does the rest. Recent improvements with SPBM provide the ability to define a data service to a policy. This will allow an administrator to provision a VM, assigning it a storage policy that perhaps defines the use of Data-at-rest Encryption, and vSAN will present a filtered listed of vSAN powered, disaggregated datastores that ONLY provide that service. The data services available for definition within a storage policy in vSAN 7 U2 and later are "Data-at-Rest Encryption" "Space Efficiency" (DD&C, or compression-only), and the storage tier type (all-flash or hybrid, relating to the OSA).

vSAN's disaggregation capabilities even allows for the mounting of a vSAN remote datastore by a cluster that is managed by another vCenter Server instance. For example, a vSphere cluster managed by vCenter server instance "VC01" can mount the remote datastore served up by a vSAN Max cluster managed by vCenter server instance "VC02."

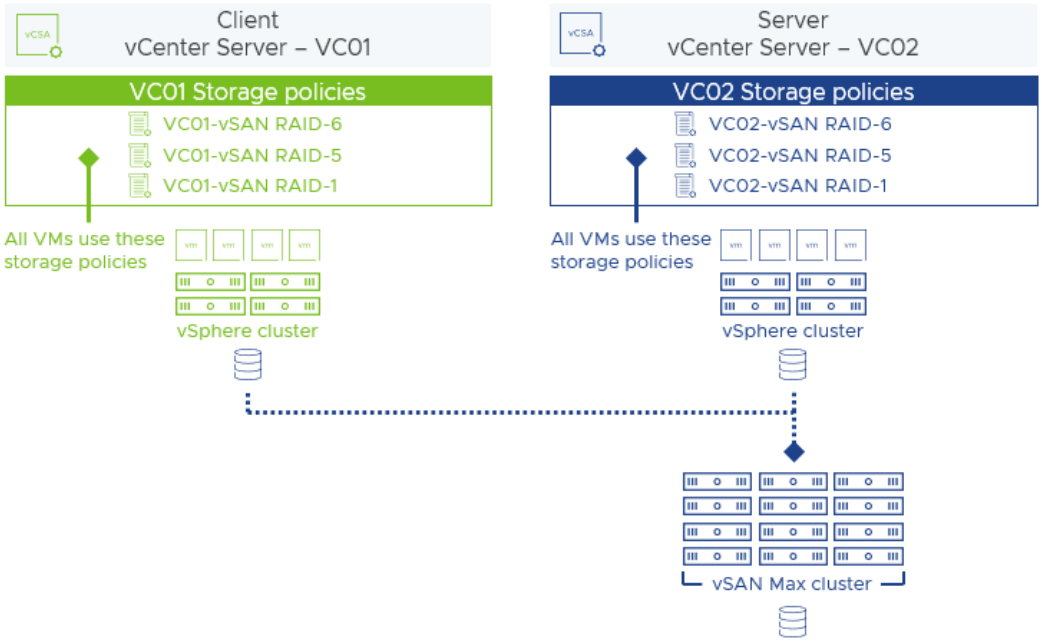

Is cross-cluster vMotion (without storage vMotion) supported with disaggregation in vSAN?

Yes. VMs can be migrated between two clusters sharing a single datastore and is fully supported with disaggregated topologies.

Does the configuration of a vSAN storage cluster require the use of DRS and HA? And what about Virtual Distributed Switches?

Even though a vSAN Max cluster does not host any user-created guest VMs, some vSphere configuration settings are necessary for proper functionality.

- Virtual Distributed Switches (vDS), which is available as a part of the vSAN Max license, should be used to simplify management.
- DRS and HA, which is available as a part of the vSAN Max license, should be enabled.
- vMotion interface should be configured to ensure mobility of management VMs.

For more information on configuration of a cluster, see the "Preparing the vSAN storage cluster for its initial configuration" section in the vSAN Storage Cluster Design and Operational Guide.

What happened to HCI Mesh?

HCI Mesh was the name for the early disaggregation capabilities of vSAN. It represented connectivity from one vSAN HCI cluster to another vSAN HCI cluster for the purpose of **cross-cluster capacity sharing**. As noted above**, this type of configuration still exists**, and is now referred to as "**vSAN HCI with Datastore Sharing**." vSAN storage clusters (previously known as “vSAN Max”) is a new first-class citizen in vCenter Server, giving the unique ability to provide a "storage-only" cluster to serve resources to vSphere clusters, and even vSAN HCI clusters. These two offerings were not well represented by the name of HCI Mesh. The term is no longer used in recent releases of vSAN.

What is the difference between vSAN storage clusters and the old HCI Mesh?

HCI Mesh represented connectivity from one vSAN HCI cluster to another vSAN HCI cluster for the purpose of cross-cluster datastore sharing. This type of capability still exists between vSAN HCI clusters but is not referred to in the UI as HCI Mesh anymore. It is now referred to as "vSAN HCI with Datastore Sharing." This option will also allow compute clusters to mount the datastore of a vSAN cluster, but for a centralized shared storage solution, vSAN Max will be the best way to achieve supreme levels of scalability and flexibility.

vSAN storage clusters is our fully disaggregated deployment option where a vSAN storage cluster provides the storage services only to vSphere clusters, and even vSAN HCI clusters (running ESA). vSAN storage clusters are integrated in the UI in a way that gives customers the ability to easily configure and manage the vSAN Max cluster and provide easy connectivity and management to vSphere clusters. vSAN storage clusters are also tuned specifically for serving storage I/O.

Can vSAN HCI with Datastore sharing be used with 2-Node clusters?

Yes. As noted above, a 2-Node cluster can be used as a client cluster that mounts the datastore of another vSAN HCI cluster, or a vSAN Max cluster. A 2-Node, vSAN HCI cluster may also be used as a server cluster. While it is possible to use this type of cluster to serve resources to other clusters, this type of cluster will have a reduced ability to tolerate against multiple host failures when compared to a vSAN HCI cluster with a larger host count.

With the introduction of vSAN storage clusters, is an aggregated vSAN HCI approach no longer preferable?

The introduction of vSAN storage clusters is about choice and flexibility to tailor systems to the specific needs of your organization. We believe aggregated vSAN HCI clusters and disaggregated vSAN storage clusters can provide a powerful combination for your enterprise needs.

I want to make vSAN storage clusters as fast as possible. How do I do this?

Use higher performance hardware to achieve faster storage performance. With vSAN, performance of a VM is derived from the host hardware and the network used to interconnect the hosts in the vSAN cluster, not the cluster host count. While increasing the host count of a cluster will increase the aggregate IOPS and bandwidth achieved by the cluster, in most cases it will not improve the discrete performance capabilities observed by the VM. VM performance will be a function of the host hardware and network connectivity. Paying attention to your network topology will be critical in delivering the highest, most consistent performance possible. Ensure that your spine-leaf topology is not oversubscribed.

Stretched Clusters and 2-Node Clusters

What is a 2-Node or 2-Host vSAN cluster and how does it work?

A vSAN 2-node cluster is a special type of vSAN cluster that consists of 2 hosts storing the data, and one host in the form of a virtual witness host appliance that provides quorum voting capabilities to determine availability and prevent split-brain scenarios.

Why do vSAN stretched clusters and 2-Node clusters need a third location for a witness?

In a stretched cluster topology, the third location that holds a witness host appliance will determine quorum for the data in the stretched cluster. This prevents failure scenarios that could result in a split-brain configuration: The same VM running and updating data in two locations simultaneously. For 2-node clusters, the same principles apply, where the third entity is simply a witness host appliance at a remote location (typically the primary datacenter) and determines quorum for the two nodes in the cluster.

What are the hardware requirements for running a vSAN stretched cluster or a vSAN 2-Node cluster?

vSAN 2-node configurations have two physical hosts, with a third witness host in a third location. A stretched cluster can have up to 40 physical hosts (20 at each site) with one witness host in a third location.

Can a witness host be shared across multiple deployments?

A witness host cannot be shared in a stretched cluster. 2-Node clusters can share a witness host, supporting up to 64 2-Node clusters per witness host. See the post "New Design and Operation Considerations for vSAN 2-Node Topologies" for a better understanding on the considerations of sharing a witness host across multiple 2-Node clusters.

Can the witness host appliance be deployed in the Cloud?

The witness host appliance is packaged as an OVA. This virtual witness host appliance must run on a physical, licensed ESXi host. More than one witness host appliance can run on a physical ESXi host. Any cloud provider issuing a licensed ESXi host for use can be used for hosting the virtual witness host appliance.

Does the ESXi host version powering the virtual witness host appliance need to be the same version as the appliance?

The vSAN Witness Host is contributing to the vSAN cluster, therefore it is recommended to be the same build as the vSAN Data Nodes. It is generally required to be the same release as vSAN.

A vSAN Witness Appliance is provided with each release of vSAN. The underlying vSphere version is the same as the version running vSAN. Upon initial deployment of the vSAN Witness Appliance, it is required to be the same as the version of vSAN.

What are my options for redundancy in a stretched cluster configuration?

In a stretched cluster configuration, data can be mirrored across sites for redundancy. Additionally, a secondary level of resilience can be assigned to data that resides within each site. Assuming enough hosts in each site to do so, the secondary level of resilience will provide additional resilience in the event of host outages. Assigning site-level protection and a secondary level of protection is all achieved in a single storage policy.

Using the ESA in a stretched cluster environment has some compelling advantages. For more information, see the post: Using the vSAN ESA in a Stretched Cluster Topology.

Can stretched clusters maintain data availability if there is a failure of one site and the witness host appliance?

If they are simultaneous failures, no, as this would not meet the levels sufficient for quorum. In vSAN 7 U3 and later, data in one site can fail or be taken down for maintenance, followed by a subsequent outage of the witness host appliance, and the data will remain available. This enhancement also applies to 2-node clusters as well.

Can 2-Node clusters provide a secondary level of resilience?

Yes, secondary levels of resilience exist for 2-Node clusters that run both ESA, as well as OSA. This can help maintain availability of data in the event of a host failure, and discrete storage device failures on the remaining host. Due to the hardware requirements, it is much more efficient and economical to use the secondary level of resilience on 2-Node clusters running ESA.

Can I use “vCenter HA” with vSAN stretched clusters?

While vCenter HA can be used with vSAN Stretched Cluster using DRS affinity/anti-affinity to pin VMs to sites a cost/benefit analysis must be done in order to understand if it will work in your environment and operational model.

VCHA requires a maximum of 5ms latency between all three VMs, the primary, secondary, and witness - as such, placing the VCHA witness on the same site as the vSAN SC witness host would mean that site can be a maximum of 5ms away, not 200ms as is supported by a vSAN stretched cluster. If your vSAN witness sites are more than 5ms away an option is to co-locate the VCHA witness with either the primary or the secondary VCHA VMs, however, this also means that if the co-located site fails, the vCenter Server will be offline.

A **vSAN stretched cluster is natively integrated with vSphere HA**, offering automated failover and startup for all VMs on it - including vCenter. The benefit of using VCHA over a single vCenter VM with vSphere HA to automate the restart is in the order of a minute or two in startup time - as such, it should be considered that if the extra startup time is acceptable, use vSphere HA and a single vCenter VM for operational simplicity.

If not, use VCHA with your chosen topology: either the VCHA witness on the same site as the vSAN Witness provided it is <5ms away or co-located with the primary or secondary VCHA VMs taking into consideration the failure scenarios that go with such a co-located topology.

Is vSAN File Services supported on a stretched cluster and 2-Node clusters?

Stretched clusters and 2-Node clusters are supported topologies as of vSAN 7.0 U2 using the OSA. File services are supported in the vSAN ESA as of vSAN 8 U2.

Does disaggregation work with vSAN stretched cluster and 2-Node clusters?

Yes, disaggregation is compatible with vSAN stretched clusters and 2-node clusters when using vSAN 8 U1 (OSA), and vSAN 8 U2 in the ESA in a vSAN HCI with Datastore sharing deployment. vSAN storage clusters support a stretched cluster topology, but at this time is limited to just protecting the data across sites, not the VM instances. See the post: "Flexible Topologies with vSAN Max." for more information.

Are there recommendations for vSAN stretched cluster network connectivity?

The vSAN Stretched Cluster guide and the vSAN Stretched Cluster Bandwidth Sizing guide contains more information and recommendations specific to stretched clusters networking. While the Inter-site link (ISL) can be a significant contributor to the effective performance of a vSAN stretched cluster, it isn't the only factor. See the post "Performance with vSAN Stretched Clusters" and "Using the vSAN ESA in a Stretched Cluster Topology" for more information.

Can a standard single site vSAN cluster be converted to a vSAN stretched cluster?

Yes, it is easy to convert a standard (non-stretched) vSAN cluster to a stretched cluster. This is performed in the "Fault Domains & Stretched Cluster" section of the vSAN UI. More details can be found in the vSAN Stretched Cluster Guide.

Miscellaneous

Where can I find technical blog posts related to vSAN?

As of January 2025, all vSAN-related technical blog posts can be found at: https://blogs.vmware.com/cloud-foundation/technical/storage/

Where can I find technical white papers and design guides related to vSAN?

As of January 2025, all vSAN-related technical white papers, design guides, operational guides, etc. can be found at: https://www.vmware.com/resources/resource-center

Networking

What are the networking requirements for running vSAN?

The networking requirements for vSAN will depend on if the cluster is running the vSAN Original Storage Architecture (OSA), or the new Express Storage Architecture (ESA) introduced in vSAN 8. Hosts participating in a vSAN cluster must be connected to the network using at least one network interface card (NIC). Multiple NICs are recommended for redundancy.

The technical minimum bandwidth needed is 10Gbps. Given the current performance capabilities of vSAN ESA, 10Gb is the accepted minimum for only the entry-level ReadyNode. A realistic minimum is 25Gbps. For more information, see "vSAN ESA ReadyNode Hardware Guidance."

Are there recommendations for vSAN network connectivity?

The vSAN Network Design Guide contains more information and recommendations for network connectivity in a vSAN cluster. The "Design and Operational Guidance for vSAN Storage Clusters" document will also have network guidance that relates to both vSAN HCI clusters, as well as vSAN storage clusters.

Does vSAN support RDMA?

vSAN supports the RoCE v2 implementation of RDMA, using qualified hardware and configurations. Other implementations of RDMA such as iWarp and InfiniBand are not supported. For more information, see the blog post: vSAN 7 Update 2 RDMA Support.

To ensure a consistent, predictable, and supported experience, strict adherence to approved RDMA adapter cards on the HCL is required. vSAN clusters using RDMA may be subject to additional limitations of supported features or functionality, including, but not limited to:

- vSAN cluster sizes are limited to 32 hosts
- vSAN cluster must not be running the vSAN iSCSI services

- vSAN cluster must not be running in a stretched cluster configuration
- vSAN cluster must be using RDMA over Layer 2. RDMA over Layer 3 is not supported.
- vSAN cluster running vSAN over RDMA is not supported with disaggregated vSAN topologies, such as vSAN Max.
- vSAN cluster must not be using a teaming policy based on IP Hash or any active/active connection where sessions are balanced across two or more uplinks.

Can multiple VMkernel ports tagged for vSAN be used to improve resilience against a vSAN network fabric failure in an air-gapped network?

Yes, this type of "air-gapped" network configuration is supported, but is, and adds some unnecessary complexity. Failover times of an active/standby arrangement using a single vSAN VMkernel port is sufficient in most cases.

Can I configure multiple VMkernel ports tagged for vSAN to help improve performance?

The use of multiple VMkernel ports tagged for vSAN is not an officially supported configuration at this time.

Does NIC teaming improve performance in vSAN?

NIC teaming aims to achieve potentially two benefits. Improved failover/resilience and improved performance. NIC teaming using LACP can offer marginal performance improvements, but is complex, and not supported at this time in VCF. NIC teaming using active/active with Load Based Teaming (LBT) was not intended for the deterministic requirements of storage traffic and is not recommended. The most reliable and robust configuration at this time is using an active/standby arrangement using “route based on originating virtual port ID.” This will provide a more deterministic path for storage I/O.

Do faster network switches and interface cards improve vSAN performance?

In general, a faster networking fabric will improve performance to a vSAN cluster when the existing network fabric is a significant contributor to contention. This is especially true for clusters running the vSAN Express Storage Architecture (ESA) introduced in vSAN 8 and enhanced in vSAN 8 U1 and U2. As storage devices become faster, this can shift the primary point of contention to the network. 25/100Gb networking is quickly becoming commonplace as a result.

Does vSAN require storage fabric host bus adapters (HBAs)?

No, vSAN uses standard network interface cards (NICs) found in nearly every x86 server platform. There is no need to provision and implement specialized storage networking hardware to use vSAN.

Can I run vSAN traffic through a network overlay, firewall, IDS, or NSX?

While front-end VM traffic can run through network overlays, we highly recommend that all VMkernel traffic (including vSAN traffic) has as simple of a path as possible. Firewalls and IDS/IPS systems can inadvertently block this mission critical storage I/O in a manner that could cause substantial impacts on the performance or availability of data.

Can vSAN support direct (switchless) connection of hosts with clusters greater than two hosts?

No. This type of switchless connection is only supported in a 2-node topology. Attempting to do this with more than two hosts would create networking loops, and violate general Layer 1/2 networking principles that apply to all environments, not just vSAN. Without a switch, there would be no spanning tree protocol (STP) to prevent network loops, and a simple node failure could easily break quorum.

Capacity

How much capacity will I need in my vSAN cluster?

The amount of capacity you need depends on the amount of capacity you plan to use for the workloads that will reside on the vSAN datastore. As noted in the post: “Demystifying Capacity Reporting in vSAN” vSAN presentas all storage capacity available in the cluster in a raw form. The cluster capacity advertised as available does not reflect the capacity available for data residing in a resilient manner. The amount of additional capacity needed depends on the type of storage policy applied, and whether or not you are using vSAN ESA or OSA. The vSAN ReadyNode sizer will help account for these factors.

How much free capacity should I maintain in a vSAN cluster?

vSAN requires additional space for operations such as host maintenance mode data evacuation, component rebuilds, rebalancing operations, and VM snapshots. Activities such as rebuilds and rebalancing can temporarily consume additional raw capacity. Host maintenance mode temporarily reduces the total amount of raw capacity a cluster has. This is because the local drives on a host that is in maintenance mode do not contribute to vSAN datastore capacity until the host exits maintenance mode.

Prior to vSAN 7, we recommended to maintain approximately 30% free capacity for vSAN clusters. In vSAN 7 and later, changes were made to the data structure that generally allowed for a lower percentage of free capacity. The recommended amount depends on many factors, including the host count of the cluster, the type of cluster, and the hardware specifications of the hosts. All of these variables will be factored in for you when using the vSAN ReadyNode sizer which will give you a precise amount needed for free capacity. The "Reserve Capacity" feature in vSAN will use these same calculations to help you manage this free space more effectively. When in doubt, maintaining 20-30% free capacity remains as a good operational practice for vSAN.

Should I enable "Host Rebuild Reserve" and "Operation Reserve" toggles in all of my vSAN clusters?

No. The Host Rebuild Reserve (HRR) and Operations Reserve (OR) toggles are optional capacity management toggles as a part of the "Reserved Capacity" capability. While it can simplify capacity management, in some cases it may not either be supported, or advised. The "Reserved Capacity" capability currently does not support clusters in stretched cluster configuraitons, 2-Node deployments, or clusters using vSAN's "Fault Domains" feature. For more information, see the post: "Understanding Reserved Capacity Concepts in vSAN."

There may be other cluster configurations where the use of the Reserve Capacity feature (particuluarly Host Rebuild Reserve) may not be ideal. Examples may be clusters with 6 or fewer hosts when paired with vSAN ESA's Auto-Policy Management feature. Using the Auto-Policy Management feature with HRR will not be able to support the use of FTT=2 using RAID-6. **For vSAN ESA clusters with 6 or fewer hosts, it is recommended to use Auto-Policy Management over the HRR feature.** For more information, see the post: "Auto-Policy Management Capabilities with the ESA in vSAN 8 U1."

How can I add storage capacity to a vSAN cluster?

Storage capacity can be increased in a few ways. Hosts containing local storage devices can be added to a vSAN cluster (scale out), or additional storage devices can be added to existing hosts (scale up).

- For vSAN clusters running the Original Storage Architecture (OSA), disk groups must be configured for the new hosts after the hosts are added to the cluster. The additional capacity is available for use after configuration of the disk groups.
- For vSAN clusters running the Express Storage Architecture (ESA), storage devices can be easily claimed and added to the storage pool that resides on each host (one and only one storage pool on each host). Storage devices added to the storage pool will all contribute to capacity resources of the single vSAN datastore presented by the cluster.

This scale-out approach is most common and adds compute capacity to the cluster. More storage devices can be added to existing hosts assuming there is room in the server's chassis to add these devices. Unlike traditional storage solutions, vSAN enables a "just-in-time" provisioning model. Storage and compute capacity can be quickly provisioned as needed.

vSAN supports the TRIM/UNMAP space reclamation options. How can this be monitored?

vCenter UI can be used to monitor IOPS and throughput generated through TRIM/UNMAP commands. This is available under [Monitor] > [vSAN] > [Performance] > Backend. Alternatively, this can also be monitored through vsantop

Space Efficiency

Does vSAN support TRIM/UNMAP space reclamation techniques?

Guest Operating systems use commands known as TRIM/UNMAP for the respective ATA and SCSI protocols, to reclaim space that is no longer in use. This helps the guest operating systems be more efficient with storage space usage. Recent versions of vSAN have full awareness of TRIM/UNMAP commands sent from the guest OS and can reclaim the previously allocated storage as free space. This is an opportunistic space efficiency feature that can deliver better storage capacity utilization in vSAN environments. For vSAN clusters running file services, vSAN will automatically perform UNMAP to reclaim space temporarily consumed by vSAN file services. TRIM/UNMAP can be enabled from the CLI, or, beginning in vSAN 8, enabled in vSphere Client. For more information, see the "TRIM/UNMAP Space Reclamation" section in the vSAN Space Efficiency Technologies document.

Can space efficiency services such as deduplication and/or compression be enabled on an existing vSAN cluster?

Yes. In OSA, one can enable deduplication and compression on an existing cluster. This does require a rolling reformat of the devices that provide the storage capacity, and can be a resource intensive endeavor. In vSAN ESA, compression is a capability provided by storage policy. Enabled by default, it can be disabled, but it will not decompress existing data. With ESA, it is recommended to leave compression on.

Can deduplication and compression in vSAN OSA impact storage performance?

Yes. The architecture of vSAN OSA performs this deduplication and compression at the time that it is destaging to its capacity tier. This can effectively slow down the destaging process, which may impact guest VM performance. For more information on DD&C using vSAN OSA, see the post: "vSAN Design Considerations – Deduplication and Compression."

With vSAN ESA, **there is no material impact on performance when using compression.**

Is deduplication available in vSAN ESA?

As of vSAN 8 U3 (VCF 5.2), deduplication is not available in ESA. Deduplication is expected in ESA as a part of VCF 9.

Operations

What is the primary user interface (UI) used to configure and monitor vSAN?

The vSphere Client is used to perform nearly all configuration and monitoring tasks. For automation and repeatable tasks at scale, PowerCLI is a good option. Additional insight can be achieved through the use of VCF Operations, which provided enhanced analytics from the data it collects from vCenter Server.

How do I monitor the health of a vSAN cluster?

vSAN features a comprehensive health service appropriately called Skyline Health for vSAN that actively tests and monitors many items such as hardware compatibility, verification of storage device controllers, controller queue depth, and environmental checks for all-flash and hybrid vSAN configurations. In vSAN 8 U1, we introduced a new Skyline Health Scoring, Diagnostics and Remediation dashboard that helps solve the challenge of prioritizing identified issues in a vSAN cluster. See the post "Skyline Health Scoring, Diagnostics and Remediation in vSAN 8 U1" for more information.

What is the Skyline Health cluster scoring dashboard, and how does it work?

Introduced in vSAN 8 U1 (OSA & ESA), the new cluster health status and troubleshooting dashboard helps answer the basic questions of, "Is my cluster and the workloads it serves in a health state?" and if not, how sever is the condition? …and should the issue be resolved? It provides that information that traditional health alerts alone cannot. For each vSAN cluster, the Skyline Health for vSAN will provide a quick at-a-glance score of the condition of a cluster so an administrator can easily determine if all is good, and if not, how impactful any identified issues are. The new mechanism used a sophisticated method of weighing triggered health checks, and aligns them with common pillars of responsibility, such as data availability, performance, capacity utilization, and efficiency and compliance. It will then provide the most important, impactful triggered health findings in an order of priority so that an administrator can resolve issues quickly and easily. See the post "Skyline Health Scoring, Diagnostics and Remediation in vSAN 8 U1" for more information.

What vSphere maintenance mode should I use in vSAN?

When a host that is part of a vSAN cluster is put into maintenance mode, the administrator is given three options concerning the data (vSAN components) on the local storage devices of that host. The option selected has a bearing on a couple of factors: The level of availability maintained for the objects with components on the host and the amount of time it will take for the host to enter maintenance mode. The options are:

- Ensure accessibility (default)
- Full data migration
- No data migration

Details on how data is handled are provided in the vSAN documentation. In summary, the default option, “Ensure accessibility," is used when the host will be offline for a shorter period of time. For example, during maintenance such as a firmware upgrade or adding memory to a host. “Full data migration” is typically appropriate for longer periods (hours or days) of planned downtime or the host is being permanently removed from the cluster. "No data migration" commonly allows the host to enter maintenance mode in the shortest amount of time. However, any objects with an FTT=0 with components on the host going into maintenance mode are inaccessible until the host is back online.

How would I know what VMs and objects would be impacted when a host enters maintenance mode?

Before moving a host into maintenance mode, an administrator can use the Data Migration Pre-Check feature to assess the impact of the maintenance mode option. This will help you determine the impact on Object Compliance and Accessibility, Cluster capacity, and Predicted Health.

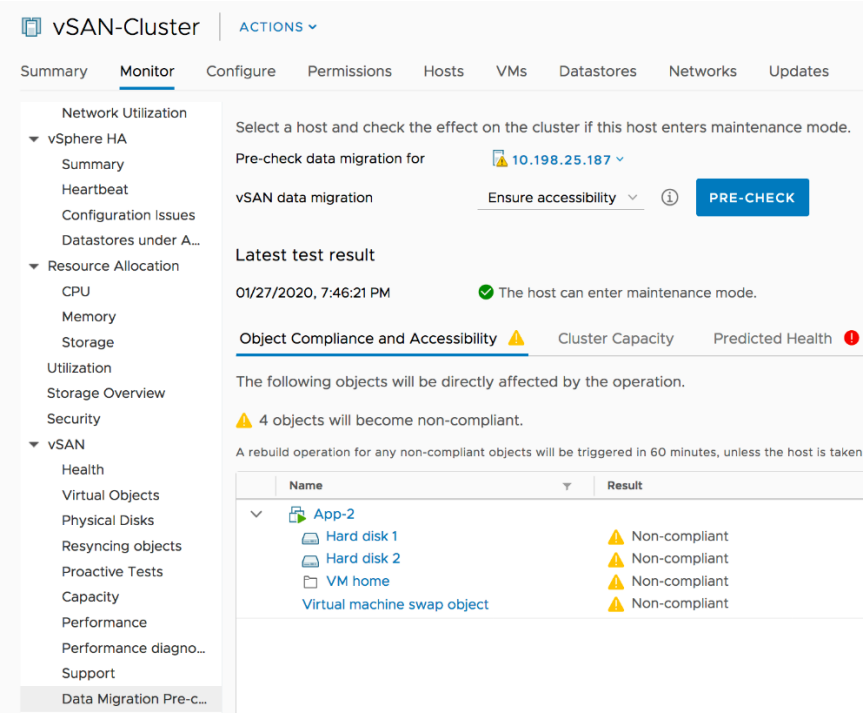

Can vSAN upload information about my environment to help improve a support case opened?

Yes, vSAN allows customers to upload anonymous information about their environments to VMware, which provides several benefits including:

Time spent on gathering data is reduced when a support request (SR) is opened with VMware Global Services (GS - previously known as "GSS"). A GS Technical Support Engineer (TSE) can utilize vSAN Support Insight <link> to view current and historic data about a customer’s environment and start troubleshooting efforts sooner, which leads to faster resolution times. vSAN

online health checks identify issues specific to a customer's environment and suggest resolutions. These online health checks can also make recommendation changes that adhere to VMware "best practices."

VMware receives anonymous data from a large number of environments that can be utilized to identify trends, potential software bugs, and better understand how products are used. Bug fixes can potentially be developed faster, and improvements are implemented to provide a better overall customer experience.

In vSAN 8 and newer, one can benefit from some of the offerings through enrolling in the CEIP without ticking the checkbox. vSAN 8 introduces "Proactive Insights" that will provide improved health awareness for environments with a vCenter Server with connectivity to the internet, but not enrolled in CEIP.

Can isolated environments use the built-in Skyline health features found in vCenter Server?

Yes. Introduced in vSAN 7 U2, the Skyline Health Diagnostics (SHD) tool allows administrators of isolated environments to manage an environment in a way that is like a fully cloud connected environment. The SHD tool can gather the latest signature libraries at a time and frequency that is best for the customer and run periodically in an environment to detect the latest alerts and updates provided by VMware. It will also allow VMware Global Support Services to resolve issues in isolated environments more effectively, as it can assist with the data gathering process.

Does vSAN work with VMware vSphere Lifecycle Manager (vLCM)?

Yes, vSAN is fully integrated with vLCM.

In stretched cluster and 2-Node environments, should I back up a vSAN virtual witness host appliance?

No. Backups, restores, clones, and snapshots of a Witness Host are not supported. Use the 'Change witness host' function in the vSphere Client to deploy a new Witness Host when there is an issue with the existing Witness Host.

How can I gracefully power down a vSAN cluster?

vSAN 7 U3 introduced an automated workflow that will guide the user through the process of gracefully shutting down a cluster. It includes several prechecks and other guidance to ensure that the shutdown and power-up process is simple and predictable. vSAN 8 enhanced the workflow to ensure there are no circular dependencies, and in vSAN 8 U1, and PowerCLI 13.1, new cmdlets are available to perform this task via PowerCLI. See the post "Automation improvements using PowerCLI with vSAN 8 U1" for more information.

Performance

What is the "Number of Disk Stripes per Object" rule in a vSAN storage policy?

Setting this rule to a number other than the default of 1 instructs vSAN to stripe one side of a RAID tree for an object across multiple capacity devices. For example, setting this rule to 4 instructs vSAN to stripe an object with this policy assigned across four devices. (Since the construct of an erasure code is different, the behavior of stripes is different as well. See the blog post: "Stripe width improvements in vSAN 7 U1" for more information.

This storage policy rule can be beneficial in certain circumstances where the capacity devices are consistently the constraining factor in the data path, when using the vSAN Original Storage Architecture (OSA). In most cases, it is best to leave the striping rule at its default setting of 1 for hybrid and all-flash vSAN configurations.

For the vSAN ESA, this storage policy rule can be ignored. For more information, see the post: Stripe Width Storage Policy Rule in the vSAN ESA.

What is the recommended way to test vSAN performance?

VMware provides a tool called HCIBench. It is essentially an automation wrapper around popular and proven synthetic test utilities. With HCIBench, you can either invoke Vdbench or Flexible I/O tester (FIO) to automate performance assessment in an HCI cluster.

HCIBench simplifies and accelerates proof-of-concept (POC) performance testing in a consistent and controlled manner. The tool fully automates the process of deploying test VMs, coordinating workload runs, aggregating test results, and collecting

data for troubleshooting purposes. The output from HCIBench can be analyzed by the Performance Diagnostics feature. See this VMware Knowledge Base article for more information: vSAN Performance Diagnostics (2148770)

HCIBench can be used to evaluate the performance of vSAN and other HCI storage solutions in a vSphere environment.

Use HCIBench to run performance tests across the cluster instead of running a workload from a single VM. HCIBench can be configured to deploy and distribute multiple VMs across the hosts in an HCI cluster to provide more realistic and accurate test results. It is also recommended that HCI bench be run on a cluster prior to introducing it into production, then save the results for reference, if needed. This can help identify if there are any issues prior to introducing the cluster into production and can serve as a nice reference point if there are any issues later, and one wants to compare current test results to test results after initial deployment.

How does vSAN minimize the impact of data resync operations when a device or host fails?

There are several mechanisms in place to dynamically balance and prioritize virtual machine and resync I/O. It is important to maintain adequate performance while providing resources for resync operations to restore resilience.

When there is contention for I/O in the hardware storage stack (e.g. Disk group found in the Original Storage Architecture, or OSA), vSAN guarantees approximately 20% of the bandwidth to resync operations while virtual machines utilize the remaining 80%. If there is no contention for bandwidth, resync operations can consume more bandwidth to reduce resync times. Virtual machines can use 100% of the bandwidth when there are no resync operations occurring. For more information, see "Adaptive Resync in vSAN."

When using the Express Storage Architecture (ESA), VMware introduces adaptive network traffic shaping for resynchronizations. This is due to the tremendous efficiency of the vSAN ESA, and its ability to push higher rates of data through the network stack. See the post: "Adaptive Network Traffic Shaping with the vSAN Express Storage Architecture" for more information.

Does vSAN require manual intervention to balance data across the cluster?

No. Users can configure proactive rebalancing as an automated action at the cluster level to let the cluster balance the data out as it sees fit. When toggled on, this will replace the previous behavior of proactive rebalancing, which included tripping an alert in the health service, followed by the need for manual intervention by the administrator to perform the rebalance.

What is the best way to troubleshoot performance issues in vSAN?

The "Troubleshooting vSAN Performance" guide will provide a framework for proper isolation of performance issues for more accurate diagnostics of the performance issue. vSAN contains several built-in tools to help gather data in the troubleshooting performance effort, including the vSAN performance metrics, I/O Insight, vsantop, and in vSAN 7 U3, the VM I/O Trip Analyzer.

How do I get more detailed performance metrics for vSAN?

vSAN 8 U1 introduces the ability to view time-based performance metrics that use a 30 second sampling interval. This is a 10x improvement over the 5-minute sampling interval used in previous versions. This will produce performance graphs that are more representative of system behavior and will help to identify when performance issues occur. This is available in both the OSA and ESA of vSAN 8 U1. See the post: "High Resolution Performance Monitor in vSAN 8 U1" for more information.

Security

Is encryption supported with vSAN?

Yes, vSAN supports Data-At-Rest Encryption and Data-in-Transit Encryption and uses an AES 256 cipher. Data is encrypted after all other processing, such as deduplication, is performed. Data at rest encryption protects data on storage devices, in case a device is removed from the cluster. The vSAN Original Storage Architecture (OSA) provides these through two discrete cluster toggles, while the vSAN Express Storage Architecture provides both capabilities with a single toggle at the time of cluster configuration.

Does vSAN encryption require special hardware?

No, vSAN Encryption does not require any specialized hardware such as Self-encrypting drives (SEDs). Some drives on the vSAN Compatibility Guide may have SED capabilities, but the use of those SED capabilities are not supported.

Should vSAN encryption be enabled when first creating a cluster, or after workloads have been migrated?

If you are interested in using vSAN encryption, it is most efficient to enable this during the initial configuration of a vSAN cluster, prior to any VM migration. The OSA and ESA in vSAN supports both options.

What are the prerequisites to enable vSAN Data-at-Rest Encryption?

vSAN data-at-rest and data-in-transit encryption require vSAN Enterprise licensing for the desired cluster that you wish to run these services. A Key Management Server (KMS) is required to enable and use vSAN encryption. Nearly all KMIP-compliant KMS vendors are compatible, with specific testing completed for vendors. vSAN 7 U2 introduced support for the vSphere Native Key Provider. This gives basic key provider functionality built directly into vSphere. The Native Key Provider does not provide KMIP services for things external to vSphere.

For the vSAN Original Storage Architecture (OSA), turning on encryption is a simple matter of clicking a checkbox. Encryption can be enabled when vSAN is enabled or after, with or without virtual machines (VMs) residing on the datastore. For the vSAN Express Storage Architecture (ESA), the decision to enable encryption must be made at the time of creating the vSAN cluster.

Always use Trusted Platform Modules (TPMs) to cryptographically store issued keys on host, for better resilience if the KMS is offline.

How does vSAN Encryption differ from vSphere VM Encryption?

vSAN encryption operates at the storage level and encrypts the vSAN datastore. VM Encryption operates on a per-VM basis and performs encryption on in-flight I/O i.e., encrypts IO as it is generated from the VM.

VMs encrypted with vSphere VM encryption can be deployed to a vSAN datastore just like other datastore types such as VMFS and NFS. However, vSAN space efficiency features such as deduplication and compression will provide little to no benefit with these encrypted VMs. **vSAN encryption services are preferred over vSphere encryption capabilities when VMs are powered by vSAN.**

How is performance impacted when using vSAN encryption services?

A detailed explanation on this matter can be found on the blog post: "Performance when using vSAN Encryption Services." The level of impact will depend also if a cluster is running the Original Storage Architecture (OSA) in vSAN, or the Express Storage Architecture (ESA). **The ESA is inherently more efficient at resource utilization for these types of data services, and will have a substantially less impact on resources when compared to the OSA.**

Does encryption in the vSAN ESA perform better than in the OSA?

Yes. Due to the architectural design of the ESA, it uses fewer CPU cycles and I/Os to encrypt a given amount of data. This translates to improved performance, and reduced utilization.

Does enabling encryption consume any additional capacity overhead?

No. There will not be any additional capacity overhead used when vSAN encryption is enabled on a cluster. This applies to both the OSA and ESA. For a better understanding of capacity overheads in vSAN ESA, see the post: "Capacity Overheads for the ESA in vSAN 8."

Does vSAN encrypt object data with different keys?

In more recent versions of the ESA, yes. While vSAN ESA uses a cluster-wide Disk Encryption Key (DEK), the object data has object-specific keys that are used. **This ensures that object data in VM uses different keys than any other VM.** Note however

that a rekey operation will perform a rekey across all data on the cluster, and does not give a choice for rekeying discrete objects.

Should I deploy a Key Management Service (KMS) server on the vSAN datastore that will use the same KMS for key management?

Deploying a third-party KMS is discouraged. When a vSAN host with encryption enabled is restarted, it requests a new Host Key and Key Encryption Key (KEK) from the KMS. If the KMS is not online to provide these keys, the host will not be able to read the encrypted data. This creates a circular dependency resulting in no access to encrypted data. For a more robust environment during failure conditions of a KMS server or the vSphere Native Key Provider, VMware highly recommends the use of Trusted Platform Modules (TPM) on each server. This will allow for the keys distributed to the hosts to be securely stored on persistent media (the TPM) and ensure keys can be used under these types of failure conditions.

What is vSAN Data-in-Transit encryption?

Data-in-transit encryption is a cluster-wide feature introduced with vSAN 7 Update 1. When using the vSAN Original Storage Architecture (OSA), administrators can choose to enable it independently or along with Data-at-rest encryption. The vSAN Express Storage Architecture (ESA) will encrypt data in transit any time that the cluster was configured for encryption.

Data-in-transit securely encrypts vSAN data traffic that traverses across hosts using FIPS 140-2 Cryptographic modules.

Does vSAN Data-in-Transit encryption require a KMS?

No, Data-in-transit encryption does not require a Key Management Server (KMS). With DiT encryption, the vSAN hosts are responsible for creating the symmetric keys. This is a process that happens transparently, without any administrative effort.

What happens when a vCenter server managing a vSAN datastore with encryption enabled is offline?

There is no impact to the virtual machines running on the vSAN datastore with encryption enabled. After vSAN Encryption is configured, vSAN hosts communicate directly with the Key Management Server (KMS) cluster. If the original vCenter Server cannot be recovered, a new vCenter Server should be deployed as soon as possible. For a more robust environment during failure conditions of a KMS server or the vSphere Native Key Provider, VMware recommends the use of Trusted Platform Modules (TPM) on each server. This will allow for the keys distributed to the hosts to be securely stored on persistent media (the TPM) and ensure keys can be used under these types of failure conditions.

What is the impact to the VMs running on a vSAN datastore with encryption enabled if the KMS is offline?

Key Encryption Key(KEK) is cached on the ESXi hosts' memory on booting. Hence there is no impact to the virtual machines till the hosts remain powered-on. If the hosts are restarted, the encrypted disk groups are unmounted and cannot be mounted until the KMS is restored. For a more robust environment during failure conditions of a KMS server or the vSphere Native Key Provider, VMware recommends the use of Trusted Platform Modules (TPM) on each server. This will allow for the keys distributed to the hosts to be securely stored on persistent media (the TPM) and ensure keys can be used under these types of failure conditions.

Doe items such as backup and recovery work with vSAN encryption services?

Yes, as vSAN encryption was designed to maintain compatibility with other vSAN and vSphere features, as well as, 3rd-party products including data protection solutions. Data is encrypted or decrypted just above the physical storage device layer. APIs such as vSphere Storage APIs for Data Protection (VADP) and vSphere APIs for IO Filtering (VAIO) that are used for data protection and other solutions are located higher in the storage stack. **Data at this layer is not yet encrypted**. Therefore, compatibility with these solutions is maintained when vSAN encryption is enabled, and VMs can be protected, migrated, or replicated without issue.

Is two-factor authentication supported in vSAN?

2-factor authentication methods, such as RSA SecurID® and Common Access Card (CAC), are supported with vSAN, vSphere, and vCenter Server.

Is vSAN part of a DISA STIG?

Yes, VMware vSAN is part of the VMware vSphere STIG Framework. The DISA STIG defines secure installation requirements for deploying vSAN on DoD networks. VMware worked closely with DISA to include vSAN in the existing vSphere STIG.

Has vSAN achieved FIPS certification?

In 2017, the VMware VMkernel Cryptographic Module achieved FIPS 140-2 validation under the National Institute of Standards and Technology (NIST) Cryptographic Module Validation Program (CMVP). In 2024, the validation for FIPS 140-3 was achieved.

vSAN consumes the VMware VMkernel Cryptographic Module when providing data-at-rest encryption thanks to the tight integration between vSAN and the ESXi kernel. When vSAN Encryption is enabled, the vSAN datastore is encrypted with FIPS Approved AES-256 utilizing the validated VMware VMkernel Cryptographic Module. This delivers FIPS compliance without the need for costly self-encrypting drives (SEDs).

How can storage devices used in a vSAN cluster be safely decommissioned, removing any residual data?

vSAN has ways to securely wipe storage flash devices decommissioned from a vSAN cluster. This is done through a set of PowerCLI commands (or API), providing an efficient and secure methodology to erase data in accordance with NIST standards. See the post: "vSAN – A Secure Fortress for your Data" for more information.

Does vSAN support the use of the vSphere Native Key Provider (NKP)?

Yes. vSAN supports the use of the vSphere NKP for key management. The vSphere Native Key Provider is a simple way for vSphere to create and manage encryption keys without using a traditional Key Management Server (KMS). It can be ideal for customers who have simple security requirements, have not already enabled encryption for their vSphere clusters, and want to do so in a secure and supported manner.

Can I still use my existing KMS for key management of a vSAN environment?

Yes! Full featured KMS solutions offer features and capabilities that are beyond the intention of the vSphere Native Key Provider. Depending on the environmental and customer security requirements, an external KMS may be the only way to achieve the specific customer security requirements.

Does vSAN support TLS?

Yes. See the link: Enable or Disable TLS Versions on ESXi Hosts for more information on how to use a specific version supported.

Should I use the vSphere NKP instead of a full-featured KMS solution?

It depends on the customer requirements, but in many cases, the specific customer requirements related to key management extend beyond the basic capabilities of the vSphere Native Key Provider (NKP). The NKP is ideal for customers who have simple security requirements and need basic key management for vSphere and/or vSAN only. The NKP is not a full featured KMS and does not support many of the capabilities found in a KMS.

How much bandwidth does a KMS introduce into an environment?

The key exchange process is relatively lightweight, introducing approximately 100KB/s between the KMS, the hosts and the managing vCenter Server. For reliable key distribution, persistent availability of keys is much more important than the bandwidth required. To ensure keys are readily available, VMware recommends the use of Trusted Platform Modules (TPM) installed on each host in a vSAN cluster. This will cryptographically store issued keys on the host should there be any connectivity issues with the KMS.

vSAN Data Protection

What is vSAN Data Protection?

vSAN Data Protection is a capability that allows customers to easily protect and recover virtual machines on a local vSAN datastore. It uses the highly efficient, high performance snapshotting capabilities of the Express Storage Architecture (ESA) paired with a new intuitive User Interface to make local protection and restoration of VM's easy.

What is required to use vSAN Data Protection?

vSAN Data Protection is a part of a VCF licensing entitlement, and does not require any extra licensing, and is a part of vSAN 8 U3 and VCF 5.2. At this time, only clusters configured in an aggregated vSAN HCI deployment are supported. vSAN Max clusters will not be able to take advantage of vSAN Data Protection, but still supports high performance snapshots taken through the UI, or using 3rd party backup applications that use VADP.

What would some typical examples of how vSAN Data Protection could be used?

vSAN Data Protection can be an ideal solution to augment your existing data protection strategies. The most common use cases might include:

- **Reverting existing VMs.** This helps address situations such as accidental VM misconfigurations, unsuccessful VM OS upgrades, or suspected malicious activity.
- **Restore existing VMs no longer registered in vCenter Server.** VMs that are accidentally deleted or moved can be easily restored to the existing cluster.
- **Clone VMs.** Cloning VMs can be an ideal way to support basic test and development workflows for all different types of teams, including Development, and IT staff.

Can vSAN Data Protection replicate these snapshots to a remote location?

As of vSAN 8 U3 (VCF 5.2), protection capabilities are limited to local protection only. The ability to replicate snapshots remotely has been announced with VCF 9.

Does vSAN Data Protection protect against ESXi hosts compromises?

No. vSAN Data Protection aims to provide an additional level of protection and flexibility to your VM data protection strategy. Concerns of host security should be addressed with the vSphere Security Hardening Guide.

What is a protection group?

vSAN Data Protection uses the concept of Protection Groups, which allows administrators to achieve two objectives.

1. Group multiple VMs for easy and repeatable snapshot creation and management
2. Define and execute a policy of outcomes, such as the frequency of protection and retention schedules.

Can VMs participate in more than one protection group?

Yes. A VM can be a part of up to three protection groups

Can protection groups consist of multiple schedules?

Yes. A protection group can have as many as 10 schedules.

How can VMs be associated with a protection group?

One or more VMs can be individually added to a protection group. Or they may be added by dynamic naming assignment, where "*" and "?" characters can be used to include all VMs that match a naming pattern.

Are the snapshots of VMs in a protection group taken at precisely the same time?

No. At this time, while a protection group will strive to take snapshots of the VM's that are a member of that protection group in accordance to the defined schedule, it will be a best effort operation, meaning that the snapshots may not be taken at precisely the same time.

What is snapshot immutability, and why does it exist?

In the context of vSAN Data Protection, snapshot immutability refers to the inability of the snapshot to be modified or deleted. Snapshots can be made immutable by an optional toggle within the settings of a protection group. They exist to preserve the integrity of the snapshot against malicious activities, and can serve as a basic way of recovering VMs during a ransomware attack.

Why not make all protection groups immutable?

While the immutability setting is extremely helpful for some use cases, it is not ideal for all situations. For example, when a protection group has the immutability setting enabled, one cannot edit or delete the protection group, change the VM membership, or edit/delete the snapshot.

How many snapshots can be crated for a VM?

vSAN Data Protection supports up to 200 snapshots per VM when snapshots are created using its UI and APIs. Note that if one uses the traditional UI or APIs, this will still be limited to 32 snapshots per VM.

What happens when VMs reach the 200 snapshot limit?

The oldest snapshots will automatically expire when the maximum number of snapshots for a VM are hit.

How does the system protect against capacity management issues when allowing for so many snapshots?

vSAN Data Protection will automatically pause snapshots if 70% of the cluster capacity is reached.

Once a VM is cloned from an existing snapshot, can it be protected using vSAN Data Protection?

No. When new VMs are created from a snapshot, these are a linked clone, not a fully autonomous, independent clone. Therefore, snapshots from the linked clone cannot be taken.

Are there any disadvantages to having a system perform a lot of snapshots?

The snapshotting engine in vSAN ESA is extremely efficient, and generally has very little impact on the performance of a VM. However, there can be cases where the frequency and number of snapshots may have an impact on the cluster. Snapshot frequency may dramatically **increase the rate of capacity usage if your data change rate is high.** As the capacity utilization of a cluster increases, it may need to perform more garbage collection processes, which may **temporarily impact performance** especially as available capacity becomes scarce.

It is recommended to choose modest levels of snapshot frequencies and retention rates so that change rate behavior can be better understood in your specific environment.

I see vSAN Data Protection uses a virtual appliance. Won't this be a single point of failure for snapshots?

No. The virtual appliance, which runs PhotonOS powering a few containers, is used to help orchestrate snapshot activities and interact with vCenter Server. Metadata about the snapshot created by the snapshot service will be persisted to disk. This allows for the retention of information about the snapshots even if the snapshot appliance is unavailable, or corrupted.

Why is this called "vSAN Data Protection" if I've always been told that snapshots are not backups?

vSAN Data Protection in vSAN 8 U3 helps you protect and recover data to a previous state using data stored locally and can augment your existing backup and recovery strategy by offering all new levels of convenience and flexibility.

The confusion in the question stems from terms such as "backup" and "protection" that have been overly generalized by the industry and can have multiple meanings depending on the context in which they are used. When paired with extremely simplified statements such as "Snapshots are not backups" this leads to more misunderstanding, as the statement lacks the level of detail necessary to be accurate in all cases.

The most accepted data protection strategy generally involves 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. Organizations generally achieve this through VADP-based backup applications triggering hypervisor snapshots so that it can capture the state of a VM at a specific point in time and copy it to a target location that lives on different media and/or different locations, which achieves the objectives of a 3-2-1 rule.

vSAN Data Protection uses the same snapshotting engine in ESA that is used by VADP-based backup vendors. vSAN Data Protection can help augment protection and recovery strategies by creating, storing, and managing locally housed snapshots. While it alone does not achieve the objectives of a 3-2-1 rule, it can provide convenient and flexible data recovery scenarios such as reverting a VM to a previous state, restoring a VM that has been deleted from inventory, or cloning a VM for other operational workflows.

Are vSAN snapshots crash consistent?

Yes, since vSAN Data Protection is integrated in the I/O path of vSAN ESA, the snapshot data is automatically committed in a crash consistent manner, without the need to stun the VM. For more details, see the FAQs in the "Availability" section of the vSAN FAQs.

Can VMware Live Site Recovery, or VLSR (previously known as Site Recovery Manager) be used with vSAN Data Protection?

VMware Live Site Recovery (VLSR), which is a part of the VMware Live Recovery (VLR), does not currently integrate or use vSAN Data Protection. VMware Live Cyber Recovery (VLCR), which is the other part of VLR, does integrate with vSAN Data Protection.