



vSAN Technology Overview

An overview of technology and
capabilities for vSAN in VMware Cloud
Foundation 9.1

May 5, 2026

Table of Contents

Introduction.....	4
Scope of Topics	4
History	4
Next Generation Storage with the vSAN Express Storage Architecture	4
Use of Commodity Hardware	5
Where vSAN can be Used	5
Distributed Architecture	5
Object Based Storage	5
Space-Efficient Erasure Coding	6
Novel Log-Structured File System for Supreme Performance	6
Integration	6
I/O Processing	6
Adaptive Throttling	7
Reporting/Visibility	7
VMware Cloud Foundation	7
Deployment Options and Topologies	7
Aggregated vs Disaggregated	8
Cluster Topology Options	8
Capabilities in Data Resilience and Availability	8
Flexible Resilience and other Capabilities using Storage Policies	8
Minimal failure domains	8
Intelligent Data Placement	9
Automated Recovery	9
Enhanced Durability during Failures and Maintenance Activities	9
Robust Failure Handling over a Variety of Conditions	9
Capabilities in Space Efficiency	9
Thin Provisioning	9
Data Compression	9
Global Deduplication	10
Space Reclamation	10
Capabilities in Data Protection and Recovery	10
vSAN Snapshots	10
vSAN Protection and Recovery	10

Remote Protection, Disaster Recovery and Cyber Recovery	10
VADP	10
Capabilities in Data Security	11
vSAN Data-at-Rest Encryption	11
vSAN Data-in-Transit Encryption	11
Secure Disk Wipe	11
Capabilities in Data Services	11
vSAN File Service	11
vSAN Native S3 Compatible Object Storage	12
vSAN iSCSI Target Services	12
Persistent Volumes for Cloud Native Services	12
Summary	12
Additional Resources	12
About the Author	13

Introduction

Since its introduction in 2014, VMware vSAN has evolved into the premier storage platform for VMware Cloud Foundation (VCF). Its extraordinary capabilities and integration into the hypervisor offer capabilities and simplicity for your private cloud infrastructure. Customers have spent over two decades trusting their hypervisor to manage the compute and memory resources that power their mission-critical workloads, so it makes sense that the hypervisor is in a unique and superior position to manage storage.

Scope of Topics

This document is intended to serve as a general overview of the technologies that are a part of vSAN. When applicable, embedded links to additional content will provide more detailed explanations. The information provided in this document will assume the use of vSAN 9.1 as a part of VMware Cloud Foundation (VCF) 9.1. VCF deployments may have additional requirements and support limitations that fall outside of the scope of this document. A long list of [vSAN FAQs](#) are available that can serve as an excellent resource for information on vSAN.

History

VMware introduced vSAN at a time in the industry in which customers were looking for something different. At that time, virtualization was a ubiquitous element in datacenters of any size. As commodity servers became more powerful, more workloads were being virtualized, which put new strains on traditional approaches to centralized shared storage. The inherent limitations of the architecture with traditional storage arrays using dedicated storage fabrics made it challenging to meet the ever-changing capacity and performance requirements. There was an opportunity for an intelligent, integrated approach to virtualized storage. Enter vSAN HCI.

VMware vSAN removed many of those challenges by aggregating storage into the same hosts that comprise the vSphere cluster. Unlike other HCI-based offerings, vSAN was built into the hypervisor. This offered performance, efficiency, and integration in ways that could not be achieved with other solutions. But it was built at a time when some hardware technology was trying to catch up. Flash was available, but it was expensive with limited capacity and used legacy protocols like SATA and SAS that prevented widespread use. vSAN was developed to help account for the hardware of the day using a 2-tier approach using a combination of flash and spinning disk, or all-flash to provide storage. This approach was wildly successful, with adoption by over 40,000 customers.

Next Generation Storage with the vSAN Express Storage Architecture

As the industry moved to high-performance NVMe storage devices, high-speed networking, and computationally dense servers, vSAN's architecture needed to be optimized to exploit the fully performance capabilities of this new hardware. The original storage architecture in vSAN solved many of the challenges inherent to distributed storage, but there was an opportunity to redesign its data path and underlying data structure to ensure that vSAN ESA could deliver high performance storage in an efficient way.

vSAN ESA provides a better way to process and store data. Its unique log-structured file system ingests data as quickly as possible while preparing the data to be stored in a space efficient and scalable manner. The result has been stunning for most customers. It is not uncommon to see 2-5x better performance than compared to vSAN OSA using the same hardware. We've even seen many customers report that [vSAN ESA beats the performance of their top performing storage arrays](#). Paired with this impressive performance is its dramatic improvement in processing and space efficiency that allows for fewer servers to run more workloads. Some customers have seen reduction in server counts by 20-40% while processing and storing the same amount of data, which results in much [lower total cost of ownership \(TCO\) when compared to vSAN OSA or other storage arrays](#). For more information, see the post: "[An Introduction to the vSAN Express Storage Architecture](#)." The minimum requirements for vSAN ESA have been lowered dramatically since its debut – and as recently as November, 2025. For more information, see the post: "[Driving Down Storage Costs with Lower Hardware Requirements for vSAN](#)."

vSAN ESA also brings operational simplicity. [vSAN in VCF 9.1 introduces Auto-RAID](#), which is an evolutionary step forward from Auto-Policy Management. Auto-RAID automatically delivers the optimal level of resilience for all of your data, based on

the characteristics of the cluster, and is fully managed by vSAN. Pair this with [Improved storage device claiming for vSAN hosts](#), simplified maintenance of storage devices, and [fewer settings to adjust for optimal performance](#), and vSAN is easier than ever to operate.

But vSAN ESA does more than drive better performance with improved efficiency. It delivers new capabilities that were simply not possible using the OSA. Scalable snapshots, vSAN Protection and Recovery, [multi-source replication](#), and vSAN storage clusters are just a few examples of powerful new capabilities in vSAN ESA that will be discussed in more detail, below.

Use of Commodity Hardware

vSAN has always delivered on the promise of running on readily available commodity hardware. This means that customers get to take advantage of the latest hardware on the market without waiting for a storage solution to incorporate these same technologies into a proprietary solution. Your data center can evolve with the industry, instead of being limited by the roadmap of a vendor providing a traditional storage solution.

At the heart of vSAN support of hardware is the vSAN ReadyNode program. These are servers from your favorite server OEMs that include hardware componentry that has been certified by the server OEM to work with vSAN. ReadyNodes come in a [variety of profiles](#) for vSAN HCI clusters, vSAN storage clusters and [vSAN Cyber Recovery clusters](#). These profiles represent different performance and capacity expectations and are preconfigured with a proportional set of hardware resources (storage devices, networking, CPU, and RAM) to ensure systems have a balanced set of resources. As of November 2025, vSAN ReadyNodes profiles certified for both vSAN HCI clusters and vSAN storage clusters have been changed, with much lower hardware minimums. See the post: "[Driving Down Storage Costs with Lower Hardware Requirements for vSAN](#)" for more details.

vSAN ReadyNodes can be purchased as a ReadyNode in the form of a single SKU, or from the same vendor as individual hardware components. This is known as a [ReadyNode emulated configuration](#), and offers another option in hardware procurement. Existing vSphere hosts [may also be eligible for retrofitting](#) so that it can be used in a vSAN cluster – an intriguing option for those faced with current market conditions.

vSAN HCI clusters and vSphere hosts mounting the datastore of a vSAN storage cluster can even work with [vSphere Advanced Memory tiering](#), which will allow you to run more workloads with less RAM – a great strategy to lower your server hardware costs.

Where vSAN can be Used

vSAN ESA is suitable for even the most demanding workloads. Whether it be your general-purpose workloads, or your mission critical ERP system, vSAN can deliver all your storage requirements using the software you already know. Its distributed architecture allows it to account for the wide variety of topologies in your environment. Whether it be isolated clusters, site-resilient stretched clusters, small 2-Node clusters for edge environments, or disaggregated topologies, vSAN lets you design your environment based on the needs of the organization instead of the constraints of a storage solution. vSAN is even the go-to architecture for cyber resilience requirements, where you can [use vSAN as the target clean room for cyber recovery scenarios, or disaster recovery requirements](#).

Distributed Architecture

Object Based Storage

Compared to traditional storage, vSAN uses a novel way to store data. Instead of storage data using a clustered file system like VMFS, it uses an approach that is most analogous to an object store. Instead of spreading a clustered file system across a series of hosts, it uses a much smaller boundary of data such as an individual VMDK, and stores this as an object. The administrator defines the desired level of resilience of the data, and vSAN automatically stores that object data to ensure it meets the requirements. This type of distributed object store provided much better availability, simplified scalability, and a granular level of management that wouldn't be possible with a traditional clustered filesystem.

For more information, see the post: "[vSAN Objects and Components Revisited](#)."

Space-Efficient Erasure Coding

Erasure coding is a common method of storing data in a resilient and space-efficient way. While it can offer a guaranteed level of space efficiency compared to that of data mirroring, it typically requires more effort and hardware resources to achieve and can impact performance. vSAN ESA eliminates any compromise on performance with erasure coding.

Erasure coding in vSAN ESA is built in. It uses data placement schemes that allow for different resilience levels using erasure codes that fit within a reasonable number of hosts. For example, a RAID-6 erasure code will spread the object data with parity across 6 hosts and will withstand up to two of those 6 hosts being offline while maintaining full availability of the data. This resilience can be achieved with just a 1.5x capacity overhead, compared to a 3x overhead for data mirroring. For clusters with fewer hosts, vSAN ESA can use RAID-5 erasure coding to store object data with parity across as few as 3 hosts. This can withstand the failure of one of the three hosts holding the object data and will remain available. It will achieve this resilience with just a 1.5x capacity overhead, compared to a 2x overhead for data mirroring.

vSAN's unique distributed storage architecture paired with space efficient erasure coding can often make it more resilient than erasure coding techniques found in traditional storage arrays. For more information, see the post: "[Erasure Codes in VMware vSAN versus Storage Arrays.](#)"

Novel Log-Structured File System for Supreme Performance

Best of all, in vSAN ESA, there is no compromise in performance with using erasure codes versus data mirroring. This is in part due to the use of a log-structured file system (LFS) with vSAN ESA. VMware has a storied history with log structured filesystems, since Mendel Rosenblum, a founder of VMware, has been credited with the proposal and implementation of the first log structured file system in the early 1990s.

vSAN uses an LFS that is uniquely designed for the characteristics of distributed storage. Unlike storage arrays that may use some variation of a single, monolithic log-structured file system in a storage array, vSAN uses a discrete LFS for every object in the datastore. This offers substantial control and tuning capabilities in vSAN, which is critical to maintaining performance at scale. Its pervasive use at the top of the vSAN storage stack is one of the ways that vSAN helps minimize I/O amplification, as it plays a role in nearly every data service, from [data-at-rest encryption](#) to [high-performance snapshots](#). An optimized I/O pathing using vSAN's LFS doesn't just improve write I/Os. [Read operations in vSAN ESA are fast and efficient regardless of how large the cluster grows.](#)

For more information on fast and efficient write operations, see the following posts: "[RAID-5/6 with the Performance of RAID-1 using the vSAN Express Storage Architecture.](#)" "[Performance Improvements with the Express Storage Architecture in vSAN 8 U1](#)" and "[Performance Improvements with the Express Storage Architecture in vSAN 8 U2.](#)"

Integration

Unlike other solutions, vSAN is a part of the hypervisor. What is the result? Much better performance, lower overheads, better I/O visibility and control, and seamless integration with the rest of the hypervisor.

I/O Processing

vSAN consists of a collection of processes that are responsible for a variety of activities. Some of its processes are running at low-level kernel space for minimal context switching and optimal efficiency. Other processes operate within a much more flexible user space. This delivers the best combination of processing I/O in a speedy and efficient manner, while maintaining high levels of flexibility. It is this integration that allows vSAN avoid the technical inefficiencies associated with controller VMs, that hairpin all storage I/O through a series of virtual machines.

This integration of I/O processing allows for supreme levels of efficiency. For example, [data compression](#), [data encryption](#), and check summing reside at an optimal location within the stack to ensure that a large amount of data processing can be performed once with the least amount of effort. This even benefits network transmission, where these optimally placed processes reduce network transmission.

Adaptive Throttling

vSAN uses its own purpose-built scheduler to manage storage I/O. This composite scheduler recognizes different vSAN I/O types, and can not only convey this information to users in the UI, but it can manage and control the I/O. For example, during times of activity where data is moving from one host to another, it will distinguish this “resynchronization” I/O as different than from the front end read and write commands from the guest to the storage system. During periods of contention, it will prioritize the front-end VM/I/O over resynchronization traffic to ensure that VM workloads can maintain sufficient levels of performance. In vSAN ESA, this adaptive throttling occurs within the storage stack of each host, along with the network used to connect all the vSAN hosts in the cluster to each other. Adaptive throttling is built into all aspects of vSAN, including [deduplication post-processing](#), and [storage resilience reconfigurations](#).

For more information, see the posts: “[Adaptive Resync in vSAN](#)” and “[Adaptive Network Traffic Shaping with the vSAN Express Storage Architecture](#).”

Reporting/Visibility

Visibility is a significant benefit to the integration of vSAN into the hypervisor.

Health

With its acute awareness of the entire stack, vSAN delivers information on the health and well-being of a vSAN cluster that is directly integrated into the product. For example, the [vSAN Health Scoring, Diagnostics and Remediation](#) capabilities help administrators quickly understand what issues exist within a vSAN cluster and their level of seriousness.

Performance

vSAN offers a distinct advantage over traditional storage solutions because it can maintain visibility and awareness of all I/O from end to end. vSAN is continually improving visibility of storage I/O through its “[Top Contributors](#)” view. Other tools built into vSAN include the vSAN VM I/O Trip Analyzer tool, as well as vSAN I/O Insight. These can be extraordinarily helpful when [troubleshooting vSAN performance](#).

It even offers high-resolution performance metrics for many of the most important storage metrics. For more information, see the post: “[High Resolution Performance Monitoring in vSAN](#).”

Capacity

vSAN provides an abundance of information on consumed and available capacity in a cluster. Since vSAN is a distributed storage system that aggregates storage devices into a single datastore for the cluster, it historically [has presented storage capacity a little differently](#) than traditional storage. In vSAN for VCF 9.1, vSAN has a new “Effective Capacity” view that renders all capacity information in a manner that is very similar to traditional storage. See the “[vSAN Space Efficiency Technologies](#)” document and the post: “[Simplifying Storage with the New Effective Capacity View in VMware vSAN for VCF 9.1](#)” for more information.

VMware Cloud Foundation

Integration with VMware Cloud Foundation (VCF) is one of the many reasons why vSAN is the premier storage solution for VCF. For example, deployment workflows within VCF give you the ability to use vSAN HCI or vSAN storage clusters as your principal storage. Performance and health metrics for vSAN can be easily collected by VCF Operations (VCF Ops) for more comprehensive analysis and trending requirements. VCF Operations for Logs can also be a great tool for interpreting complex log data sent to vCenter Server.

Deployment Options and Topologies

vSAN distributed architecture translates to flexible deployment options for customers. Based on your geographic constraints, your availability requirements and your personal preferences in resource provisioning, you can easily choose what works best for your environment. Regardless of what you choose, vSAN is managed in the same way.

Aggregated vs Disaggregated

- **Aggregated vSAN HCI Clusters.** vSAN HCI clusters is a deployment model that aggregates compute and storage resources into the same hosts that comprise the vSAN cluster. The VM instances will run on the hosts in the cluster, and unless its datastore is shared with other vSAN clusters, the storage will be treated as an exclusive resource of the cluster.
- **Disaggregated vSAN Storage Clusters.** vSAN storage clusters (previously known as vSAN Max) is a deployment model that provide centralized shared storage to vSphere clusters. Unlike vSAN HCI clusters, they do not run any VM instances on the cluster itself. vSAN storage clusters have [impressive capabilities that cannot be matched by traditional modular storage arrays](#). They allow customers to [start small and grow incrementally](#) as their needs evolve. See the post, [“vSAN HCI or Storage Clusters – Which Deployment Option is Right for You?”](#) to help sort through what may be best for you.

Cluster Topology Options

- **Standard Single-Site vSAN Cluster.** This cluster type typically refers to a vSAN HCI or vSAN storage cluster in a single geographic site. It assumes network bandwidth and latency are the same between all hosts in the cluster.
- **Standard Single-Site Cluster using vSAN Fault Domains.** This cluster type represents a single site vSAN cluster where the vSAN “fault domains” feature is used to define specific fault domains. These fault domains can be racks, or rooms. For more information, see the post, [“Using Fault Domains in vSAN ESA.”](#)
- **vSAN Stretched Clusters.** This cluster type will consist of hosts in a single vSAN cluster stretched across two geographic sites. They are connected by an inter-site link (ISL) and a third location that communicates with a witness host appliance. Stretched clusters provide site-level resilience of data – maintaining the availability of data in the event of a site failure. For more information, see: [“vSAN Stretched Cluster Guide.”](#)
- **vSAN 2-Node Clusters.** This cluster type consists of two hosts that are typically connected directly to each other, along with a virtual witness host appliance at a third location. 2-Node clusters provide simple host-based resilience for small edge and remote office locations.

Any of these deployment options and topologies can be used for your production workloads, DR sites, or cyber recovery environments.

Capabilities in Data Resilience and Availability

Availability of data is a core competency of an enterprise storage system. The unique traits of vSAN allow for flexible levels of data resilience to account for a wide variety of conditions in the data center. While much of this is covered in significant detail in the document [“vSAN Availability Technologies”](#) the following will highlight some of the key traits of data resilience and availability in vSAN.

Flexible Resilience and other Capabilities using Storage Policies

vSAN is unique in its ability to prescribe storage characteristics on a per-VM basis. This is achieved using storage policies. The storage policies applied to VMs will ensure that the VMs adhere to the prescribed policy, regardless of the conditions of the cluster. This “outcome oriented” approach to administration makes administration easy and consistent. [vSAN for VCF 9.1 introduces Auto-RAID](#), the successor to [Auto-Policy Management](#) that automatically manages the optimal level of resilience for all data stored in a vSAN cluster, based on the characteristics of the cluster. Storage policies remain for elements such as workload throttling and capacity reservations, but now lets the system manage the optimal data resilience settings.

Minimal failure domains

vSAN ESA aggregates storage devices across hosts in a cluster and presents it as a single datastore. In the event of a discrete failure of a storage device, the boundary of failure will only be that specific storage device. This minimizes the amount of data that needs to be repaired during a device failure, and allows vSAN to self-heal faster, regaining prescribed

levels of resilience in minimal time. For more information, see the post: “[The Impact of a Storage Device Failure in vSAN ESA versus OSA](#)” and the document: “[vSAN Availability Technologies](#).”

Intelligent Data Placement

While vSAN uses [distributed object store](#) as a model for storing data, it must place shards of this data across hosts in a manner that makes it resilient to the level prescribed by the associated storage policy.

Automated Recovery

Intelligent data placement also means automated recovery if vSAN recognizes that the data isn't in full compliance with the prescribed storage policy. If a host in a cluster is unavailable for a period of time, or fails, vSAN will recognize the absent data, and heal itself so that the data regains its prescribed level of resilience.

Enhanced Durability during Failures and Maintenance Activities

In times of planned host maintenance, or unplanned host outages, vSAN will create what are known as “durability components.” These components will include an additional copy of all updated data from the time of the host outage. This helps ensure that if subsequent failures occur, the most recently updated data will be present in more than one location.

Robust Failure Handling over a Variety of Conditions

Data resilience is built throughout vSAN's storage stack. Here are a few other capabilities that automatically make your data more resilient and durable.

- **DDH.** [Degraded Device Handling](#) (DDH) is a monitoring technique within vSAN that will detect high device latency thresholds and patterns to these thresholds that suggest imminent failure of storage device in the capacity tier. This symptom-based approach attempts to detect impending failures of storage devices.
- **Low-Level Metadata Resilience.** vSAN incorporates low-level metadata redundancy to assist with Unrecoverable Read Errors (UREs) which are a common and expected occurrence during the life of a storage device. This low-level metadata is mirrored on devices to prevent the need to recover unreadable data from other devices, improving the performance and efficiency of vSAN under these types of failure conditions.

More detail on these features, as well as other capabilities including NVMe device endurance tracking and proactive hardware management can be found at: “[vSAN Availability Technologies](#).”

Capabilities in Space Efficiency

vSAN uses multiple techniques to store data in a space efficient manner. These methods are discussed in detail in the document: “[vSAN Space Efficiency Technologies](#)” but a few are highlighted here.

Thin Provisioning

Thin provisioning is the act of initially provisioning a minimal amount of storage capacity while presenting a much larger amount of storage capacity for use by a VM. This makes for unclaimed space to be available for other provisioning activities and VMs, making the storage system more efficient.

Data Compression

vSAN uses data compression not only to make the data stored more space efficient, but to reduce the amount of data needed for further processing and transmission. [Data compression has been enhanced in vSAN for VCF 9.1 using a new highly tuned compression algorithm](#). Positioned high in the storage stack, the data will be compressed prior to it being encrypted and transmitted across the network. Not only does this reduce the amount of CPU processing needed, but it also reduces the amount of bandwidth needed for data transmission. The benefits of data compression is applicable to both standard vSAN clusters, as well as [stretched vSAN clusters](#).

Global Deduplication

[Global deduplication was first introduced under limited availability in vSAN for VCF 9.0, and generally available in 9.1.](#) It is a revolutionary step forward in reducing the total cost of storage. This feature provides a cluster-wide deduplication domain, meaning that it will look for duplicate blocks across the entire vSAN cluster and deduplicate the data when it finds a match. These processes occur asynchronously so that they do not interfere with the performance of the VMs. They are adaptively throttled to occur more often during quieter times and will deduplicate less frequently written data first. vSAN global deduplication will perform as well or better than deduplication found in leading storage arrays. For more information, see the posts: “[Save Costs and Scale Efficiently with vSAN Deduplication in VMware Cloud Foundation 9.0.](#)” and “[More Capacity with VMware vSAN Compression and Global Deduplication in VCF 9.1.](#)”

Space Reclamation

vSAN fully supports [space-reclamation techniques such as TRIM/UNMAP](#) to reclaim previously used capacity that is no longer used by a VM. This pairs nicely with vSAN’s thin provisioning capabilities for VMs that inflate in consumed capacity followed by deleting data inside the guest VM.

Capabilities in Data Protection and Recovery

vSAN offers a comprehensive set of technologies to help protect your VM workloads and data in an easy, flexible, and scalable manner.

vSAN Snapshots

The architecture of vSAN ESA allowed for the development of an [entirely new snapshot engine](#). ESA’s log-structured file system (LFS) allows for a metadata-based approach to snapshots. They are fast, scalable, and efficient. Crash consistent snapshots can be taken without stunning the VM, and snapshot deletions or consolidations occur 100x faster than the traditional redo log-based snapshot mechanism found in the original storage architecture (OSA).

vSAN Protection and Recovery

vSAN ESA snapshots are the foundation for a new method to automatically protect VMs in a manner that is easy, scalable, and integrated into the software you already know. vSAN Protection and Recover (previously known as “vSAN Data Protection”) provides local protection of VMs through automatic scheduling and retention policies. They can be used to revert existing VMs, clone VMs, and even recover VMs that have been deleted from vCenter Server’s inventory. The local protection capability of vSAN Protection and Recovery is a part of a VCF license. New capabilities have been introduced in vSAN for vCF 9.1, as described in the post: “[VMware vSAN Protection and Recovery Enhancements for VCF 9.1.](#)”

For more information, see the posts: “[Superior Snapshots using VMware vSAN Protection and Recovery](#)” and “[vSAN Protection and Recovery in VMware Cloud Foundation – The Solution You Already Own.](#)”

Remote Protection, Disaster Recovery and Cyber Recovery

Remote protection of VMs can be easily achieved using VMware Advanced Cyber Compliance (ACC) or VMware Site Recovery. Previously known as “VMware Live Recovery” (VLR) this VCF add-on extends the functionality of vSAN Protection and Recovery to easily protect data residing on vSAN clusters, VMFS datastores, or NFS datastores to another vSAN cluster using the very same interface as with local protection. Data is asynchronously replicated to a remote/target cluster which can be used for disaster recovery. The VMware ACC add-on license includes both DR protection and orchestration that comes with SRM, as well as the full cyber recovery capabilities to serve as a safeguard against ransomware attacks.

VADP

Third-party backup solutions can be used with vSAN courtesy of VMware APIs for Data Protection (VADP). These APIs will allow data protection vendors to integrate their backup solution with vSAN ESA snapshots. When these solutions perform a backup on one or more VMs, VADP will issue an ESA snapshot to create a point-in-time state of the VM at the time the backup is desired. The backup solution then protects the VM(s) at this point, removing the possibility that subsequent data is changed prior to the completion of the backup. When the backup is complete, VADP will remove the temporary snapshot

from the VM. Data protection vendors may also use Changed Block Tracking (CBT) that are compatible with vSAN ESA snapshots.

Capabilities in Data Security

Data security is an intrinsic part of vSAN. vSAN offers several ways to ensure that data can be properly secured to meet your organizational or regulatory requirements. These features are discussed in detail in the document: “[vSAN Encryption Services](#)” but a brief overview is provided below.

vSAN Data-at-Rest Encryption

vSAN can provide full [data-at-rest encryption capabilities](#) through a simple toggle that can be applied at the cluster level. All data written will be fully encrypted to storage devices. vSAN provides this encryption capability through VMware vSphere’s VMkernel cryptographic modules and **does not require the use of self-encrypting devices (SEDs)**. Encryption occurs high in the storage stack, after data compression so that all subsequent traversal down the storage stack remains encrypted.

vSAN’s Data-at-Rest Encryption is compatible with many KMIP compliant KMS solutions, as well as VMware vSphere’s Native Key Provider (NKP). vSAN also offers [multiple key rotation options](#) for Data-at-Rest Encryption. This can be helpful in maintenance and regulatory compliance procedures. Best of all, [vSAN Data-at-Rest Encryption allows a cluster’s datastore to be fully encrypted, without any compromise on the space efficiency gained from vSAN global deduplication](#).

vSAN Data-in-Transit Encryption

vSAN also offers Data-in-Transit (DiT) Encryption for extra levels of security of data transmitted across the network. While vSAN Data-at-Rest Encryption does encrypt the data prior to transmission across the network, its focus is on encrypting the data as it resides on the storage devices. Enabling DiT Encryption will ensure that no duplicate blocks are using the same cryptographic hash.

DiT Encryption is also a cluster-based feature and can be enabled with a single toggle. It manages its own keys between hosts in the vSAN cluster and does not need or use a KMS or the vSphere NKP. And as of vSAN in VCF 9.1, [it can even encrypt data transmitted to and from vSphere clusters mounting the datastore to a vSAN storage cluster](#).

Secure Disk Wipe

vSAN offers a simple and secure way of [erasing residual data](#) on a storage device per the standards defined by the National Institute of Standards (NIST). This is an optional feature when enabling encryption services and is especially useful in conditions where encryption is enabled for an existing vSAN cluster where the devices already store unencrypted data, or adding new devices to an existing cluster.

Capabilities in Data Services

vSAN File Service

vSAN file services is an optional cluster-based toggle that allows vSAN to serve file shares to VMs and cloud-native applications. Instead of purchasing a dedicated array providing file services capabilities (known as a “filer”), you can provide this same functionality anywhere you have a vSAN cluster. vSAN File Services supports NFS protocols for use with Linux VMs and containers, as well as SMB for use as file shares.

vSAN File services are rendered through a collection of stateless containers automatically deployed on the cluster where the capability is enabled. The containers are what present the specific protocol (NFS or SMB) and passes that data down to an underlying “Virtual Distributed File System” (VDFS). This is a highly scalable substrate that allows for file shares and exports to be distributed across the hosts efficiently. The file share data resides in data objects, that are protected at a level of resilience prescribed by storage policy assigned by an administrator.

vSAN Native S3 Compatible Object Storage

vSAN in VCF 9.1 is offering a technology preview of new S3 compatible object storage built directly into vSAN. The new native S3 compatible object storage capabilities will give vSAN the ability to provide scalable, S3 storage to all of your consumers on premises using the storage you already use for block and file needs in your VCF environment. Provisioning and management of S3 object stores and buckets is achieved using VCF Automation or APIs through vSphere Supervisor, which helps make it fully capable of multi-tenancy. Since this capability is fully integrated with vSAN, it reflects the flexibility and scalability traits of any other vSAN cluster. Perhaps best of all, native S3 compatible object storage is included in your VCF license, with no need to purchase or license additional software. For more information, see the post: "[Native S3 Compatible Object Storage in VMware vSAN for VCF 9.1.](#)"

vSAN iSCSI Target Services

vSAN offers the ability to connect legacy workloads that require block-based iSCSI volumes. It allows one to provision a volume within the guest, where the volume is housed on a vSAN cluster. It allows you to continue to operate your legacy workloads while shifting some of the storage requirements to vSAN. See the document: "[vSAN iSCSI Target Usage Guide](#)" for more information.

Persistent Volumes for Cloud Native Services

Cloud Native Applications (CNA) often need persistent volumes to retain for further processing or access by other workloads. vSAN can easily present persistent volumes for container-based applications. These volumes provide persistent storage to vanilla Kubernetes, as well as VMware vSphere Kubernetes Services (VKS). These persistent volumes are fully integrated with Kubernetes Storage Classes, which allow persistent volumes to be provisioned by a tenant using a specific level of protection and will translate to the appropriate storage policy.

Summary

Simply put, vSAN is the premier storage platform for VMware Cloud Foundation. Its comprehensive set of capabilities, flexibility, integration, and performance allows you to provide enterprise storage using the software you already know.

Additional Resources

The following are a collection of useful links that can be helpful for vSAN customers.

[vSAN Frequently Asked Questions \(FAQs\)](#). The vSAN FAQs are an assembled list of common questions received on all aspects of vSAN.

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

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

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

[vSAN Network Design Guide](#). This network design guide applies to environments running vSAN 8 and later.

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

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

[Official vSAN documentation](#). The location for all "how to" documentation on vSAN.

About the Author

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

