vSAN File Services

An overview of vSAN File Services in VMware Cloud Foundation 9.0

November 17, 2025

vmware by Broadcom

Table of Contents

Introduction

vSAN File Services provides the ability for administrators to present and share data using SMB shares to Windows systems or NFS exports to Linux systems or Cloud Native Applications (CAN) anywhere a vSAN cluster is present. It's ease of use and integration with the hypervisor eases the burden of management for vSAN environments that require file-level access. Instead of using a legacy physical storage array, or deploying VMs to provide file services, an administrator can simply enable this cluster level service on a vSAN cluster.

Its thoughtful design allows vSAN File Services to provide file services capabilities across a broad range of conditions with the ease of management using software you already know.

Scope of Topics

The information provided in this document will assume the use of vSAN 9.0 and/or VMware Cloud Foundation (VCF) 9.0. VCF deployments may have additional requirements and support limitations that fall outside of the scope of this document.

Feature Overview

vSAN File Services offers a wide variety of features and capabilities that are commonly associated with file sharing solutions. It makes it ideal for customers who want to provide file services in a location without the additional expense of a dedicated array providing similar functionality. The capabilities below will list many of the capabilities that are like other solutions, while highlighting other capabilities that are unique to vSAN File Services.

Easy to Enable and Configure

vSAN File Services is an optional cluster-based feature that can be enabled easily in the vSphere Client UI. The initial configuration will walk you through the information that it needs to configure the service for your environment. For more information on what information is needed to configure vSAN File Services, see "Enable vSAN File Service" in Techdocs. Once installed and configured, one can easily see and revise this configuration information in the UI.

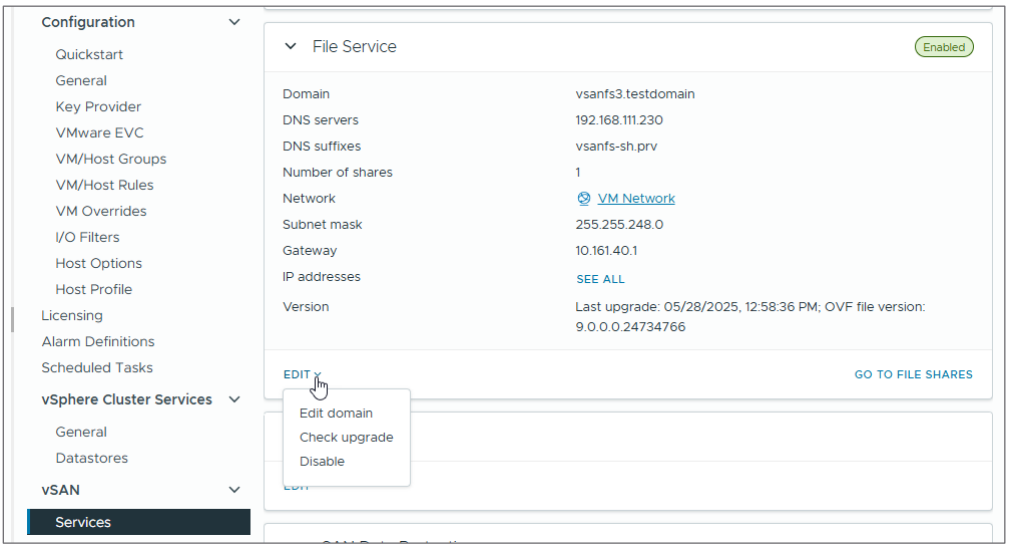

Figure. Viewing the current configuration of vSAN File Services.

vSAN File Services can use Kerberos-based authentication to provide a common, secure method of authentication for vSAN file shares. It supports multiple Kerberos flavors (KRB5, KRB5I, and KRB5P) for maximum flexibility. Kerberos can be used directly with NFS shares, or through configuration of vSAN File Services with Active Directory for SMB shares.

Flexible Presentation of Shares

vSAN File services can provide up to 500 file shares per cluster. This can provide a sufficient number of shares for a large majority of circumstances, including environments running cloud native workloads that use NFS-based persistent volumes for persistent storage.

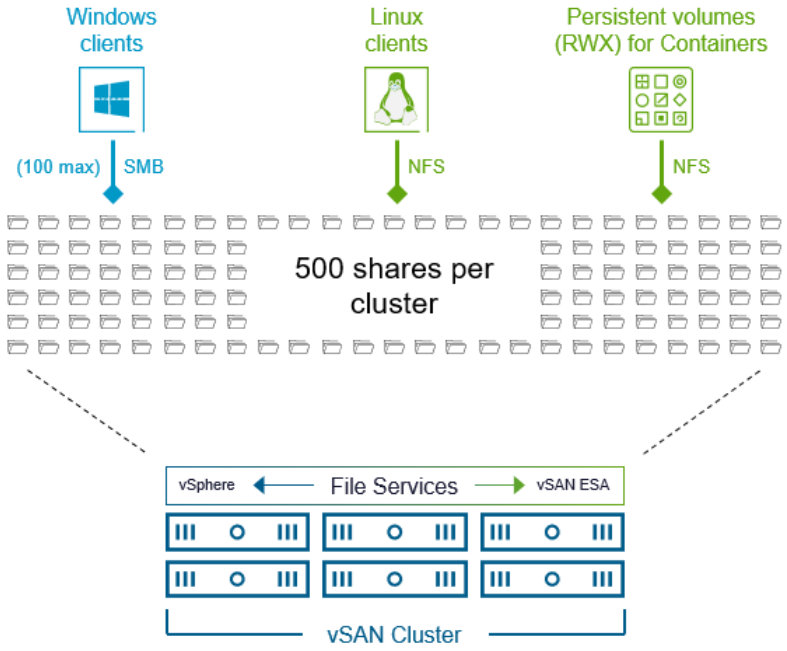

Figure. Support of up to 500 shares per vSAN cluster.

While vSAN File Servers does support up to 500 file shares per cluster, the limit on SMB shares used by Windows systems remains at 100.

Automatic Load Balancing and Scalability

vSAN uses its existing Cluster-Level Object Manager (CLOM) to distribute file share data as objects across the cluster, just like it does for VM object data. When vSAN File Services is enabled, it uses and additional "protocol services" layer to ensure that the presentation of shares is fairly distributed across the hosts within a cluster. This combination makes the storage of the file share data, and the access points for each share evenly distributed.

Support of Commonly used Protocols

vSAN File Services supports two common protocols that can be used to connect with virtual guest VMs, or legacy physical servers.

- **NFS (v3 and v4.1).** NFS shares are used in environments where Linux systems may need to share content from a common location. An NFS share will be "exported" or made available to mount from Linux machines. These could be used for user home directories, or common shared repositories for files and other unstructured data. Cloud Native Applications (CNA) may use these NFS exports in similar ways, where the ephemeral container may need a location to persistently store and retrieve data. This is commonly known as a Read-Write-Many (or RWX) persistent volume for cloud native environments.
- **SMB (v 2.**1 and v3). SMB shares are used in environments where Windows systems may need to share content from a common location. These SMB shares will be connected to Windows systems as drives, or UNC addresses. They are ideal for common directory locations of project collaborations, unstructured data, and home directories.

Guided syntax for mounting.

vSAN helps you to easily understand and use any share that is created. vSAN will present an enumerated list of shares that have been created. Highlighting any share followed by the "Copy URL" button will provide you the exact syntax needed for the exported mount point for Linux, or UNC address for Windows. This syntax can be used in Linux fstab files or Active Directory Group Policy settings or scripting for Windows systems.

This guided mounting syntax is especially helpful for Linux environments, as the version of the NFS protocol will determine the style that is used. For example, in NFS v3, the connection string will present a simple IP address of the container responsible for rendering the share, followed by the exported share name. But in NFS 4.1, it will use a single "Primary IP" defined in the initial configuration of vSAN File Services. NFS 4.1 will then manage the protocol redirection to the container responsible for the exported share name. The primary IP is used for the initial mount, followed by subsequent NFS v4 referrals.

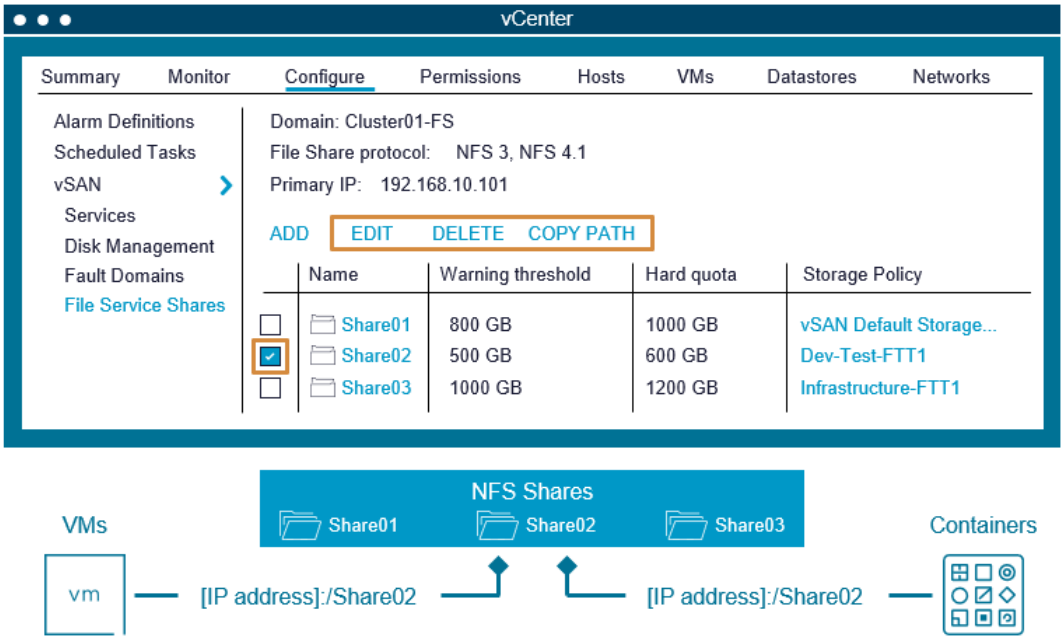

Figure. Helpful guidance in file share paths for existing file shares in vSAN File Services.

Support of Multiple Topologies and Deployment Options

One of vSAN's great capabilities is that it can be designed to accommodate multiple different topologies and deployment options. vSAN File Services benefits from this flexibility by supporting a wide range of topologies and deployment options.

Aggregated vSAN HCI Clusters and Disaggregated vSAN Storage Clusters

File Services is supported in both aggregated vSAN HCI clusters as well as disaggregated vSAN storage clusters. vSAN File Services can be configured on vSAN HCI clusters as large as 64 hosts, and vSAN storage clusters as large as 24 hosts. Although you may find whether you are running vSAN File Services or not, that sizing clusters to fit within a rack is preferable to much larger clusters that cannot fit within a rack. For more information, see the post: "vSAN Networking – Optimal Placement of Hosts in Racks."

Stretched clusters

vSAN File Services can even be deployed in a stretched cluster arrangement. An administrator can specify a parameter that determines the placement of the file share and file server. **This helps provide the proper site affinity between the clients, the share, and the backing services**. Even if the file share is protected across sites via site level protection the file server is a single entity that is responsible for the connection from the client system, and thus, one would want this to be on the same site as the file server connection to provide an optimal data path. This mechanism will maintain colocation of the client to the protocol services used, the VDFS proxy, VDFS server, and at least one of the backing vSAN objects.

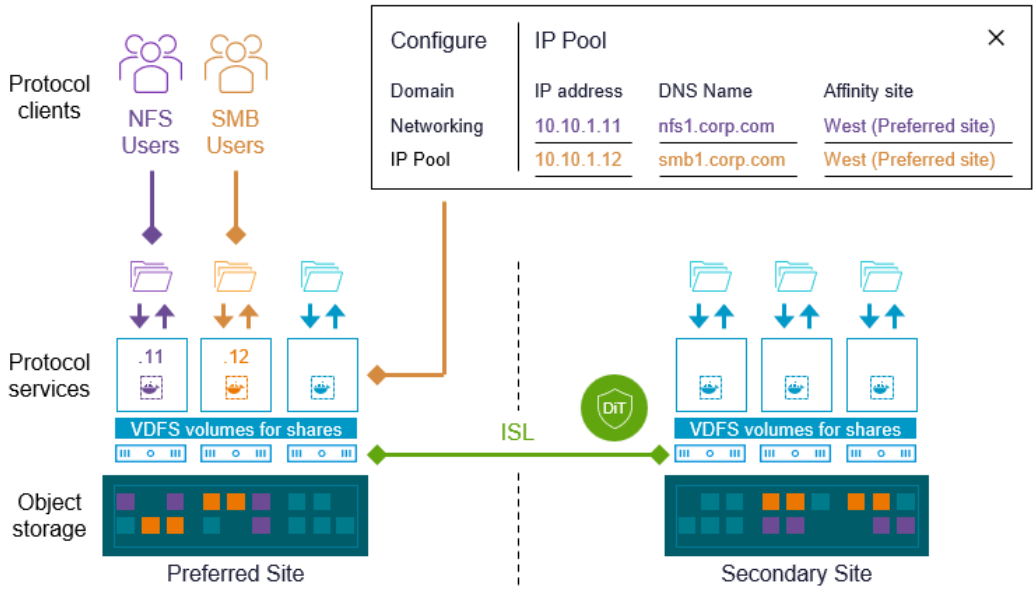

Figure. TBD

2-Node Clusters

vSAN File Services supports 2-Node clusters. This provides a powerful and easy way to serve NFS or SMB shares in edge sites or remote offices, where a traditional filer would be cost prohibitive and wasteful.

Centralized Administration

Since vSAN File Services is part of the hypervisor, it makes for easy management courtesy of vCenter Server and the vSphere client. One can adjust share names, associated storage policies, storage quotas, and network access restrictions.

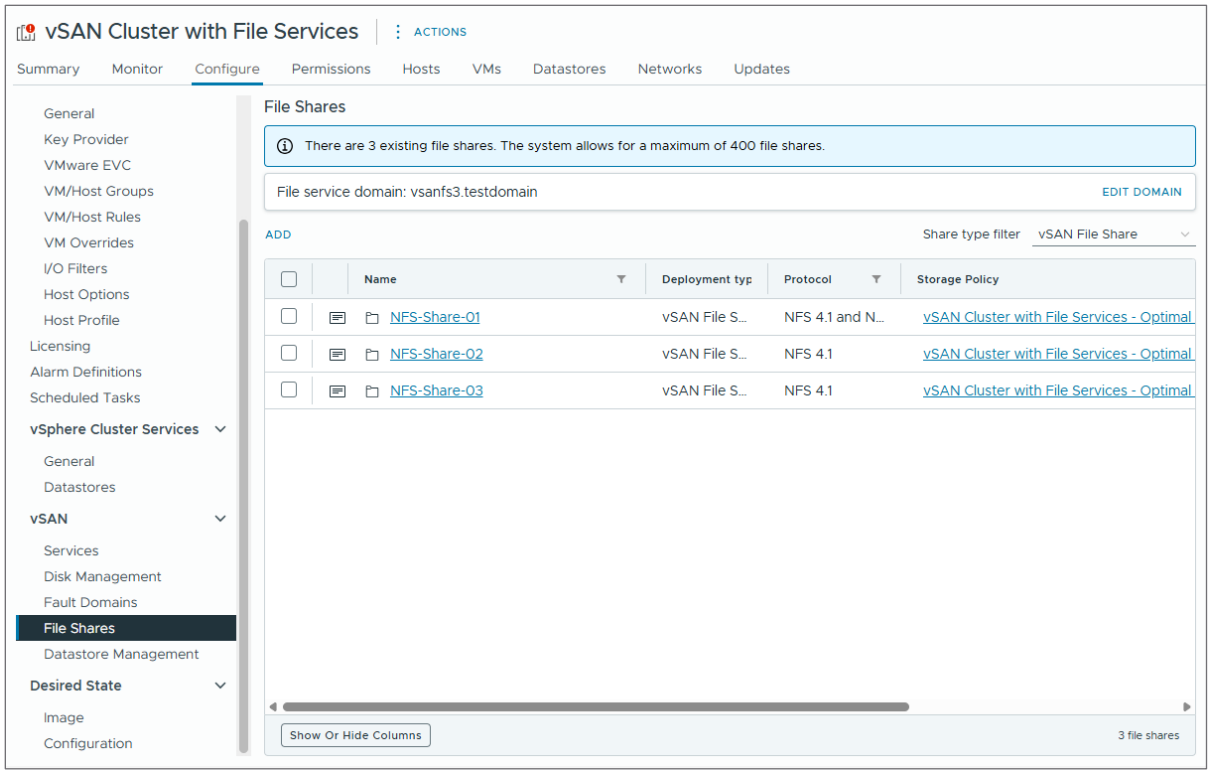

Figure. Administering share-specific settings in vCenter Server.

Integrated Health, Performance and Capacity Reporting

Since vSAN File Services are built into the hypervisor, many of the same health and performance monitoring capabilities are available with file shares as they would be with any other aspect of vSAN.

Health Monitoring

vSAN File Services uses health findings integrated into Skyline Health for vSAN. When vSAN File Services are enabled, three additional health findings are added to the long list of checks that vSAN makes to ensure the proper health of a cluster. These additional findings include:

- **File Server Health**. Check if file server is in good state. The column NFS Daemon shows if NFS daemon process is running or not.
- **Infrastructure Health.** Check file service infrastructure health state per ESXi host. The column 'vSAN File Service Node' checks if the vSAN file service node VM is in powered on.
- **Share Health.** Check if the file service share is in good state.

These health findings are proportionally weighted in the Skyline Cluster Health Score. Any triggered health checks will impact the score based on the severity of the issue, and the relative priority against other health checks. For more information, see the post: "Skyline Health Scoring, Diagnostics, and Remediation in vSAN 8 U1."

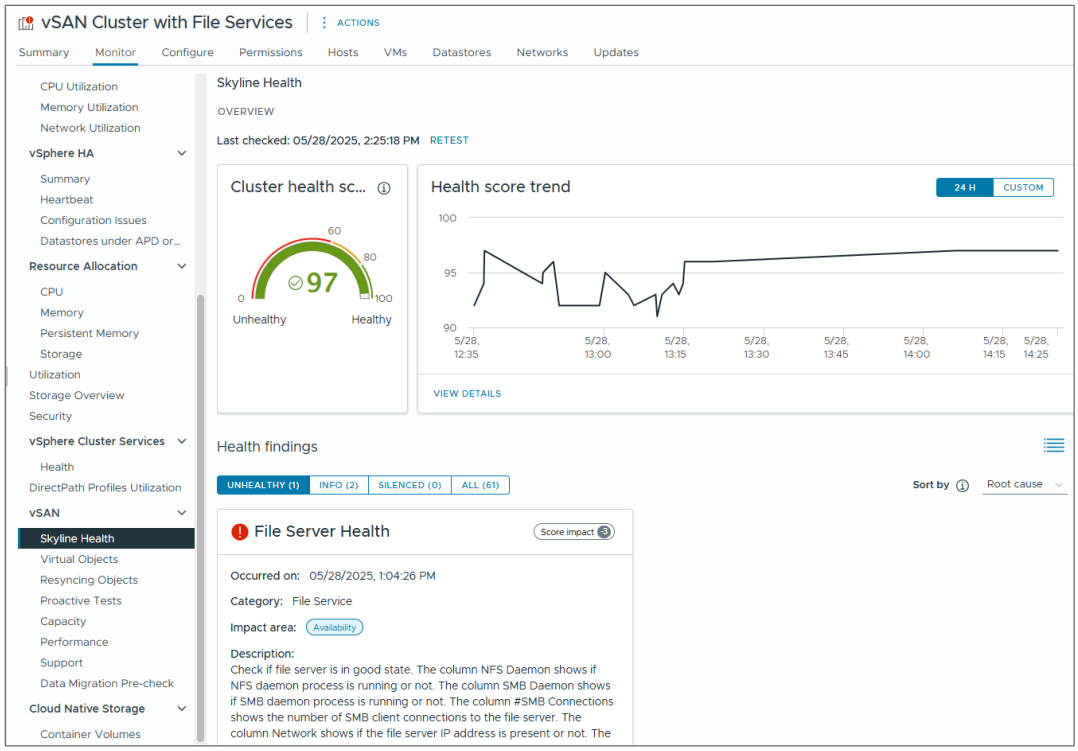

Figure. Skyline cluster health scoring and vSAN File Services health findings.

Performance Monitoring

Performance monitoring for vSAN File Services is very similar to monitoring other aspects of vSAN performance. For example, one can view time-based performance metrics on a per share basis, monitoring common metrics such as IOPS, throughput, and latency for reads and writes.

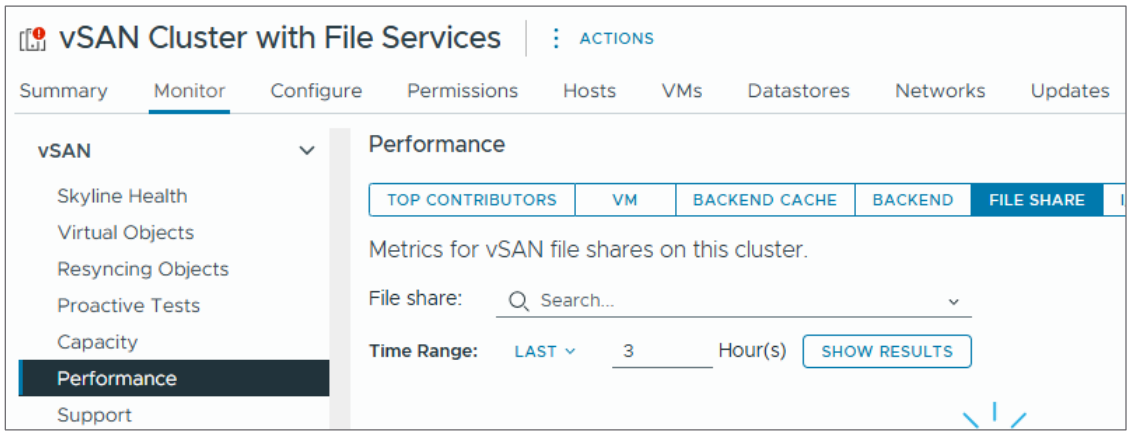

Figure. Time-based performance monitoring of file shares in vSAN File Services.

Capacity Reporting

vSAN also provides capacity reporting specific to File Services. One can view capacity consumption on a per-share basis or view the overall cluster capacity consumption of file shares served by the cluster.

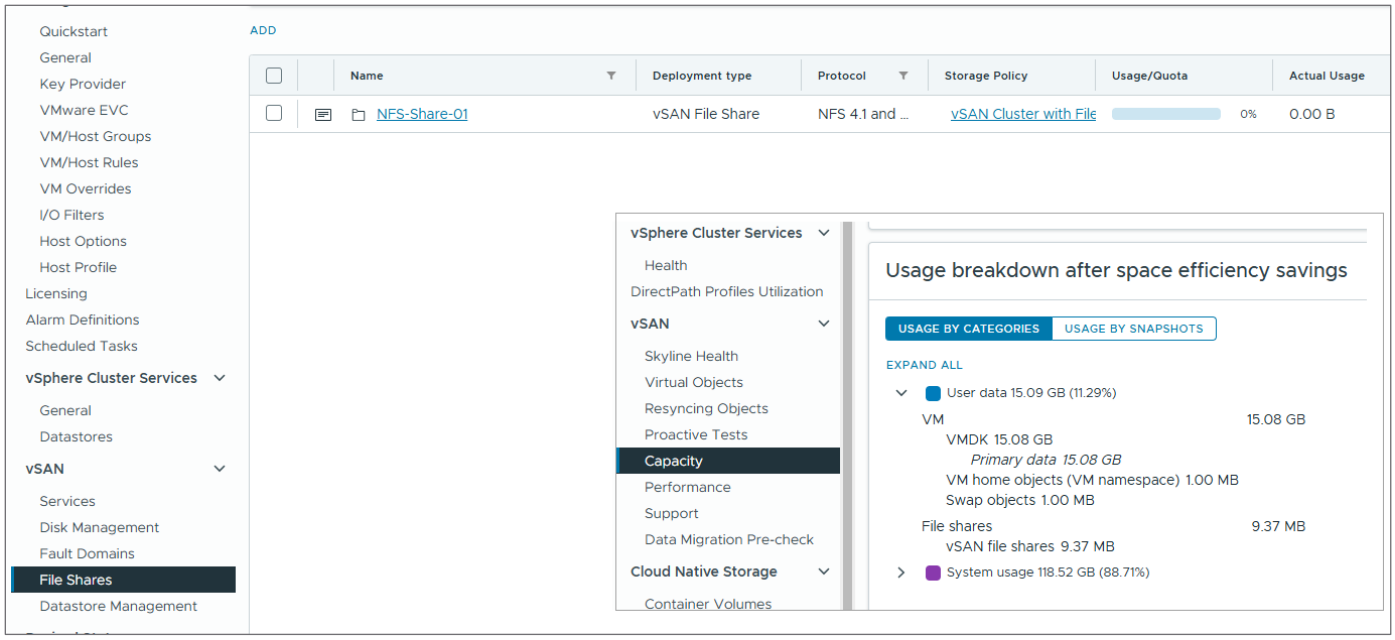

Figure. Per-share capacity and per-cluster capacity information in the UI.

Enforcement through Quotas and Access Based Enumeration (ABE)

File share quotas can be enabled on a per-file share basis. A "share warning threshold" will warn when quota capacity is getting near its limit, and the "share hard quota" will alert when full quota capacity has been met.

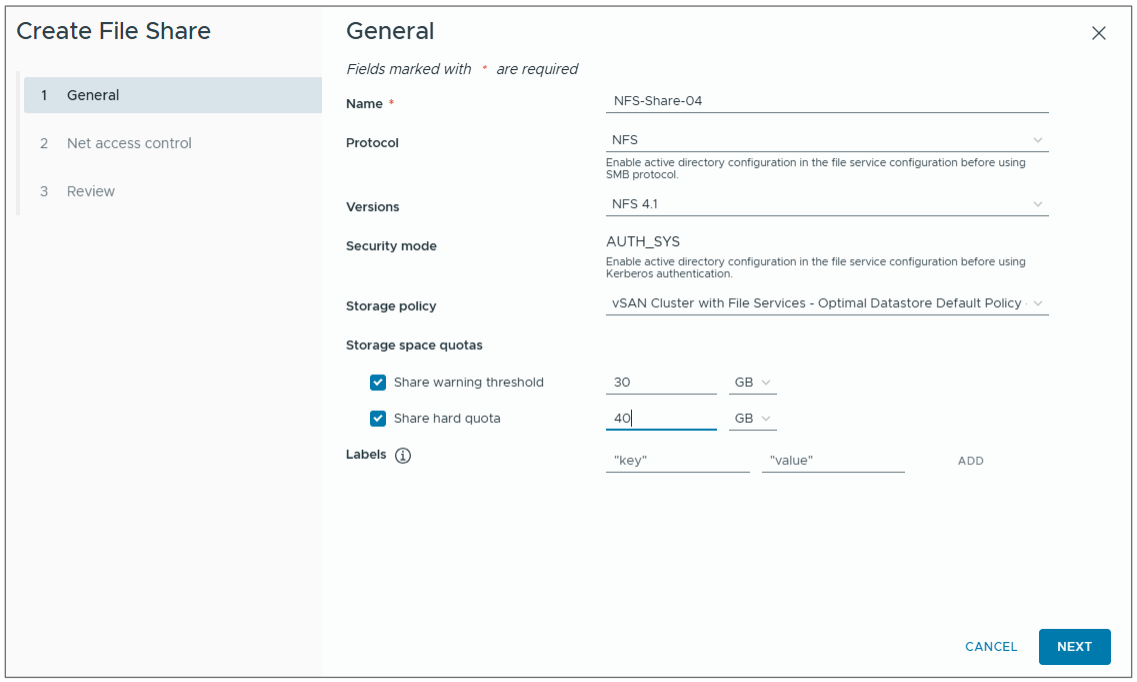

Figure. Establishing share quotas for a specific file share.

vSAN File Services also supports "Access Based Enumeration" (ABE). This is a security mechanism for SMB shares that restricts file and directory visibility within a share. Visibility will be based on user/group read and write permissions of files or directories. This can help reduce unnecessary disclosure of content to those with appropriate permissions.

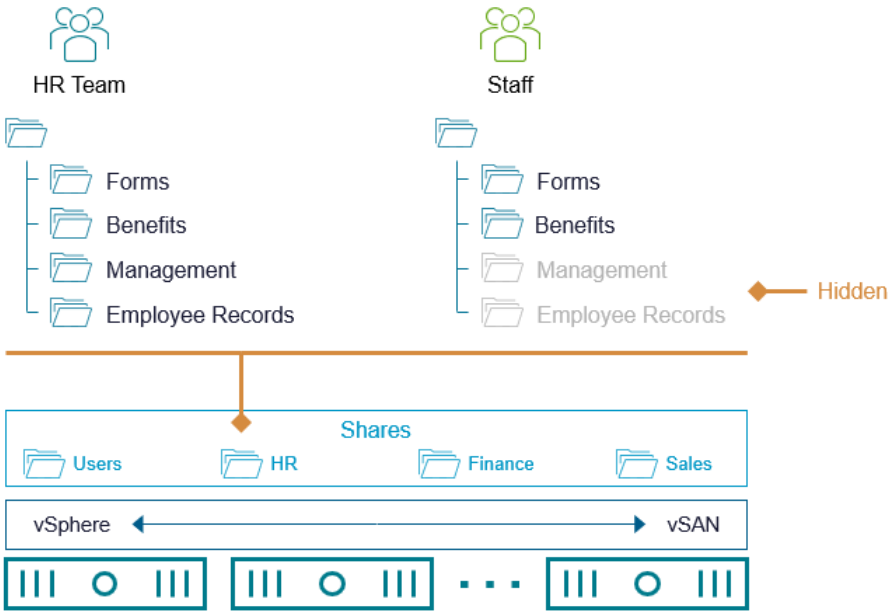

Figure. Support for Access Based Enumeration for SMB file share directories.

Architectural Overview

Distributed Protocol services

When vSAN File Services is enabled, the protocols provided are rendered through a protocol services layer that resides on top of vSAN. This protocol services layer consists of a service or agent VM on each host in the vSAN cluster. The VM is not directly responsible for the processing of I/O. It in fact does not even have an IP address assigned to it. It simply serves as a shell for the instantiation of stateless containers. Agent VMs may need an update after a major version update in vSAN to reflect a new version of PhotonOS, Docker, Ganesha or Samba, but this can be achieved easily in the UI.

These stateless containers running inside of the agent VMs serve as the presentation layer for vSAN File Services. It is the container that binds itself to an available IP address and provides the presentation of the share and the initial processing of I/O for the desired protocol. For example, NFS and SMB services are provided by an installation of Ganesha and Samba respectively in each container. This is a production-quality open-source project that enables vSAN File Services to stay current with a modern distribution of NFS.

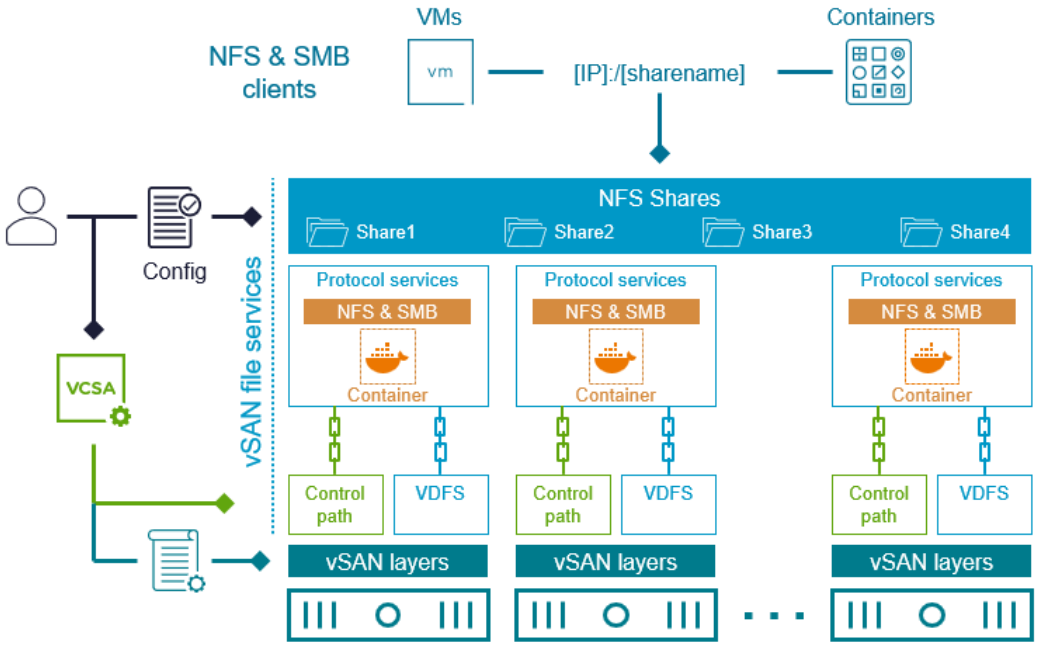

Figure. High-level architecture of vSAN File Services.

When a designated container presents a given file share, all communication to and from the clients that mount that share will be through that specific container. While this can be thought of as a single ingress/egress point for a share, the role of the container is to simply pass the commands and the data through to vSAN's fully distributed file system for vSAN File Services. **The container does not store any of the file share data.** vSAN evenly distributes containers across the hosts that make up the vSAN cluster, and when new shares are created, a share will bind to a given container. This helps provide increased performance as new shares and new hosts are added.

As vSAN File Services are deployed with a small number of shares, each container will typically be responsible for one share, but as share count accumulates, a container may be responsible for many shares. In vSAN File Services for VCF 9.0, each container may be responsible for as many as 50 shares.

Relatively small clusters with fewer hosts may limit the overall share count below the documented maximum supported limit of shares per cluster. This is because there are not enough hosts to instantiate enough containers to present the documented maximum number of supported shares.

Integrated Distributed File System

vSAN File Services uses its own purpose-built, distributed file system specifically for the purpose of file sharing. Known as "vSAN Virtual Distributed File System" (VDFS), it was designed to provide a highly efficient and scalable data path, automated load balancing, high performance. VDFS provides a common platform that is independent from the presentation layers above that provide the protocol services. It will distribute the data path as the number of shares grow and will shard the metadata for improved utilization of resources.

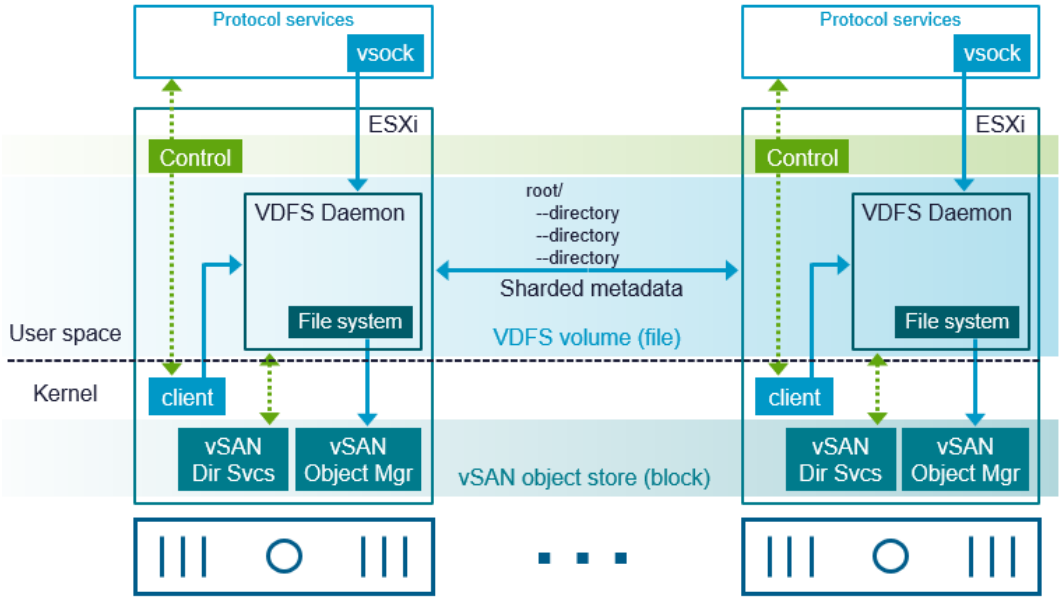

Figure. The role of VDFS in vSAN File Services.

This file system is part of the hypervisor, and not part of the virtual appliances on each host, or the stateless containers. It translates the commands received from the presentation layers into block requests to the lower vSAN layers as it is stored in a vSAN object that represents the share.

File shares live as vSAN objects, with resilience defined by the storage policy associated with the share. As with any other vSAN object, placement decisions of the file share object are made by vSAN's cluster-level object manager (CLOM). This means that the design allows for vSAN to manage the placement of the file shares, where the protocol services layer is only responsible for the presentation of the share itself.

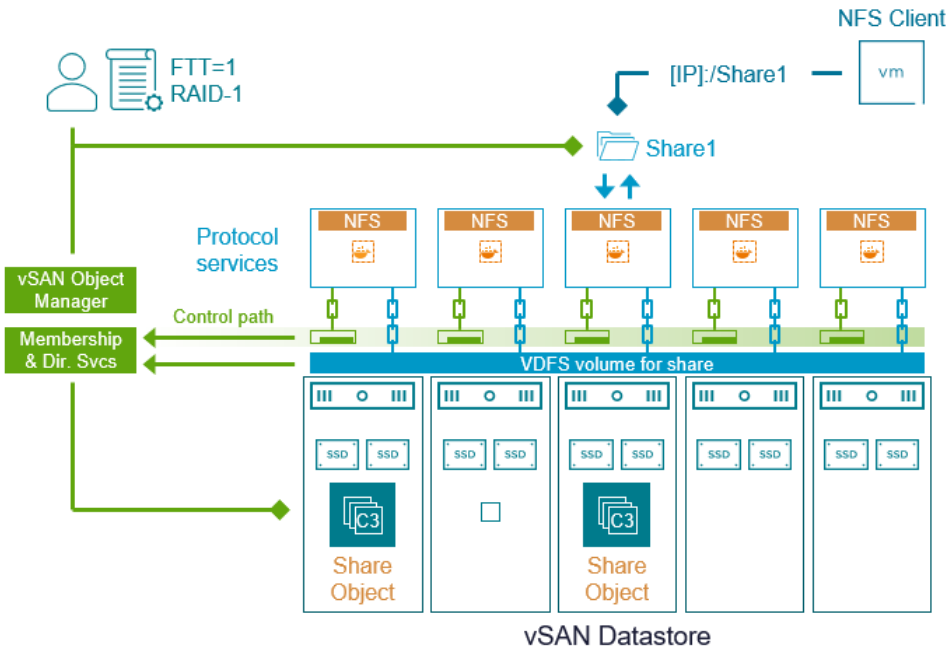

Figure. Data placement of an object representing a file share in vSAN File Services.

Zero-Copy Data Path for Optimal Performance

The architecture of the vSAN file services uses a technique throughout the stack known as "zero-copy." Typically, when a solution is designed in layers, there is inherent latency of I/O because of the "hops" that must be taken to traverse through the stack. vSAN file services is constructed in a way that allows for the use of "zero-copy" techniques, where say, a single write will "pass through" to the destination in only on hop. This means that in the case of VDFS, the protocol service containers, as well as the hypervisor itself can effectively interact directly with VDFS. This technique reduces overhead and queuing of I/O which lowers latency to the VDFS volumes, thereby delivering a fast data path for the clients.

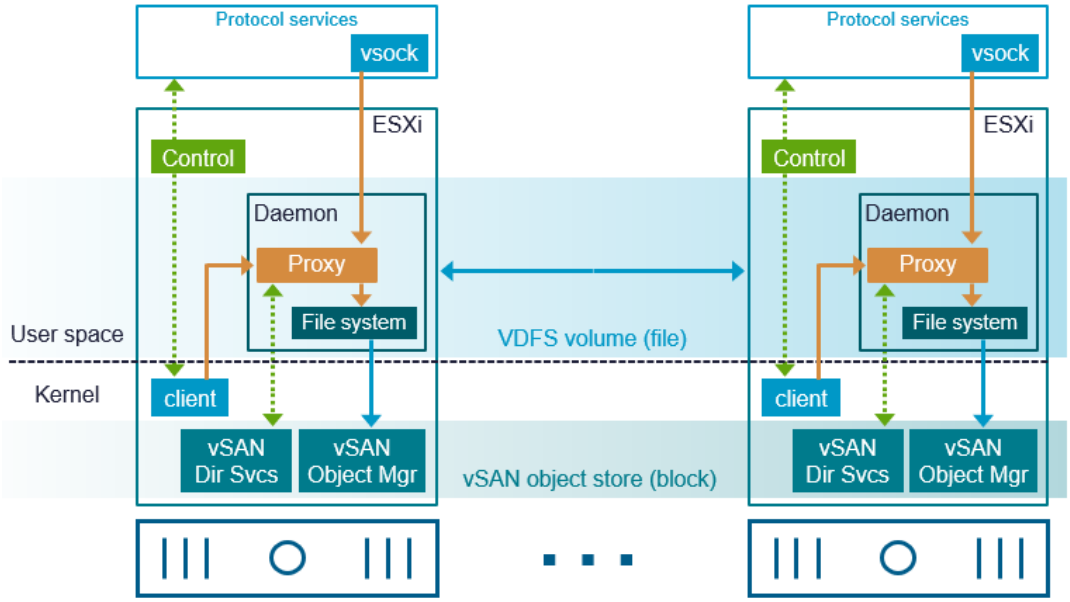

Figure. Data path efficiency using zero-copy data path between protocol services layer and VDFS.

Sophisticated Control Path for Failover and Maintenance

Since vSAN file services uses the same object-based storage that is used with VM's, the logic for maintaining availability of the data remains the same. vSAN file services introduces techniques that provides resilience of the stateless containers that render the service.

Resilience of the share connections – whether they be SMB or NFS – are provided through the use of stateless containers. Under an initial configuration, each host will have one protocol services VM that will instantiate a single container that will present the respective protocol. Note that it is the container itself that has an IP address assigned to it, and not the protocol services VM hosting it.

Under failure or maintenance conditions, containers providing a share (and its associated IP address) will instantiate itself on another host. When there is a reduction of potential hosts for the container to be instantiated elsewhere, to another host that may already be providing another container (as shown in the illustration). The host selected for the container to be instantiated be based on host having the least number of shared served. Once the hosts become available again, a rebalancing of containers may occur.

A rebalance of containers is based on the share count served by the container. If a container does not have a share served, it will not be considered in a rebalance. This is a conservative approach to rebalancing that can be found through many features (vSAN data rebalancing, vSphere DRS, etc.). Rebalancing of the containers will be checked every 30 minutes, and if it needs to occur, may show up in one of the health findings specific to vSAN File Services.

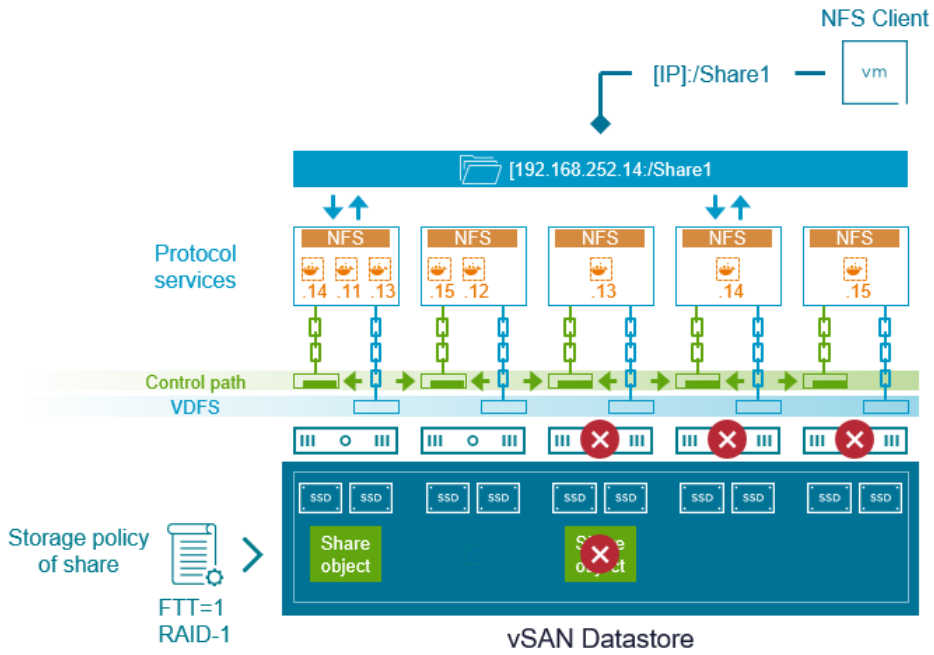

Figure. The instantiation of containers to available hosts during a multiple host failure scenario in vSAN.

During maintenance and failover conditions, it may take some time to instantiate a container from one host over to another. The amount of time this takes can vary depending on several factors. However, with more recent versions using vSAN ESA, this failover time has been reduced significantly.

Scaling Mechanisms

The unique architecture of HCI allows administrators to scale out by adding more hosts, in addition to scaling up (adding more capacity to hosts). For vSAN file services, not only can scaling out increase available capacity, but it can also potentially improve the performance and resilience of the service as it is able to distribute the services across more hosts.

Scaling out vSAN File Services

File shares in vSAN file services will automatically distribute the shares across a cluster. For example, if a 16 host cluster using vSAN file services was using 32 file shares, each host would have approximately two shares per host. This uses the presentation layer (aka the protocol services layer) in the most resource efficient way possible, where each host can be responsible for presenting its fair share of file shares. This applies only to the presentation layer. vSAN already distributes the file share object data and the VDFS metadata across the hosts efficiently.

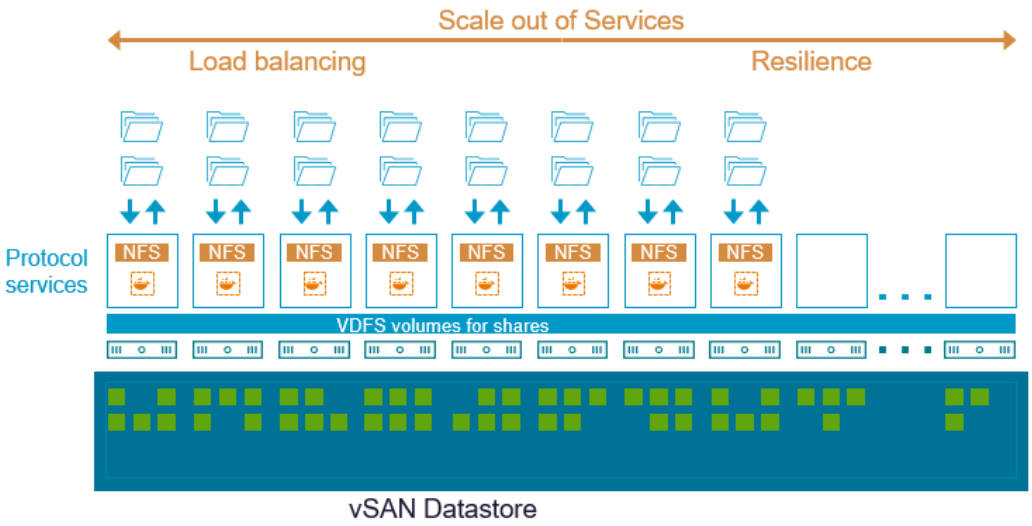

Figure. Scaling out vSAN File Services in a cluster.

Scaling Up vSAN File Services

File share capacities will scale automatically. As additional capacity for a share is needed, vSAN will automatically create additional objects to support growth of the share. These concatenated share objects will be completely transparent to the administrator.

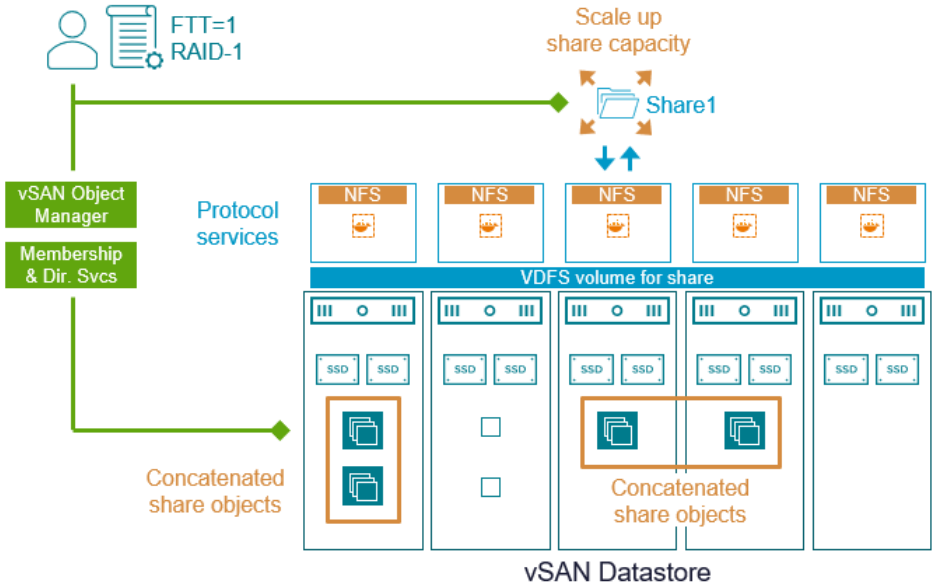

Figure. Scaling up file shares in vSAN File Services.

Deployment Considerations and Limitations

As with any feature, the use of vSAN File Services will introduce additional considerations in design and operation of your environment. See "Limitations and Considerations of vSAN File Services" on Techdocs for a comprehensive list of items to consider when deploying vSAN File Services. Typically, the primary considerations to be aware of include:

- vSAN File Services will require the use of additional IP addresses for the containers that are used to present the file shares. These IP addresses should be reserved for this specific use, and each IP address should have both a forward and reverse-lookup DNS records.
- The container IP addresses for a cluster with vSAN File Services enabled must live on a single network, and the port groups for file services require the use of promiscuous mode and forged transmits.
- A single share cannot be used with both SMB and NFS protocols. It must be one or the other.
- Since the total count of file shares is based on the cluster, consider using relatively smaller clusters to present more file shares than one large cluster.
- vSAN File Services cannot be used to present an NFS datastore to be mounted by an ESXi host and run VMs.
- vSAN File Services has its own snapshot capabilities that are independent from vSAN ESA snapshots that are used as a part of vSAN Data Protection, and VMware Live Recovery (VLR). The snapshot mechanism for vSAN File Services is accessible via API only.
- vSAN File Services do not currently have a native way to replicate the data elsewhere. Replication of file share contents would need to occur inside the guest VMs using tools like rsync for NFS, or Robocopy for Windows.

The vSAN Operations Guide has a section titled "*File Services: Introducing it into an Existing Environment*." This can serve as a helpful starting point for deploying vSAN File Services in your own environment.

Summary

vSAN File Services can help you deliver file share services to any environment where a vSAN cluster resides. It can be especially useful in environments that require the use of file shares while minimizing their hardware and software investments.

Additional Resources

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

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.