



vSAN Cluster Design – Large Clusters versus Small Clusters

Recommendations for vSAN in VCF 9.1

May 5, 2026

Table of Contents

Introduction.....	3
Scope of Topics	3
Why Clustering in vSAN is Different, and Better!	3
Defining Small and Large in the Context of Cluster Sizes.....	3
Considerations with vSAN Cluster Design	4
Needs of the Business	4
Data Resilience requirements	4
Performance Requirements	4
Physical Constraints	4
Applications Behaviors	5
Larger Mixed-Use versus Smaller Purpose-Built vSAN clusters.....	5
Benefits from Smaller Purpose-Built Clusters	5
Benefits from Larger Mixed-Use Clusters	5
Recommendations.....	6
For most standard single site vSAN clusters, aim for 7 to 16 hosts per cluster.	6
If you prefer or have requirements that dictate the use of a larger cluster, stick with no more than 32 hosts per cluster.	7
For relatively small clusters, strive for 4 hosts instead of 3 hosts.	7
Keep all hosts make up a single vSAN cluster in the same rack.	7
For stretched clusters, aim for a minimum of 8 hosts across two data sites (4+4+1).	7
When using the vSAN Fault Domains feature, use at least 12 hosts across 4 racks.	7
Evaluate when you would like to split a larger cluster into two vSAN clusters.	7
Summary	7
Additional Resources	8
About the Author	8

Introduction

Is it better to have larger mixed-use vSAN clusters using many hosts, or do smaller purpose-built clusters using fewer hosts make more sense? This document will lay out considerations to help the reader understand the tradeoffs between environments that use fewer vSAN clusters with a larger number of hosts, versus a larger number of vSAN clusters with fewer number of hosts.

The information shared is primarily targeted at vSAN deployments using an aggregated “vSAN HCI” deployment model. It does not go into detail on the considerations in sizing disaggregated [vSAN storage clusters](#). Information on the sizing of vSAN storage clusters can be found at: “[Design and Operational Guidance for vSAN Storage Clusters](#).”

Scope of Topics

The information provided in this document will assume the use of vSAN 9.1, and/or VMware Cloud Foundation (VCF) 9.1. VCF deployments may have additional requirements and support limitations that fall outside of the scope of this document. **For VCF environments, please refer to the Administration Guide for VMware Cloud Foundation for guidance as it relates to VCF.** The release of vSAN in VCF 9.1 marks the 6th release of vSAN ESA, and as a result, this document will primarily focus on the capabilities in vSAN ESA. For considerations related to the vSAN Original Storage Architecture (OSA), see earlier versions of this document.

Why Clustering in vSAN is Different, and Better!

VMware vSphere helped introduce and popularize the concept of virtualization clusters to modern data centers, where it provides a collection of physical resources for VMs. A cluster is the defined boundary of these physical resources that can be used for a common set of workloads. vSphere clusters focused on the sharing of compute and memory resources, and does so in an incredibly efficient way.

Traditional external monolithic storage used with vSphere was typically shared across many vSphere clusters. This means that **storage is not treated as a cluster resource, as defined by the hypervisor**. VMs from multiple clusters may arbitrarily use and contend for the same shared storage resources, with the storage solution having no understanding of priority. Block-based storage arrays are unable to discern the real characteristics of I/O traffic, as most intelligence is lost once it leaves the host. This can lead to cluster-wide noisy neighbor problems, which makes performance difficult to predict. This architecture also tends to funnel in all I/O to hotspots such as storage array controllers, which may have finite limits in performance that cannot be surpassed. Administrators were typically forced to design and operate their environments around the constraints of monolithic arrays in a three-tier architecture.

vSAN provides storage capabilities as a part of the hypervisor. It aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from using the vSphere Client and vCenter Server. In aggregated vSAN HCI clusters, storage and compute for VMs are delivered from the same server platform running the hypervisor. A vSAN cluster presents storage as a single datastore per cluster. This data store automatically increases in capacity if the cluster is scaled up by adding more storage devices, or scaled out by adding more hosts.

When using vSAN in an aggregated vSAN HCI deployment, **storage is treated as an exclusive resource of the cluster, just like memory and compute resources**. Storage I/O will not cross cluster boundaries, unless you mount the datastore of one vSAN cluster to another vSAN cluster.

Defining Small and Large in the Context of Cluster Sizes

To better understand what is meant by “small” versus “large” cluster host counts in vSAN, let's review the supported minimums and maximums. A single site, standard vSAN cluster can have as few as 3 hosts, and as many as 64 hosts. The most common configurations we see are clusters starting at 4 hosts and going up to as high as 24 to 32 hosts. Exceeding 32-host clusters typically introduces other design considerations and generally does not yield any practical advantage with a cluster that large.

vSAN ESA has continued to deliver higher levels of efficiency with each new release, and as a result, and as noted in the post “[Driving Down Storage Costs with Lower Hardware Requirements for vSAN](#)” allows customers to do more with what

they already have. With modern [vSAN ReadyNode Profiles](#) available for vSAN HCI clusters, vSAN storage clusters, and [Cyber Recovery Clusters](#), customers have the ability to offer a tremendous amount of storage capacity with very few hosts. This can be a highly effective way to lower storage costs, but since vSAN is a distributed storage system, it is wise to ensure you have a recommended minimum host count to meet your data resilience objectives. The information provided in this document will help you determine what may work best for your environment.

Considerations with vSAN Cluster Design

Data center infrastructures exist for one primary reason: Serve the needs of the applications that run on it, for the benefit of the business and consumers of those applications. Clustering in vSphere and vSAN is a resource management construct that provides the ability to allocate workloads to a collection of resources with infrastructure-level resilience. Application owners and consumers should be unaware of the makeup of a cluster. Yet, if designed poorly, cluster sizing decisions can impact on resource efficiency, application performance, and data resilience.

Placing considerations such as capacity and performance aside, **the following five factors can influence vSAN cluster host count and placement decisions.**

Needs of the Business

Shared environments often overlook considerations regarding the **priority of the business**. Business needs may dictate the importance of discrete workloads, and delivering an infrastructure that can provide flexibility to meet a level of predictability is a top priority. Most organizations have their respective "most important consumers" and "least important consumers." Departments or business units in an organization may have different initiatives, levels of funding, and **unique considerations in cost accounting**.

Data Resilience requirements

vSAN is a cluster-level distributed storage solution that requires several hosts to provide the resilience of data. While the maximum size for a standard vSAN cluster is 64 hosts, the minimum number of hosts required in a vSAN cluster depends on your resilience requirements. With standard, single site vSAN clusters, this means a minimum of 3 hosts ensure a failure to tolerate (FTT) level of 1 (FTT=1), and a minimum of 6 hosts to ensure an FTT=2. Although we do recommend more hosts in a vSAN cluster than required by the desired level of resilience. But understanding the desired levels of resilience will help you determine the minimum host count for a vSAN cluster. Thanks to [Auto-RAID in vSAN for VCF 9.1](#), the optimal level of resilience is determined and managed for you! See the "[vSAN Availability Technologies](#)" document for more information.

Performance Requirements

In any type of shared environment, one workload may contend for resources demanded by others, potentially impacting other services. vSphere has incredible resource schedulers that aim to use the shared resources in the most efficient way, but contention may occur based on overwhelming demand. A single noisy neighbor VM can impact many VMs in an environment. It is not uncommon to find that many **critical applications in an environment may not be the busiest**. The end-to-end clustering abilities of vSAN can help mitigate this matter, as it will prevent VMs demanding high storage I/O from impacting VMs in other clusters. Note that discrete performance of a VM on vSAN is derived from the hardware specifications of the hosts, and the network capabilities that connect them. Only the aggregate performance increases (cluster IOPS and throughput) increases as the host count increases. For more information, see the post: "[Performance Recommendations for vSAN ESA](#)."

Physical Constraints

Business requirements are often the primary driver in a cluster design, but physical constraints should not be overlooked. For example, since vSAN is a cluster-based solution, it communicates with all the hosts in the cluster to ensure data is resilient. The decision on how many hosts to place in the same physical rack or different racks will change the network usage across the datacenter. For more information, see the post: "[vSAN Networking – Optimal Placement of Hosts in Racks](#)."

Applications Behaviors

Applications and the demands placed upon them can vary substantially. The demand placed on a physical infrastructure by workloads can be unique to every environment. How much an application reads or writes data, how frequent it does so, and other subtle but important traits of the application can make simple categorization of workloads a challenge. For more information, see the post: "[What is a Workload?](#)"

Larger Mixed-Use versus Smaller Purpose-Built vSAN clusters

With an understanding of how vSAN enhances vSphere's already powerful clustering capabilities, let's review the tradeoffs typically associated with larger mixed-use vSAN clusters versus smaller purpose-built vSAN clusters.

Benefits from Smaller Purpose-Built Clusters

More compartmentalized resource usage

One of the most immediate advantages to smaller clusters using fewer hosts is complete isolation of storage resources from other storage activity outside of the cluster. This can reduce resource contention between VMs, as all I/O is kept within the boundary of the cluster itself. This makes it much easier to identify and manage the requirements of the applications and tailor the resources to needs of the application. By isolating storage, it greatly reduces the impact of noisy neighbors.

Tailor cluster services and limit operational and maintenance domains

Smaller clusters allow you to **apply cluster level services prescriptively**. This would include data services such as deduplication, data-at-rest encryption, data-in-transit encryption, and vSAN File Services. For example, instead of enabling vSAN data-at-rest encryption for all VMs, an environment may choose to only enable that service on one or more clusters where the requirement needs to be met.

Multiple **smaller clusters may make maintenance and unplanned events easier to manage**. Imagine running through the upgrade process for firmware devices on the servers, or perhaps a hypervisor update. Smaller clusters with fewer hosts make the maintenance domain smaller, and more agile. Run with the latest version of vCenter, and then phase in new versions of vSAN on a cluster-by-cluster basis, or upgrade clusters in parallel for a faster upgrade experience. Smaller clusters limit the area of impact on unplanned events as well, such as clusters that have reached a capacity "full" condition. Smaller more agile maintenance domains are a significant advantage over a "largest possible cluster" design.

Better Cost Accounting - Targeted Spending and Easier Approvals

It is common to find organizations where **some departments have higher degrees of financial influence than others**. A company may need high performing or large amounts of storage for their R&D team, but traditional architectures meant that the accounting would hit all departments – which may have negative tax implications. Smaller purpose-built clusters can help target more specific needs for an organization, and specifically, the groups inside of an organization who are requesting more resources. With this model, a vSAN cluster for general office administration could be configured with very modest resources, while the cluster intended for the development group could target their performance-focused goals. Those vSAN clusters can be built with different hardware to reflect those needs but maintain symmetry within the cluster.

When cost accounting is simplified and targeted, this often makes internal purchasing approvals easier and more cost-effective for the administrator.

Can accommodate available rack space more easily.

Smaller clusters tend to be easier to find a home all within a single rack more easily. This can help reduce network resource usage across a network spine, as the intra-vSAN traffic remains within the Top of Rack (ToR) switches

Benefits from Larger Mixed-Use Clusters

More Protection and Space Efficiency Options

With vSAN, data **resilience capabilities are dependent upon how many hosts are in the vSAN cluster**. While a typical vSAN cluster can be as small as three hosts, this limits the resilience capabilities of a cluster to FTT=1. A cluster of 6 or more hosts would provide FTT=2. For more information, see the post: "[Erasure Codes in VMware vSAN versus Storage Arrays](#)."

Additionally, one should always factor in **at least** one additional host (N+1) beyond the minimum required for the prescribed level of protection. This allows vSAN to automatically heal the VMs affected by a failed host, and regain the level of compliance without waiting for an offline host to come back online. Ensuring a cluster is sized for N+1 in **compute and storage capacity** should be included in any cluster sizing process. See the “[vSAN Availability Technologies](#)” document for more information.

Improved Shared N+1 Efficiency

Clusters with more hosts have the potential to use a **larger percentage of resources, with a smaller percentage allocated for failure conditions**. This reduces the impact of a node failure in terms of both compute, storage performance, and storage capacity.

Note that the benefit of increased N+1 efficiency has diminished returns as the cluster size increases. Just as the example shows, as the cluster grows in host count, each host contributes a smaller percentage of resources to the cluster. Therefore, adding an additional host to an 8-host cluster will benefit this N+1 efficiency more than adding a host to a 30-host cluster.

Potentially faster rebuild times in the event of a sustained host failure.

In a sustained host failure, a cluster with a relatively higher number of hosts in the cluster can reconstruct the data faster than a cluster with a relatively few number of hosts in the cluster.

More Cost-Effective for Smaller Environments

For organizations that are smaller, a single mixed-use cluster is often the most cost-effective, and will improve the likelihood of meeting the minimum host requirements for your data resilience requirements. Splitting clusters up too small would reduce data resilience capabilities and tends to leave more stranded capacity. Starting with a single, relatively larger cluster will allow for the most flexibility for these smaller environments.

As cluster size grows in host count, an administrator can evaluate the option of splitting up into multiple clusters later. Management clusters are always a good design practice, whether one is using a traditional three-tier architecture, or environments powered by vSAN. [Management clusters can easily be powered by vSAN](#), and are a good candidate when a small single cluster environment eventually grows in host count, and cluster sizing is revisited.

Note that while management clusters are ideal, they aren't always practical given the size of the environment. vCenter Server can easily run on the same vSAN cluster that it manages.

Improved Space Efficiency when using Global Deduplication in vSAN ESA

vSAN Global Deduplication is now generally available in vSAN for VCF 9.1. This is a cluster-wide capability that will look for duplicate blocks across the entire cluster. Thus, with [vSAN Global Deduplication, the larger the cluster size, the more potential for duplicate blocks to be found and deduplicated](#). [Global Deduplication is generally available in vSAN for VCF 9.1](#).

Recommendations

The following recommendations work well for most of our customers regardless of their size and sophistication.

For most standard single site vSAN clusters, aim for 7 to 16 hosts per cluster.

Starting with at least 7 hosts allows for vSAN to protect the data using FTT=2 and can regain the prescribed level of resilience in the event of a sustained host failure. See the “[vSAN Availability Technologies](#)” document for more information. Limiting the cluster size to 16 helps keep a single vSAN cluster within the same rack (assuming 2U servers in a 42U rack). This helps keep vSAN intra-cluster traffic within the ToR switches. For more information, see the post: “[vSAN Networking – Optimal Placement of Hosts in Racks](#).”

If you prefer or have requirements that dictate the use of a larger cluster, stick with no more than 32 hosts per cluster.

Limiting cluster sizes to no more than 32 hosts will help in the overall data storage capabilities of vSAN. There is little practical benefit to cluster sizes exceeding 32 hosts. **If you have 64 hosts that you wish to provide vSAN storage, create 2, 32-host clusters, or even perhaps 4, 16-host clusters.**

For relatively small clusters, strive for 4 hosts instead of 3 hosts.

Always using 4 hosts instead of 3 hosts in a vSAN cluster will ensure that in the event of a sustained host failure, the data can regain the prescribed level of resilience without waiting for a host to come back online. See the [“vSAN Availability Technologies”](#) document for more information.

Keep all hosts make up a single vSAN cluster in the same rack.

For modern spine-leaf topologies, keeping all the hosts in a vSAN cluster in the same rack will keep vSAN intra-cluster traffic within the ToR switches. This reduces the burden on the network spine, which may inadvertently hinder performance because it is oversubscribed. For more information, see the post: [“vSAN Networking – Optimal Placement of Hosts in Racks.”](#) Note that vSAN storage clusters, and vSAN HCI clusters using datastore sharing will typically communicate to the other clusters across a network spine, as the client and server clusters will reside in different racks.

For stretched clusters, aim for a minimum of 8 hosts across two data sites (4+4+1).

An 8-host stretched cluster will allow for data to be mirrored across both sites, with a secondary level of resilience within each site of FTT=1. With 4 hosts in each site, vSAN can automatically regain the prescribed level of resilience in the event of a sustained host outage. See the [“vSAN Stretched Cluster Guide”](#) for more information.

When using the vSAN Fault Domains feature, use at least 12 hosts across 4 racks.

The vSAN Fault Domains feature is typically employed to provide rack-level resilience. In these configurations, it is advised to have at least 4 racks, with 3 or more hosts in each rack. This will allow for an entire rack to fail, and the data to remain available. A total of 6 racks (and 18 hosts) would be the appropriate design if you need the ability to tolerate two rack failures. For more information, see the post: [“Using Fault Domains in vSAN ESA.”](#)

Evaluate when you would like to split a larger cluster into two vSAN clusters.

Sometimes you may have scaled out an existing cluster to the point that splitting the cluster may be a consideration. As the cluster host count exceeds some of the recommendations above, splitting the cluster may provide some operational agility that is lost with larger clusters.

This same principle applies to stretched clusters. For example, you may be considering a 40-host vSAN stretched cluster (20+20+1), which is the maximum supported cluster size in a stretched cluster topology. However, you may find it more manageable to have two 20-host stretched clusters (10+10+1) for better manageability.

Summary

vSAN offers flexibility for your business in ways that were previously unachievable given normal cost constraints. Now the power of choice is given to administrators of vSAN environments and takes clustering to the next level. So with that, here are some simple takeaways from the information provided above:

- There can be a trade-off between workload isolation and cost/feature efficiency.
- Extremely small clusters may limit vSAN features.
- Extremely large clusters may offer limited benefit, and unnecessarily enlarge the management domain.
- vSAN allows for clustering that is dictated by the needs of the business and the consumers of those services.
- Evolving needs of a business may dictate new requirements of a cluster. Thanks to vSAN, this can be accommodated easily.

Additional Resources

The following are a collection of useful links that relate to cluster design and sizing.

[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 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.

