



vSAN Cluster Design - Large Clusters Versus Small Clusters

Table of contents

vSAN Cluster Design - Large Clusters Versus Small Clusters	3
Executive Summary	3
Why Clustering with vSAN is Different, and Better!	4
Considerations with vSAN Cluster Design	7
Applications Behaviors	7
VM Priority	7
Needs of the Business	8
Data placement scheme and resilience settings	9
Larger Mixed-Use versus Smaller Purpose-Built vSAN clusters	10
Benefits from Smaller Purpose-Built Clusters	10
Benefits from Larger Mixed-Use Clusters	12
vSAN Cluster sizes when using datastore sharing	16
Disaggregation with the Express Storage Architecture	16
Consume Storage Externally using vSAN Stretched Clusters	16
Consume vSAN Storage across vCenter Server Instances	17
Example #1 - Organization with 12 hosts	18
Example #2 - Organization with 16 hosts	19
Example #3 - Large Organization with 128 hosts	20
Example #4 - Organization with 16 hosts across two sites	21
Takeaways	22
About the Author	23

vSAN Cluster Design - Large Clusters Versus Small Clusters

Executive Summary

Data center administrators often struggle in making design and implementation decisions for their virtual environments that meet the needs of their organization, but fit within the constraints of their infrastructure. vSphere cluster sizing is one of those decisions, but when using a traditional three-tier architecture, vSphere is limited to the types of resources that it can cluster. vSAN extends the capabilities of vSphere clustering in ways that allow for all new levels of opportunity to design an infrastructure that is custom-tailored to best meet the needs of the organization. It is this flexibility that raises questions as to the best approach to cluster sizing. Is it better to have larger mixed-use vSAN clusters using many hosts, or does smaller purpose-built clusters using fewer hosts make more sense.

This document will discuss the characteristics that make vSAN clustering different than typical vSphere clusters in three-tier architectures. It 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. This document will also discuss considerations of cluster sizing when used with the new vSAN HCI Mesh feature introduced in vSAN 7 U1.

Four examples are included that will serve as a way to understand the considerations and decision process in determining the cluster host count that may work best for an organization's vSAN powered environment. These examples use fictitious circumstances and offer up recommendations based on assumptions. Other cluster configurations not listed do not constitute a poor or incorrect design.

Why Clustering with 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 cluster-specific resources for VMs. Unfortunately, monolithic shared storage systems are plagued with challenges inherent to the architecture, especially when demand and/or scale increases.

The nature of a monolithic, externally shared storage architecture 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. As shown in Figure 1, 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. Storage arrays cannot be cut in half, and thus, administrators are often forced to design and operate their environments around the constraints of monolithic arrays in a three-tier architecture

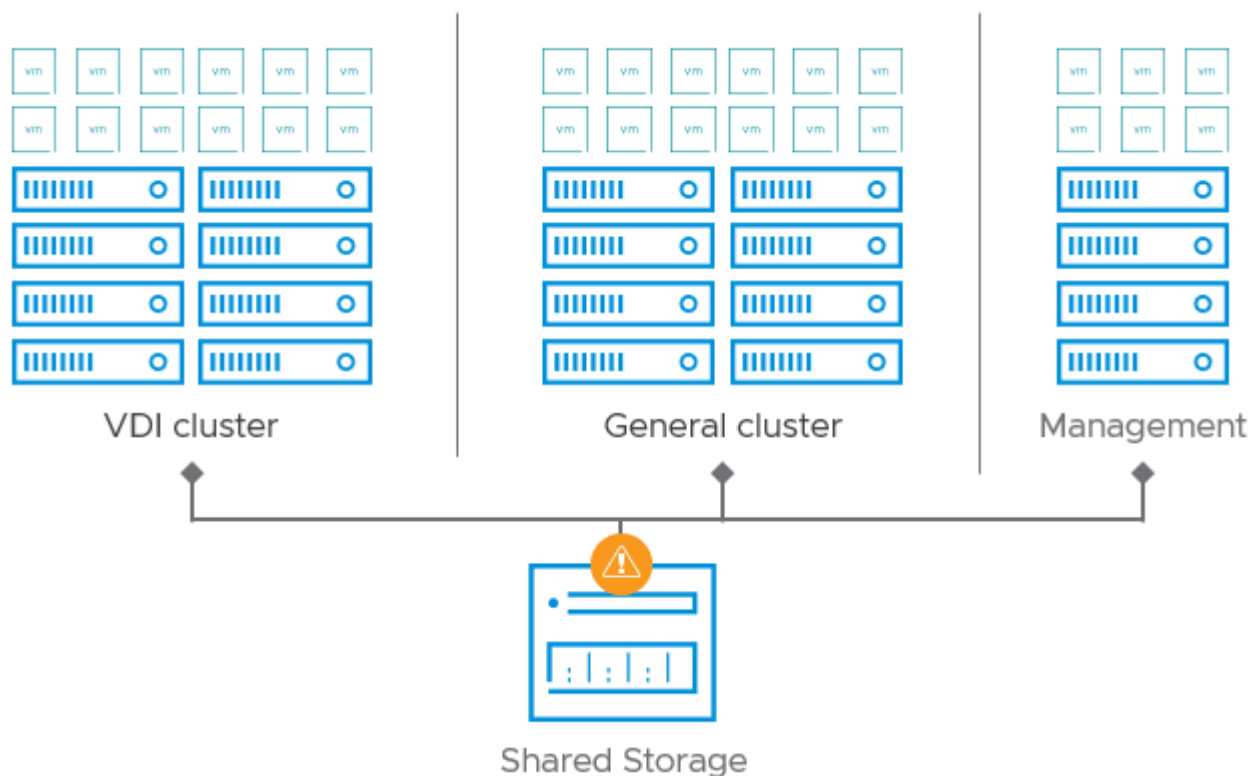


Figure 1. Understanding the lack of relationship between clustering, and traditional shared storage

vSAN is VMware's software-defined storage solution. It abstracts and aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from vCenter. Storage and compute for VMs are delivered from the same server platform running the hypervisor. It integrates with the entire VMware stack, including features like vMotion, HA and DRS. VM storage provisioning and day-to-day management of storage SLAs can all be controlled through VM-level policies that can be set and modified on-the-fly.

Integration into the hypervisor means that vSAN can run more VMs per host than HCI systems not integrated into the hypervisor, and can do so with much more consistent levels of performance. This not only gives vSAN the ability to scale up by adding more storage devices, but also scale out by adding more hosts to a cluster, as shown in Figure 2. This flexibility means that users can design their vSAN powered clusters to accommodate growth in multiple ways to best meet the needs of the organization.

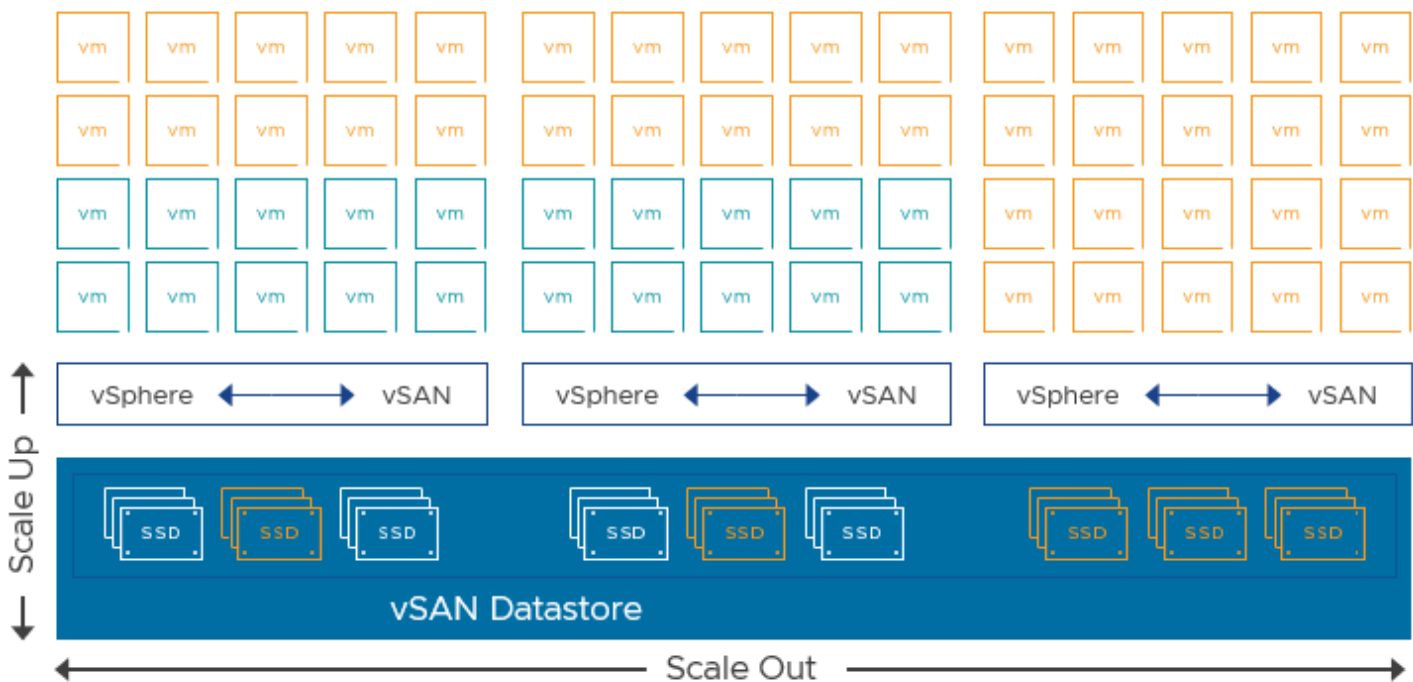


Figure 2. Integrated, distributed storage with VMware vSAN allows for scaling up or scaling out

When using vSAN, vSphere clusters treat block storage as a cluster resource. It completely isolates storage traffic to just storage I/O being served by the VMs living in that cluster. As Figure 3 illustrates, **I/Os will not cross cluster boundaries**, avoiding the problem of traditional storage architectures. vSAN allows the hypervisor to have complete end-to-end control of storage.

Administrators can easily create modular, cluster-aware storage environments that separate workloads while being **managed through a common control plane: VMware vCenter**. This ultimately allows administrators and architects to deploy arrangements based on the needs of the applications and organization.

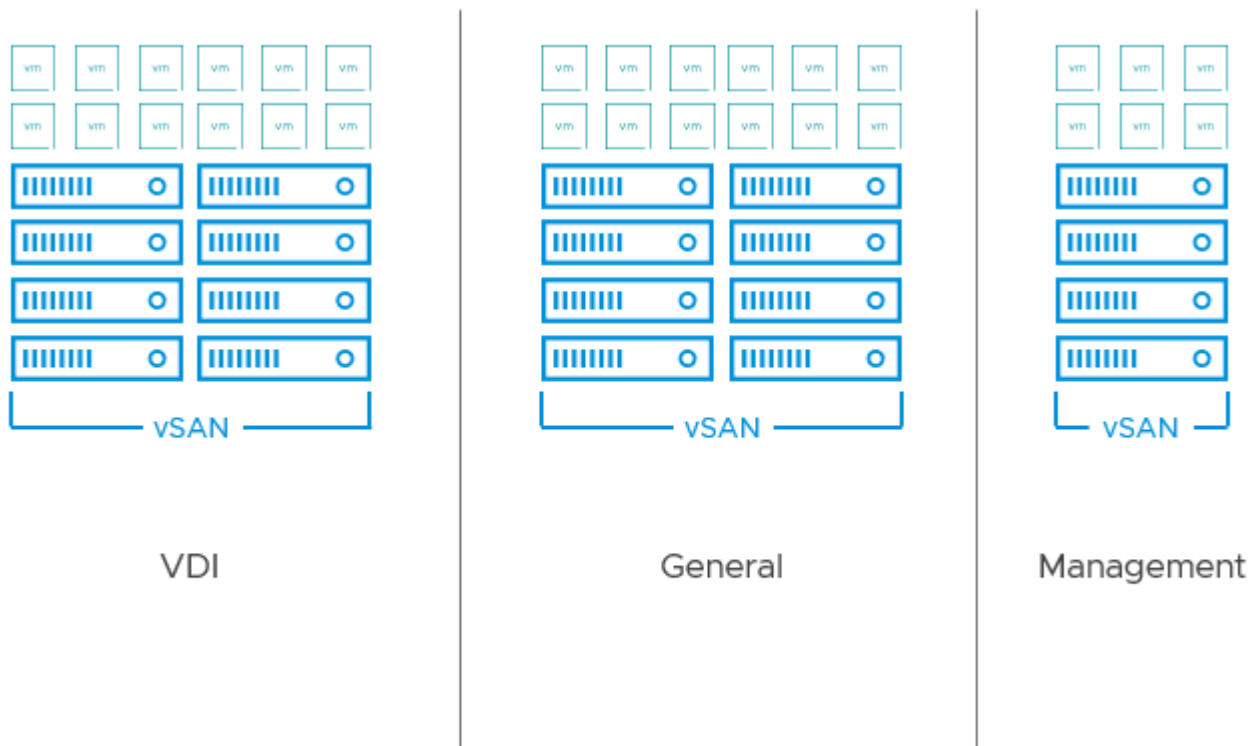


Figure 3. vSAN treating storage as a cluster resource

vSAN clusters can easily connect up to existing three-tier architectures, as shown in Figure 4. This means that a vSAN cluster will have the single, cluster-wide datastore served by vSAN, but the hosts can also connect to traditional array-based storage, whether it be block, file, or vVols. This is something that is often not supported by other HCI solutions.

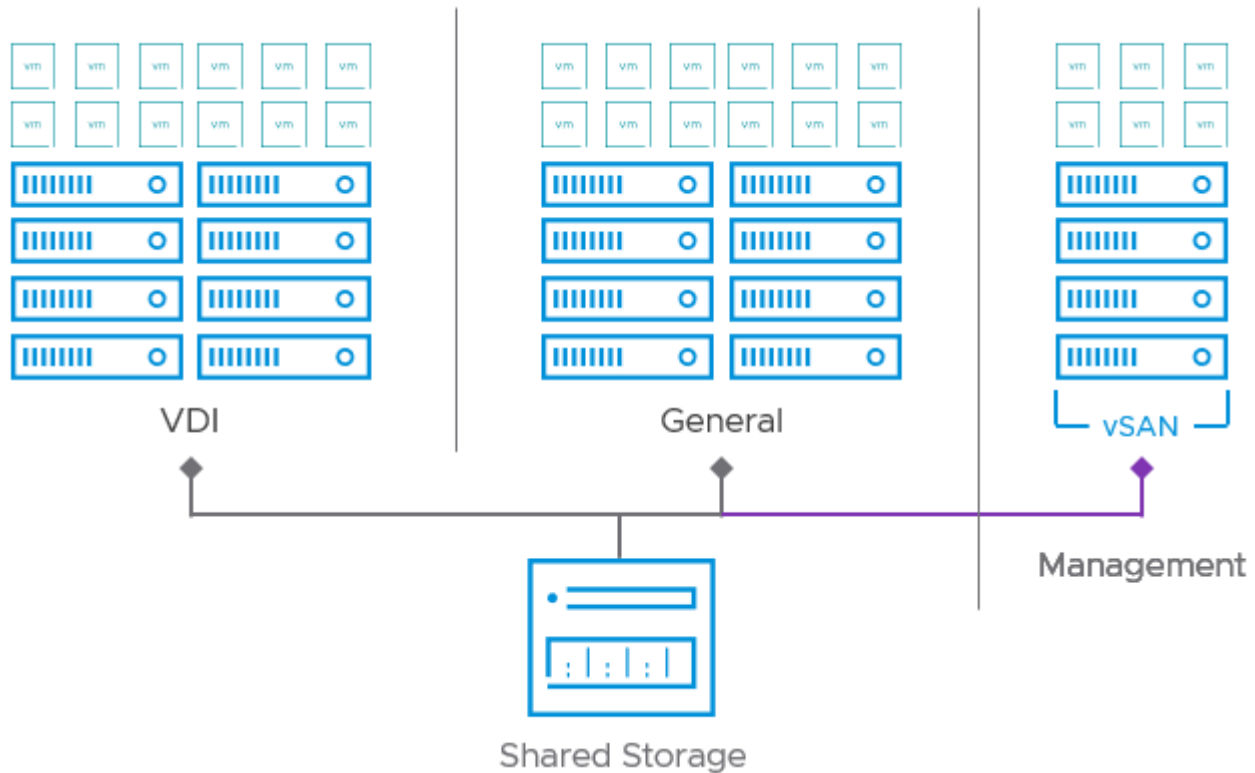


Figure 4. vSAN clusters can enjoy exclusive cluster-based storage and shared storage at the same time

vSAN's ability to extend the capabilities of vSphere clustering, yet easily work with traditional infrastructures, gives the reader an opportunity to revisit cluster strategies for vSAN powered environments.

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 application performance and resilience. Applications share resources in a cluster, and if the cluster design results in a significant amount of resource contention, this can result in poor application performance and processes that take longer to complete for the users.

Placing considerations such as capacity and performance of discrete devices aside, the following four factors can influence vSAN cluster host count decisions.

Applications Behaviors

Applications and the demands placed upon them can vary substantially. The characteristics of a specific application may place a higher level of burden on the infrastructure than other applications. Application and workload behaviors of one application can impact other applications. They may have unique tendencies on their own, but workload demands 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.

Workload characteristics in terms of storage I/O patterns are most often the result of:

- OS and application type
- OS and application demand by users or processes
- Workload duty cycle - the length in time and effort of a workflow, and the frequency of repetition over a given period of time

Reads and writes have different impacts on any type of storage system, as the effort to write data is typically more significant than to read data. Furthermore, most workloads will use a broad distribution of storage I/O sizes in which one I/O can be 256 times the size of another I/O. This can result in workloads committing a relatively few large I/O writes impacting the performance of smaller I/O reads.

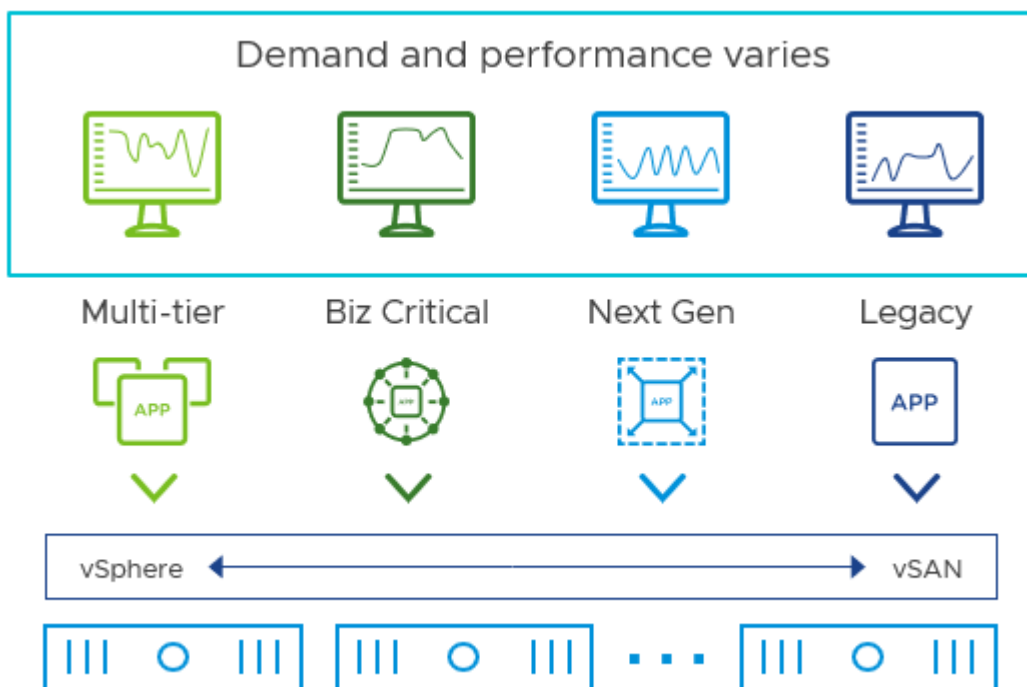


Figure 5. Different application types demanding unique needs from the infrastructure

VM Priority

In any type of shared environment, one application may contend for resources demanded by others, and impact other applications

in the same shared environment. This is inherent to the idea of sharing resources. The hypervisor does allow for some levels of control around this, but it is often left for the administrator to decide what is best, which may lead to non-optimal settings. A single noisy neighbor VM can impact a large number of 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. This will deliver more consistent performance across all workloads.

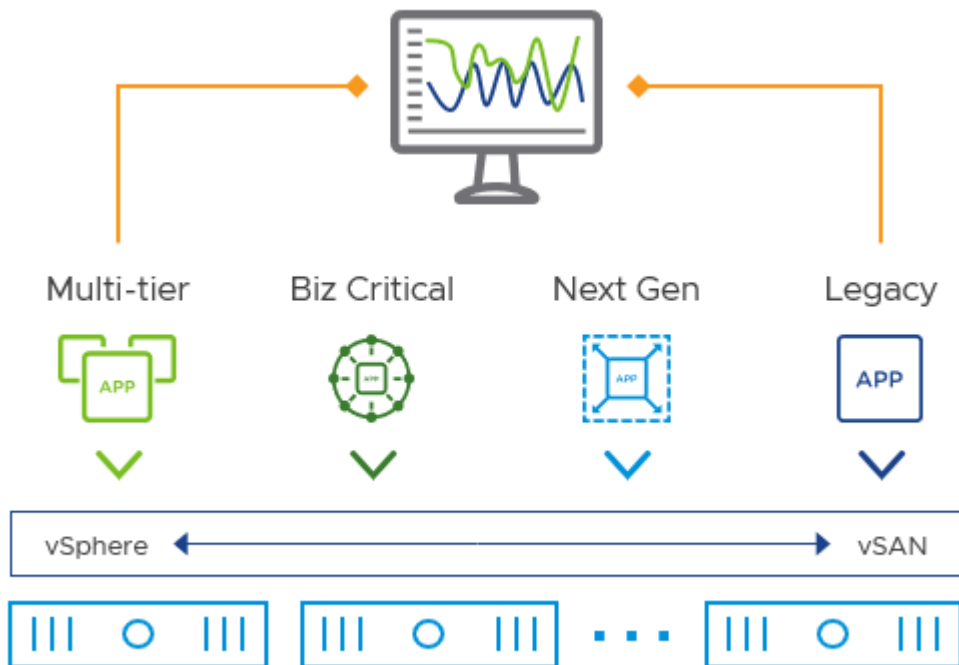


Figure 6. The workload of one VM may impact the workload of others

Needs of the Business

Shared environments often overlook considerations regarding the **priority of the business**. Business needs may dictate the importance of these discrete workloads, and delivering an infrastructure that can provide the flexibility to meet a level of predictability is a top priority. Most organizations have their respective "most important consumers" and "least important consumers." Unfortunately, three-tier architectures cannot treat storage as a cluster resource, and arbitrarily processes I/Os with no understanding of priority when contention occurs.

Departments or business units in an organization may have different initiatives, levels of funding, and **unique considerations in cost accounting**. Traditional shared storage often prevented the ability to take advantage of these potentially dramatic cost savings. Figure 7 represents these considerations in cost allocations and levels of importance.

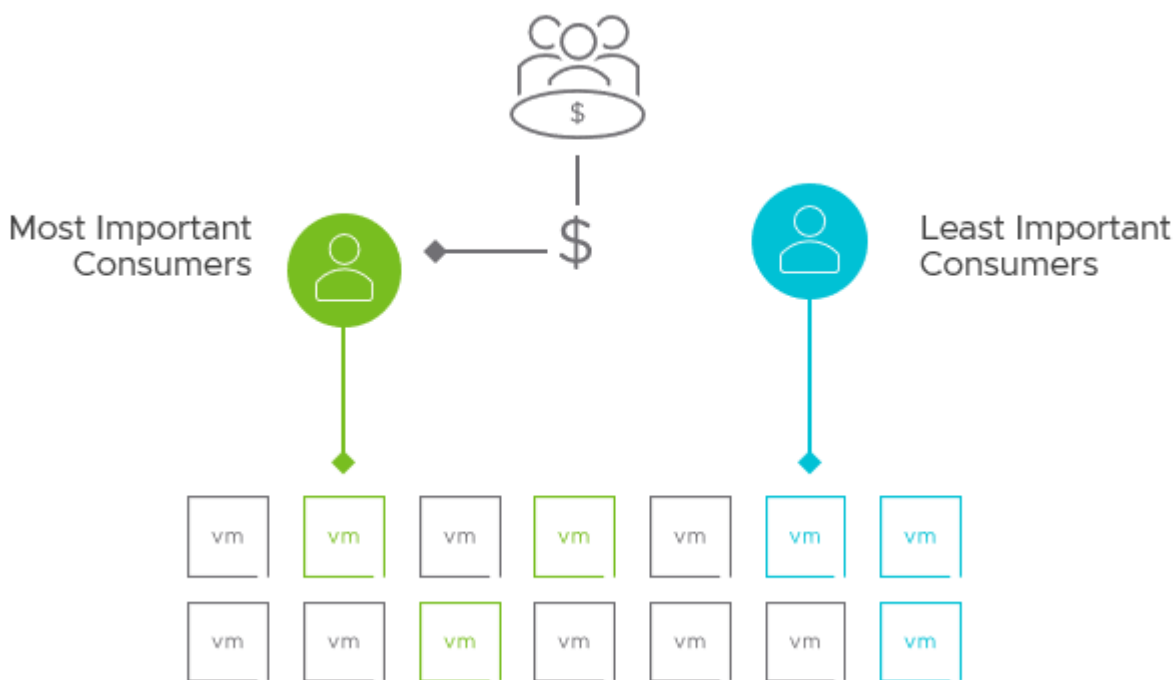


Figure 7. Most important consumers, the least important consumers, and cost allocation

Data placement scheme and resilience settings

vSAN is a cluster-level distributed storage solution that requires a number of 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 the data placement and resilience settings assigned to the objects using storage policies. Understanding the desired protection levels for your VMs will help in determining the minimum host count for a vSAN cluster. See Figure 11 later in this document for more information.

Note that vSAN is an object-based storage system that does not stripe data across all nodes in a cluster. The object data is only spread across the number of hosts required to comply with the assigned storage policy. **Therefore, in most cases, increasing the number of hosts in a cluster will not improve the performance of a discrete VM, unless there is contention from other VMs in that cluster.** In other words, a single VM running a tier-1 application living on a 12 node vSAN cluster will not see improved performance if the same cluster has been increased to 18, 24 or more hosts in a cluster. For this reason, extremely large vSAN clusters offer limited material benefit, while imposing some limitations for the data center administrator. This document describes these trade-offs in more detail.

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

1. Reduce Noisy Neighbor Scenarios

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.

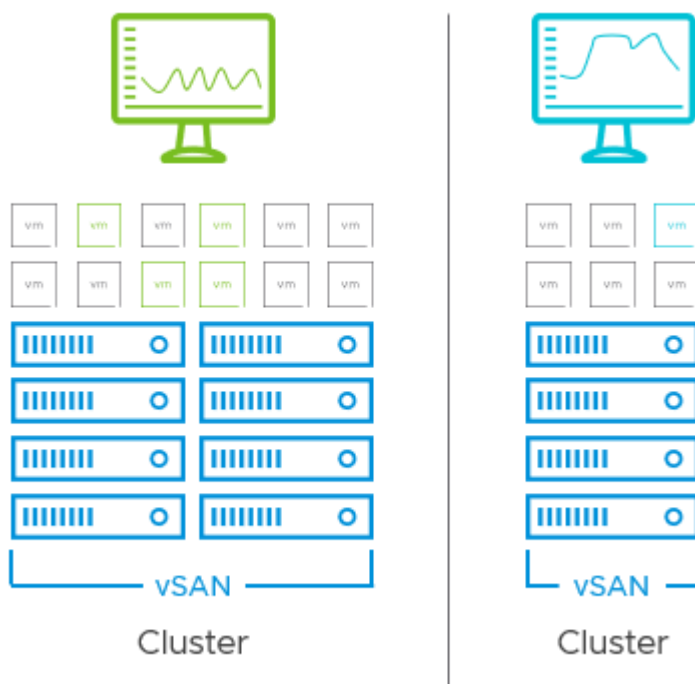


Figure 8. Isolation storage I/O to reduce noisy neighbors.

2. 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 and compression, compression only, data-at-rest encryption, data-in-transit encryption and vSAN iSCSI 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 is one of the biggest advantages over a "largest possible cluster" design.

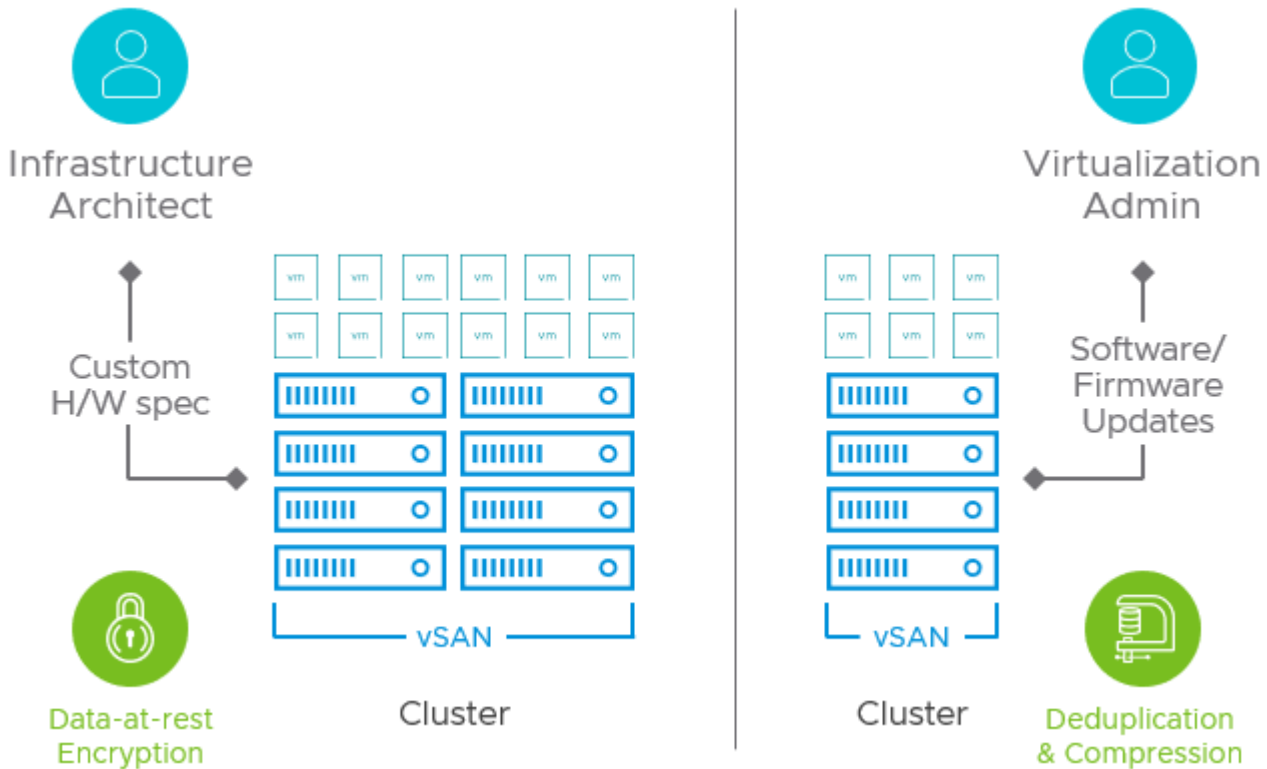


Figure 9. Tailoring cluster services, and limiting maintenance domains

3. 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 to 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 data center administrator.

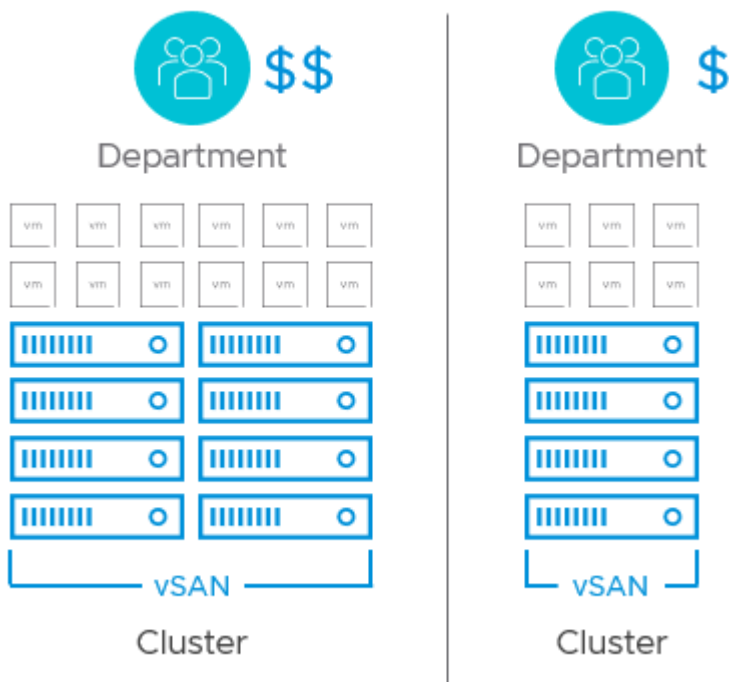


Figure 10. Simplified cost accounting, and targeted purchasing

Benefits from Larger Mixed-Use Clusters

1. More Protection and Space Efficiency Options

Each host in a vSAN cluster will serve up a respective amount of storage capacity, and storage performance for VMs. Unlike traditional three-tier architectures, **some data resilience functionality is dependent upon how many hosts comprise the vSAN cluster.** While a typical vSAN cluster can be as small as three hosts, this does not allow for some of the advanced space efficiency capabilities like RAID-5/6 erasure coding, nor does it allow for higher levels of failure to tolerate. Figure 11 details the minimum number of hosts (in gray) needed to achieve a level of failure to tolerate (using RAID-1 mirroring, and RAID-5/6 erasure coding) assigned to a VM object using a storage policy.

Additionally, one should always factor in **at least** one additional host (N+1) beyond the minimum amount needed for the given levels of protection: Illustrated as the host in green in Figure 11. This allows vSAN to automatically heal the VMs affected by a failed host, and regain the level of compliance assigned by the storage policy, without waiting for an offline host to come back online. Ensuring a cluster is sized for N+1 or N+2 in **compute and storage capacity** should be included in any cluster sizing process.

Note that with clusters that exceed 32 hosts, additional capacity considerations will come into play. vSAN has a per host component limit of 9,000 per host, thus a 32 host cluster would have a maximum component limit for the cluster of **288,000**. For clusters that exceed 32 hosts, the maximum component limit of the cluster remains at 288,000. For environments that have a high density of object components (such as VDI), it will be most efficient to keep cluster sizes at 32 hosts or smaller. For more information on what components are, why the limitation exists, and ways to minimize hitting this limit, see: "[Monitoring and Management of vSAN Object Components.](#)"

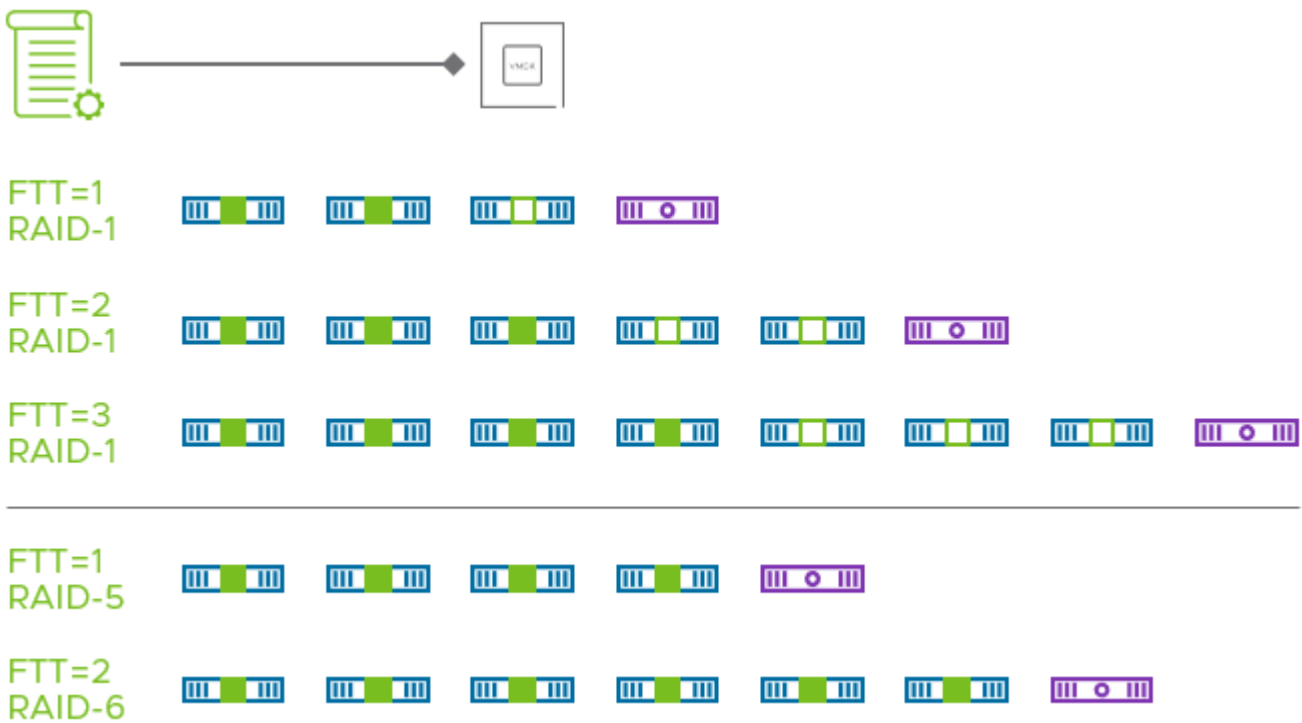


Figure 11. More protection and space efficiency options with larger cluster sizes

Note that a user may have 99 VMs in a vSAN cluster that only need RAID-1 mirroring with a level of failure to tolerate of 1 (FTT=1). If one VM requires a higher level of failures to tolerate (FTT=2 or greater), they will need more hosts in the vSAN cluster to satisfy the requirement for just that one VM.

2. Improved Shared N+1 or N+2 Efficiency

Larger clusters 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. Many storage administrators are familiar with this concept in terms of a storage array using RAID, and similar concepts apply to a hyper-converged environment.

For example, assuming the need to absorb resources in other hosts to accommodate a single node failure, **a 4 node cluster would commit 25% of resources for this N+1 redundancy**, while a **16 node cluster, as illustrated in Figure 12, would provide just over 6% of its resources for an N+1 redundancy**. This provides improved cost efficiencies and is easier to provide spare resources to tolerate failure. This characteristic is also now built into vSAN's reserved capacity calculations (in vSAN 7 U1 and later). The new "Reserved Capacity" calculation is comprised of two elements: An "Operations Reserve" and a "Host Rebuild Reserve." It is the Host Rebuild Reserve that allows vSAN to reserve a smaller percentage of capacity as the cluster size increases. These calculations are accommodated in the [vSAN ReadyNode Sizer](#).

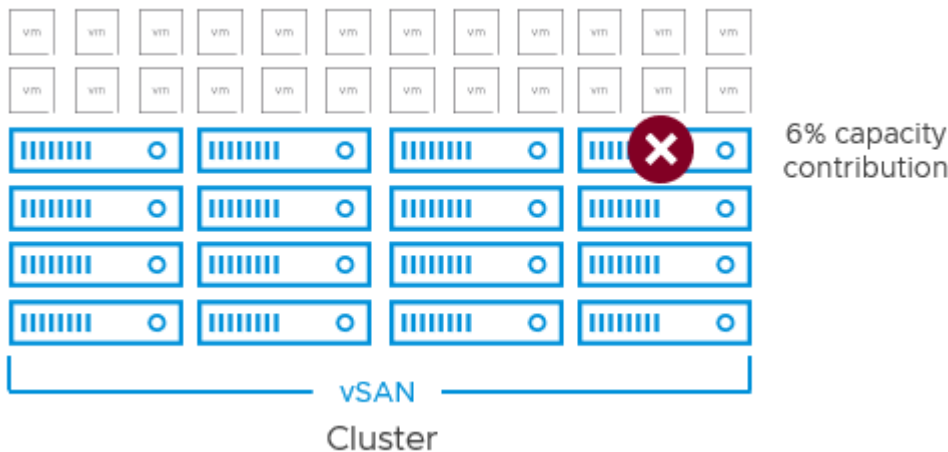


Figure 12. Larger clusters, and reducing the percentage of impact on compute, storage performance, and storage capacity

Note that the benefit of increased N+x efficiency has diminishing 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 node cluster will have much more of an impact improving N+x efficiency than adding a host to a 30 node cluster.

Unlike other solutions, vSAN scales down risk as the cluster increases. This is because vSAN is an object-based storage system that does not stripe data across all nodes in a cluster. Therefore, adding more hosts to a vSAN cluster does not enlarge the failure domain, as can happen with other solutions. Some solutions that automatically distribute data across every potential node in a cluster can introduce more points of potential failure. With vSAN, this isn't the case. Figure 13 illustrates the chance of impact (in percentage) of a double failure impacting VM availability when protecting an object using RAID-1 (FTT1), RAID-5 (FTT-1), and RAID-6 (FTT2)

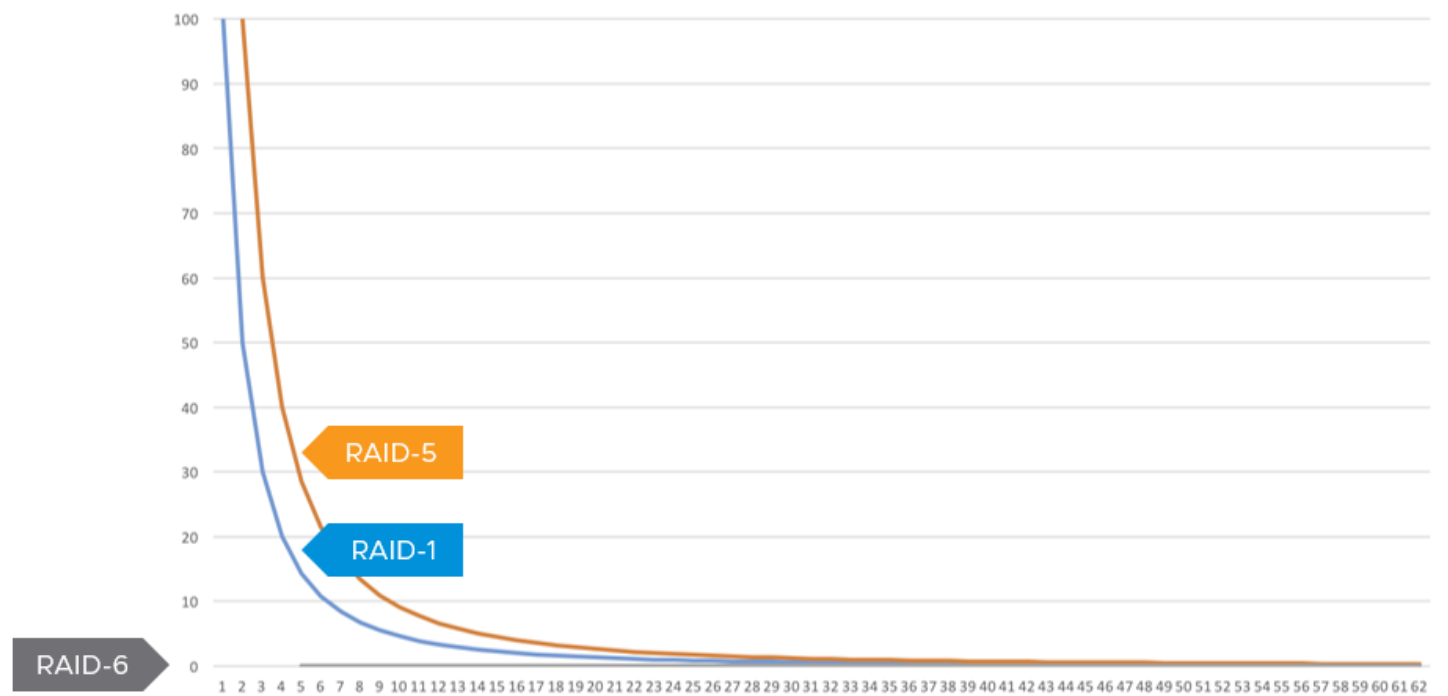


Figure 13. Chance of impact with two host failures vs cluster size (shown as a percentage)

3. More Cost-Effective for Smaller Environments

For organizations that are smaller, a single mixed-use cluster, as shown in Figure 14, is often the most cost-effective, and will improve the likelihood of meeting the minimum host requirements for various data placement and protection options. Splitting clusters up too small would reduce storage policy assignment options for levels of failures to tolerate, and limit space efficiency settings. Beginning with a single larger cluster will allow for the most flexibility for these smaller environments.

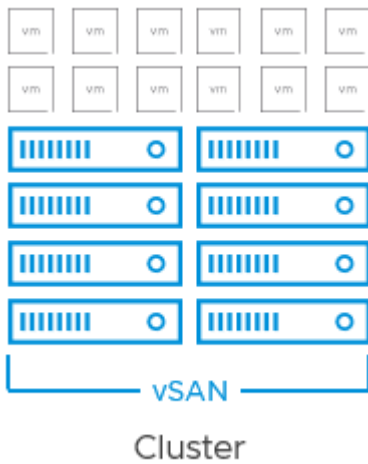


Figure 14. Single mixed-used cluster for smaller organizations

As a host cluster size grows, an administrator can evaluate the option of splitting up into multiple clusters at a later time. The process is relatively straightforward. 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.

vSAN Cluster sizes when using datastore sharing

VMware HCI Mesh is a new feature introduced with vSAN 7 U1 that allows a VM instance to run off of the hosts in one vSAN cluster while using the storage resources of another vSAN cluster. This offers tremendous flexibility for administrators to "borrow" or use capacity across these independent vSAN clusters should their workload demands be unable to meet their current cluster design. See the [VMware vSAN HCI Mesh Tech Note](#) for more details.

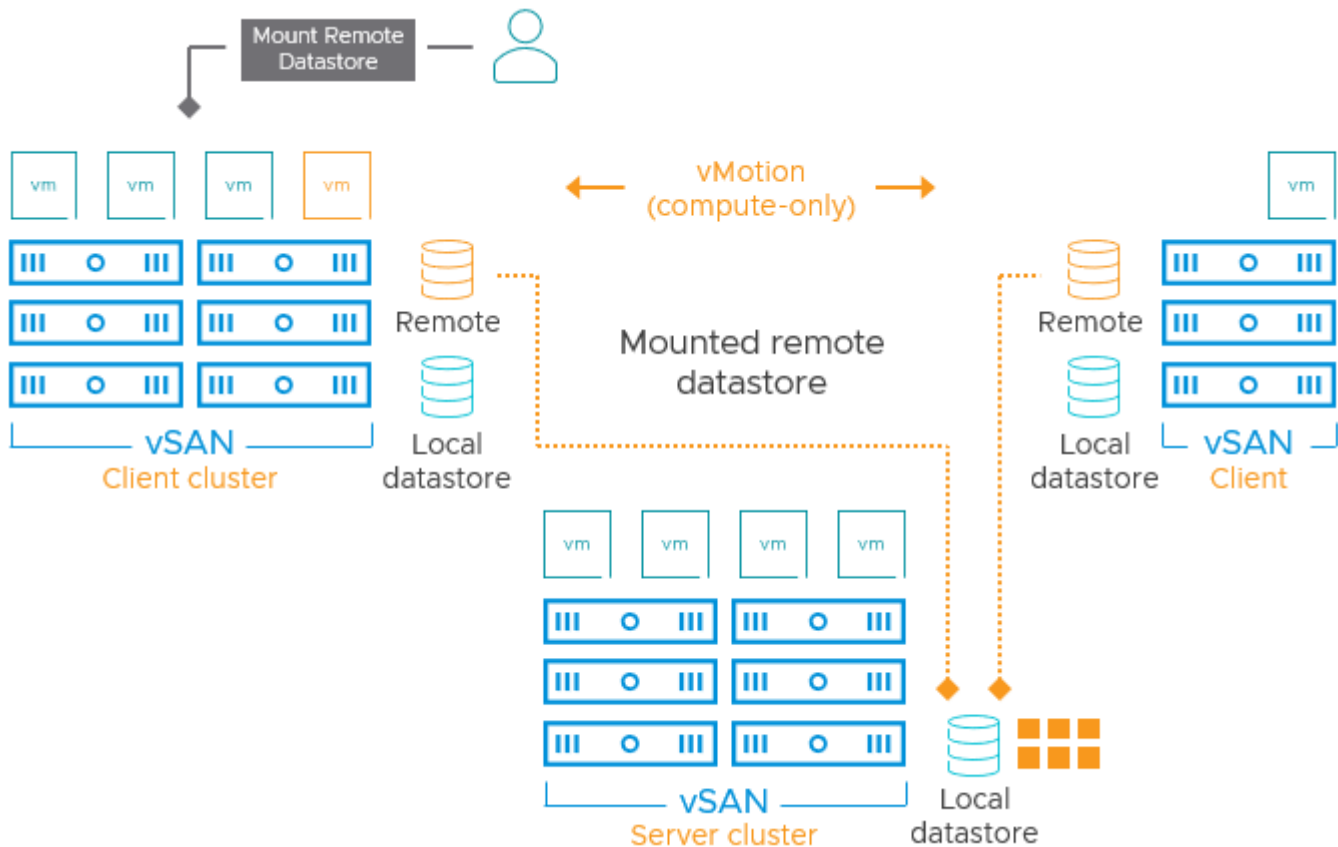


Figure 15. An example of a VMware HCI Mesh topology

How does HCI Mesh change the decision process for large clusters versus small clusters? For the most part, the decision process should remain the same. vSAN has tremendous technical advantages to treating the block storage it serves as a resource of the cluster, and is the default behavior of a standard vSAN cluster. HCI Mesh is a complementary option to cluster designs and topologies that do not meet existing needs. It provides more flexibility in the use of storage resources already deployed in other clusters. HCI Mesh is targeted at addressing a specific need and set of circumstances, and not intended to be a design default. HCI Mesh has inherent **limits and prerequisites** that are covered in the Tech Note described above.

While an HCI Mesh configuration may be borrowing storage capacity from other highly resilient clusters, it does increase a dependency domain beyond that of a typical vSAN cluster. For example, in a traditional vSAN cluster, availability issues are typically limited to the hosts and switching fabric that contributes to the cluster. With an HCI Mesh configuration, availability issues **may include the hosts and the switching fabric from any of the vSAN clusters contributing to the meshed topology**. While HCI Mesh is a flexible and efficient way to consume resources, properly sizing vSAN clusters will be the best way to maintain topologies with minimal dependency domains.

Disaggregation with the Express Storage Architecture

In vSAN 8 U1, the flexibility of disaggregated storage is fully compatible with vSAN Express Storage Architecture (ESA), enabling customers to maximize resource utilization of next-generation devices by sharing storage across non-stretch clusters. Users can mount remote vSAN datastores living in other vSAN clusters and use a vSAN ESA cluster as the external storage resource for a vSphere cluster.

Consume Storage Externally using vSAN Stretched Clusters

In vSAN 8 U1, we introduce support for storage disaggregation when using vSAN stretched clusters powered by the vSAN Original Storage Architecture (OSA). Users can now scale storage and compute independently across non-stretch and stretched clusters.

Consume vSAN Storage across vCenter Server Instances

vSAN 8 U1 also supports disaggregated storage across environments using multiple vCenter Servers when using the vSAN OSA. This enhancement permits customers to use vSAN datastores across different vCenter environments.

Example #1 - Organization with 12 hosts

In this example, a software development company has 12 vSAN hosts. Upon an evaluation of workload demand, budget allocation by the user and the number of hosts available, it is determined that two clusters may offer a better use of resources.

- **Software Development Cluster.** The extremely resource intensive development-based systems supporting CI/CD workflows will use 8 hosts. No deduplication and compression will be used, and VMs in this cluster will largely use RAID-1 based storage policies to optimize performance. Higher performing storage devices can be introduced to just this cluster which would maintain uniformity across the cluster: A recommended (but not required) practice.
- **Multi-Department Cluster.** The general-purpose cluster will support a mixture of workloads and will consist of 4 hosts. In this cluster, deduplication and compression will be used, and VMs will use RAID-1 based storage policies. While the cluster would support VMs using a RAID-5 based erasure coding, using RAID-1 policies allows for the 4-host cluster to have one additional host for automatic rebuilds in the event of a full host failure.

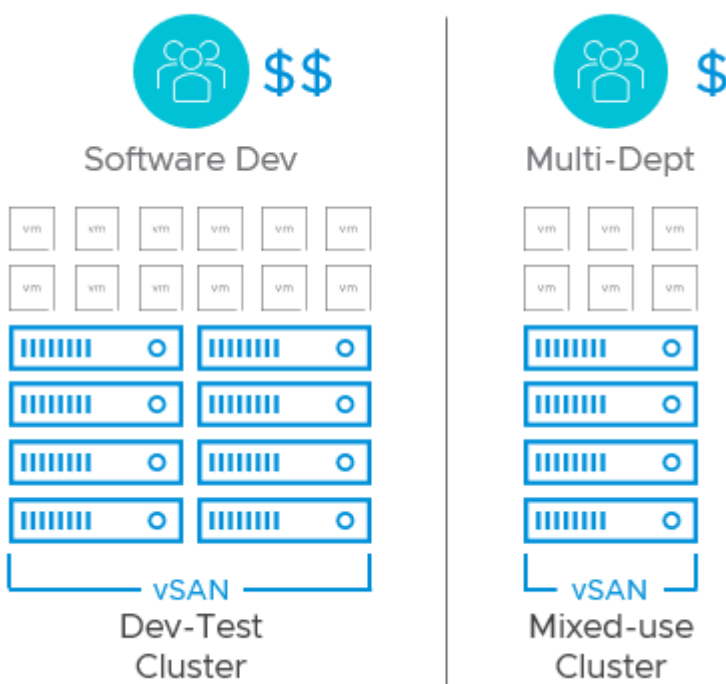


Figure 16. Two clusters created from 12 hosts

The two-cluster topology would achieve the following:

- More predictable performance outcomes for both development, and non-development workloads
- Clusters are large enough to apply required data services
- Prescriptive upgrades to Development cluster would enable faster builds
- Isolation of Dev/Test cluster would limit capacity full conditions to within the cluster
- Easily apply cost accounting benefits to the software development group
- Easily adjust cluster sizing later, if needed
- Reduce the management domain size to a more granular level for easier cluster upgrades in the cluster (firmware, drivers, hypervisor)

The two-cluster approach of 7 hosts + 5 hosts could also be used, depending on resources. This would allow VMs in the general-purpose cluster to use storage policies that use RAID-5 erasure coding for even greater space efficiency.

Example #2 – Organization with 16 hosts

In this example, an organization running business-critical analytics VMs has 16 vSAN hosts. Based on the organizational model of the environment and the number of hosts available, it is determined that three clusters will make for the best use of resources.

Mixed-use Cluster. The mixed-use (general purpose) cluster will carry the bulk of the general production workloads. At 8 hosts, the cluster is sufficiently large enough to protect VMs using storage policies that an FTT level of 3, as well as RAID-5 and RAID-6 erasure coding options.

Analytics Cluster. At just 4 hosts, this cluster is intended to keep resource intensive analytics VMs isolated from general infrastructure. VMs would use RAID-1 mirroring with an FTT level of 1, allowing for the additional host to perform automatic rebuilds to regain compliance in the event of a host failure. Isolating these workloads not only helps contain the I/O traffic but potentially improve communication speeds of the analytics VMs as they communicate with each other, as they would be more likely to live on the same hosts.

Management Cluster. This 4-host management cluster is responsible for all infrastructure services such as vCenter, DNS, alerting, and log aggregation. In this cluster, deduplication and compression will be used, and VMs will use RAID-1 based storage policies. While the cluster would support VMs using a RAID-5 based erasure coding, using RAID-1 policies allows for the 4-host cluster to have one additional host for automatic rebuilds in the event of a full host failure.

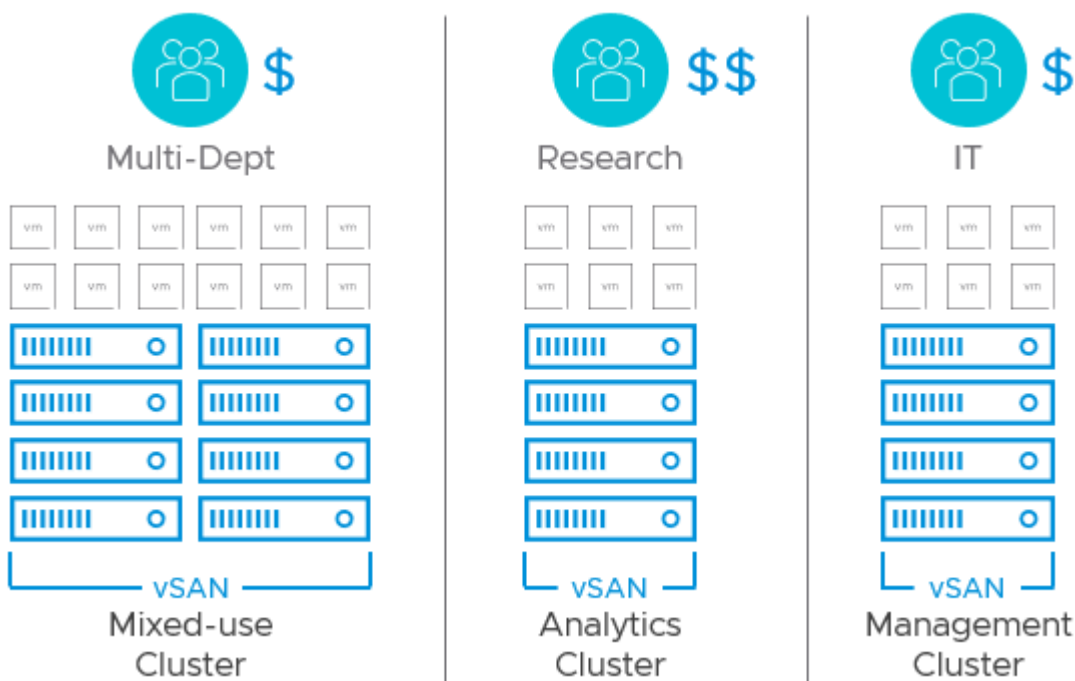


Figure 17. Three clusters created from 16 hosts

The three cluster topology would achieve the following:

- Provides a mixed-use cluster for high efficiency, various levels of protection, and easy growth
- Isolate busy analytics cluster from other resources
- Allows for high-end hardware to be used in resource-intensive analytics cluster
- Isolates management and infrastructure VMs for effective administration, and resilient operations
- Reduce the management domain size to a more granular level for easier cluster upgrades in the cluster (firmware, drivers, hypervisor)

Additional hosts could be added to each cluster as demands increase, requirements in resilience increase or targeted optimizations are needed.

Example #3 - Large Organization with 128 hosts

In this example, a larger organization with a single location has a broad mix of workloads spanning from business-critical ERP systems to B2B applications, to VDI. Knowing that a standard vSAN cluster can be up to 64 hosts, the organization could create two clusters at their maximum size. Given the overall host count, and the needs of this organization, a more prescriptive approach may look like the following:

- **Mixed-use Cluster(s).** A quantity of 5 clusters consisting of 20 hosts per cluster is the initial standard size that would recognize the various needs of the organization, such as business-critical apps, and VDI, as well as target projects by various departments. Cluster-based services could be turned on where needed.
- **Secure Mixed-use Cluster.** A separate 20 host cluster is running vSAN data-at-rest encryption could be used for data that must meet specific security requirements.
- **Management Cluster.** An 8-host management cluster is used to house all infrastructure related VMs, for proper isolation and more effective management.

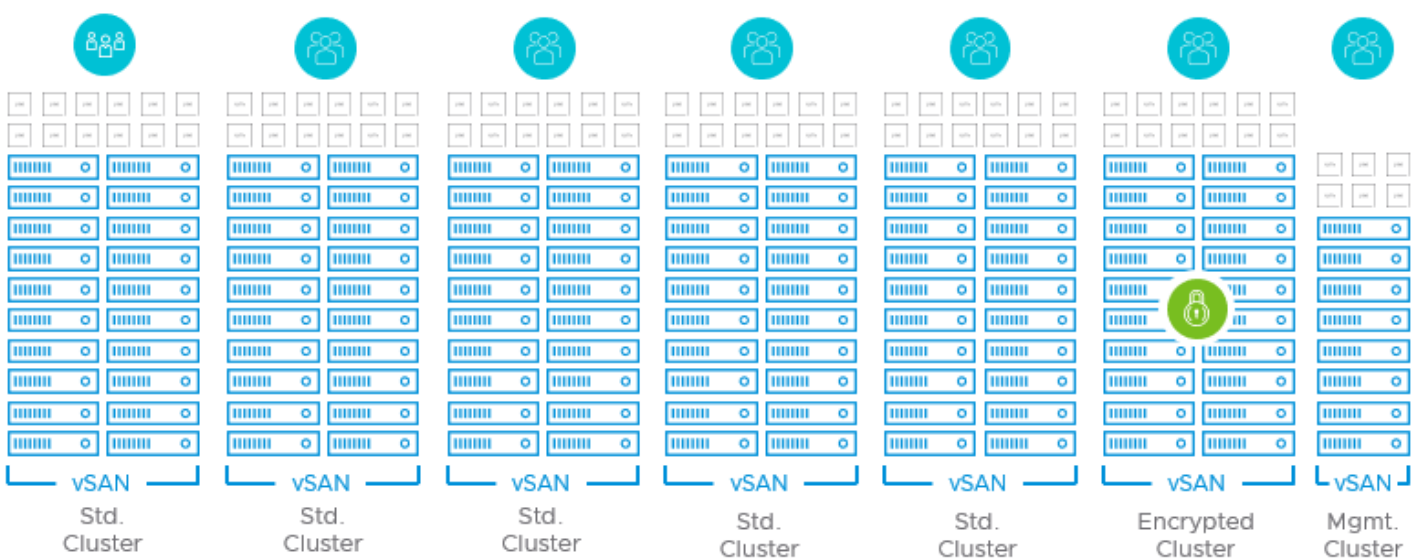


Figure 18. Seven clusters of various sized created from 128 hosts

The seven-cluster topology would achieve the following:

- Clusters small enough to make maintenance domains easier and more agile to manage
- Clusters small enough to target specific needs of the business
- Clusters small enough to contain impacts of potential noisy neighbor scenarios
- Clusters small enough to isolate events such as cluster "full" conditions
- Clusters large enough to take advantage of all data services and protection capabilities of vSAN
- Clusters large enough to have good N+1 efficiency ratios

In this example, the suggestion of a 20-host cluster size would provide clusters large enough to take advantage of all data services and protection capabilities of vSAN, but small enough to reduce impacts of noisy neighbors, maintenance domains, as well as the better address of the needs of the business. In this example, a cluster size of 20 hosts would not be a strict requirement but simply serve as a starting point for a cluster design.

Example #4 - Organization with 16 hosts across two sites

Cluster host count decisions can impact different topologies as well. This example represents an environment that has a DR/CoLo facility used for replication purposes. In a traditional model using array-based replication, opportunities to use the CoLo facility for purposes other than a replication target were somewhat limited, due to the inherent constraints of array-based replication. With vSAN, stand-alone vSAN clusters can be created and used for independent purposes such as a web farm, or some multi-tenant solution. The CoLo facility can be used for other initiatives without compromising the resources of the vSAN cluster acting as the replication target.

- **Mixed-use Cluster.** The mixed-use (general purpose) cluster will carry the bulk of the general production workloads and is what will be replicated to the DR/CoLo site. At 8 hosts, the cluster is sufficiently large enough to protect VMs using storage policies that an FTT level of 3, as well as RAID-5 and RAID-6 erasure coding options.
- **Replication Target Cluster (CoLo).** At just 4 hosts, this cluster is intended to serve as the target for select B2B VMs. The size of the replication target cluster can be smaller than what exists at the primary site, since there are a limited number of VMs replicated, and the user chose to run them with a lower level of FTT than at the primary site.
- **Secure Web Farm/Portal (CoLo).** This 4-host management cluster is responsible for all infrastructure services such as vCenter, DNS, alerting, and log aggregation. In this cluster, deduplication and compression will be used, and VMs will use RAID-1 based storage policies. While the cluster would support VMs using a RAID-5 based erasure coding, using RAID-1 policies allows for the 4-host cluster to have one additional host for automatic rebuilds in the event of a full host failure.

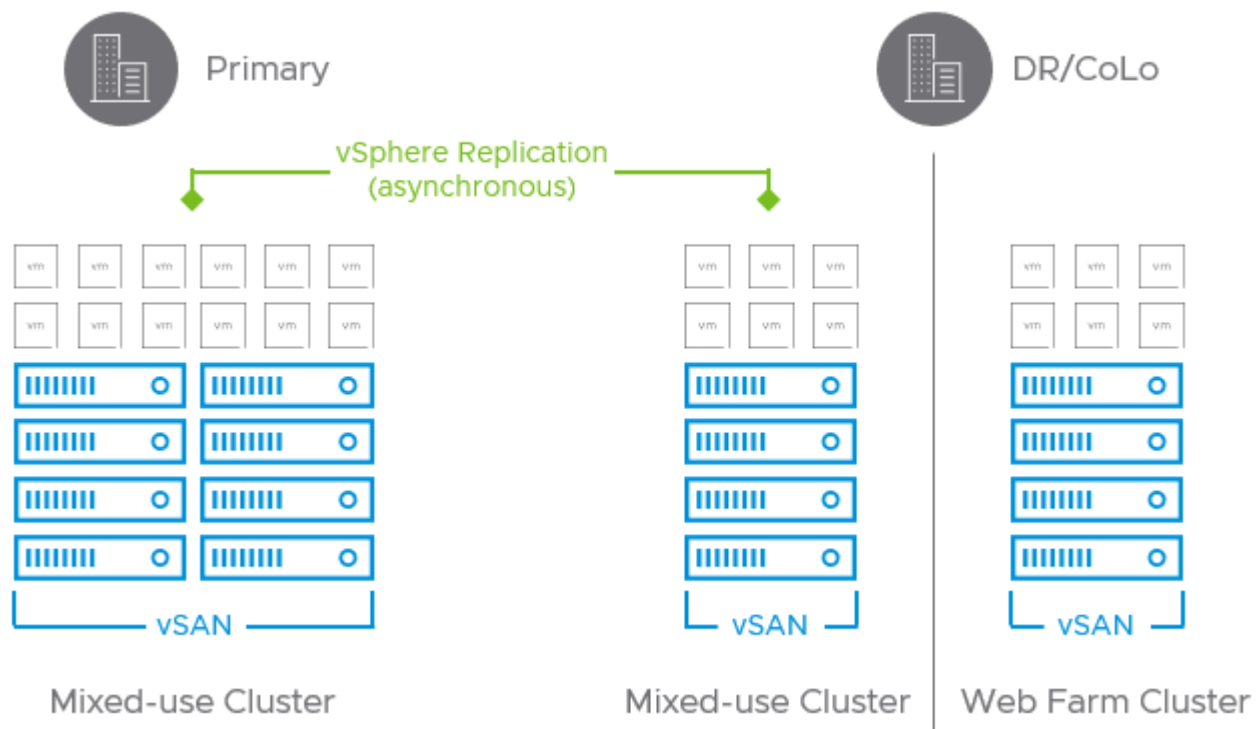


Figure 19. Dedicated cluster created in CoLo for independent purposes

This cluster topology would achieve the following:

- Improve cost-effectiveness of CoLo site (beyond DR)
- Provide services in DR site isolated from cluster used as the replication target cluster
- Take advantage of vSAN cluster related services for these independent clusters
- Optimize sizing requirements for targeted needs and scale more accurately as needs grow
- Provide a mixed-use cluster

Much like the other examples, additional hosts could be added to each cluster as demands or requirements change.

Takeaways

vSAN offers flexibility for your business in ways that were previously unachievable given normal cost constraints. Now the power of choice is given 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.

About the Author

This documentation was written using content from various resources in the Cloud Platform Business Unit, including vSAN Engineering, and vSAN Product Management.

Pete Koehler is a Sr. Technical Marketing Architect, working in the Cloud Platform Business Unit at VMware. He specializes in enterprise architectures, data center analytics, software-defined storage, and hyperconverged Infrastructures. Pete provides more insight into the challenges of the data center at blogs.vmware.com/virtualblocks and can also be found on twitter at [@vmpete](https://twitter.com/vmpete).

