

# VMware Virtual SAN

## Hyper-Converged Management Cluster


### Hyper-Converged Infrastructure for Management Clusters

A VMware vSphere® management cluster is a dedicated group of hosts reserved for running virtual machines that provide management services to infrastructure environments. These services might include, but are not limited to, directory services, DNS, DHCP, VMware vCenter Server™, VMware vRealize® Operations™, and VMware Site Recovery Manager™.

A management cluster separated from other vSphere clusters helps guarantee sufficient compute resources are always available to these critical infrastructure services for improved reliability and performance. In the event of a site failure and subsequent recovery, it is typically easier and faster to bring the foundational infrastructure services back online since they are located in a smaller, dedicated cluster.

In addition to dedicating compute resources, separate storage should be provided to the management cluster for the same reasons we dedicate compute resources. In addition to performance and reliability, having storage dedicated to management resources makes it easier to architect proper data protection and disaster recovery for the workloads in the management cluster.

Compute and storage resources for a management cluster are ideally simple to implement, cost-effective, and provide the performance and availability features necessary to protect critical infrastructure workloads. vSphere and VMware Virtual SAN™ provide the Hyper-Converged Infrastructure (HCI) best suited to meet these requirements.



### Why Virtual SAN for a management cluster?

Virtual SAN is VMware's software-defined storage solution for HCI. Virtual SAN is natively integrated with vSphere and it is configured with just a few mouse clicks. Since disks internal to the vSphere hosts are used to create a Virtual SAN datastore, there is no dependency on external shared storage. Virtual machines can be assigned specific storage policies based on the needs of the management application. Management workloads benefit from dependable storage with predictable performance characteristics.


Virtual machine-centric storage policies can be created and assigned for various workload types. Policies are based upon the available services provided by the storage platform. These policies can be modified and reassigned, as needed, with no downtime.

A dedicated HCI with vSphere and Virtual SAN help ensure predictable performance and availability of management applications and services. Running management workloads in a separate compute and storage facilitates more flexible maintenance schedules. If maintenance is required in the management cluster, there is typically no need to coordinate planned downtime across all of the business applications running on other storage platforms.

Virtual SAN includes a health check plugin for the vSphere Web Client, which automatically monitors and alerts on items such as overall disk health, hardware compatibility list (HCL) compliance, network connectivity issues, and utilization.

## Virtual SAN with vSphere Availability

Virtual SAN features a storage policy rule called “Number of failures to tolerate” or “FTT”, which defines the number of copies of a virtual machine component to distribute across hosts in the cluster. For example, if a policy contains the rule FTT=1 and this policy is assigned to a virtual disk, two copies of the virtual disk are maintained and placed on separate hosts. If one host is offline, the virtual machine disk is still accessible on another host. In other words, that virtual disk can tolerate one host or disk failure and still be available.



Virtual SAN is tightly integrated with VMware vSphere High Availability™. If a host fails, virtual machines that were running on that host are automatically rebooted on other hosts in the cluster to minimize downtime. vSphere HA can also monitor guest operating systems and automatically reboot a virtual machine in the event of a failure such as a Windows blue screen.

vSphere Fault Tolerance™ is also compatible with Virtual SAN and provides continuous availability for applications with up to four virtual CPUs in the event of a host failure.

A variety of data protection solutions are available to back up virtual machines and applications in a Virtual SAN cluster including VMware vSphere Data Protection™. VMware vSphere Replication™ works seamlessly with Virtual SAN to enable rapid recovery with per-virtual machine RPOs as low as five minutes.

## Virtual SAN Performance

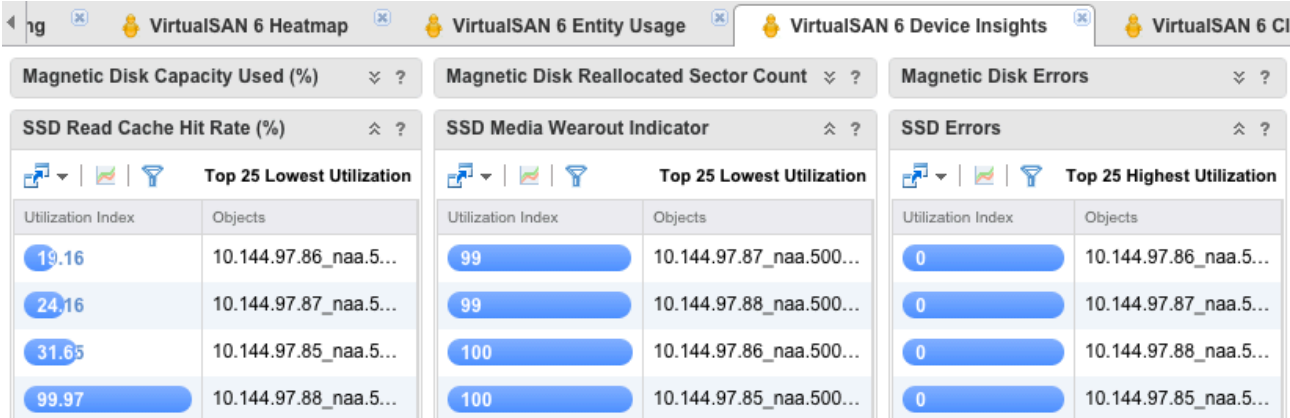
Virtual SAN is uniquely embedded in the vSphere hypervisor kernel and sits directly in the I/O data path. Virtual SAN is able to deliver the highest levels of performance without taxing the CPU or consuming high amounts of memory resources, as compared to other storage virtual machine appliances that run separately on top of the hypervisor. A combination of magnetic and solid state disks is used to enable flash-accelerated hybrid architectures. An all-flash configuration is also possible to deliver even higher levels of performance.

Specific rules such as “Number of disk stripes per object” and “Flash read cache reservation (%)” can be used to further accelerate read-intensive management workloads such as directory services and the vCenter Server database. With Virtual SAN, it is possible to apply policies with precision. For example, database servers are commonly deployed with the guest OS on one virtual disk and databases on other virtual disks. A storage policy that reserves a higher percentage of flash read cache could be assigned specifically to the virtual disks containing databases to help guarantee performance.

The screenshot shows the vSphere storage policy configuration interface for VSAN. The 'Rules based on data services' section is expanded to show three rules: 'Number of failures to tolerate', 'Number of disk stripes per object', and 'Flash read cache reservation (%)'. A help window is open for the 'Flash read cache reservation (%)' rule, providing the following information: 'Flash capacity reserved as read cache for the storage object. Specified as a percentage of the logical size of the object. To be used only for addressing read performance issues. Reserved flash capacity cannot be used by other objects. Unreserved flash is shared fairly among all objects. Default value: 0%, Maximum value: 100%.'

## Visibility and Proactive Notifications with vRealize Operations

vRealize Operations features predictive analytics and smart alerts to help ensure optimum performance and availability of applications and infrastructures. The vRealize Operations Management Pack for Storage Devices includes Virtual SAN dashboards enabling administrators to monitor several factors such as throughput, latency, capacity, and cache hits. Device/hardware-specific reports and alerts are provided using SMARTS.



If there is an issue in the environment, vRealize Operations makes it easy to review symptoms and recommendations for remediation. The vRealize Operations Management Pack for Storage Devices can also be used to monitor, alert, and remediate issues for other storage types such as NFS and FC.

## Reduced Cost and Complexity

Virtual SAN simplifies managing and automating storage for management workloads by eliminating traditional, purpose-built storage systems and by letting IT use familiar vSphere Web Client tools rather than proprietary storage management interfaces. Virtual SAN integrates storage policies into the virtual machine creation workflow, ensuring each workload has the type of storage it needs. There are no additional virtual appliances to install and you can enable Virtual SAN with just a few clicks.

Virtual SAN is a distributed architecture that allows for elastic, non-disruptive scaling. Capacity and performance can be scaled at the same time by adding a new host to the cluster – scale out – or capacity and performance can be scaled independently by adding new drives to existing hosts – scale-up, adding SSD for performance or HDD for capacity. This “grow-as-you-go” model provides predictable, linear scaling with affordable investments spread out over time.

In hybrid deployments with solid state and magnetic disks, Virtual SAN delivers top-end performance out of industry-standard hardware by using SSDs for read caching/write buffering and HDDs for data persistence, balancing performance and cost requirements. If higher performance is needed, an all-flash Virtual SAN architecture can deliver up to 90K IOPS per host with consistent sub-millisecond response times.

## Summary

Virtual SAN and vSphere are the ideal HCI platform for running management workloads where predictable performance and high availability are strict requirements. Virtual machines running important applications such as directory services, DNS, and vCenter Server do not have to compete for compute and storage resources shared by business workloads. Virtual machine-centric storage policies can be created, assigned, and modified, as needs change in the environment without incurring downtime. Maintenance windows are easier to schedule and there are features such as vSphere HA and vSphere Replication to enable rapid recovery from unplanned downtime. Virtual SAN health monitoring is included and, optionally, vRealize Operations Management Pack for Storage Devices provides multiple Virtual SAN dashboards for proactive alerting, heat maps, device and cluster insights, and streamlined issue resolution.

## Learn More

[Virtual SAN Product Page](#)

[vRealize Operations Product Page](#)

[Virtual Blocks Blog](#)

[Customer Stories](#)

