VMware vCloud® Architecture Toolkit™
for Service Providers

Introduction to VMware vSAN™ for VMware Cloud Providers™

Version 2.9

January 2018

Derek Mitchell

VMware, Inc.
3401 Hillview Ave
Palo Alto, CA 94304
www.vmware.com

Contents

Overview 5
Target Audience 6
Use Cases 6
3.1 Management Cluster 6
3.2 Tiered Storage 7
3.3 Hosted vSphere 8
3.4 Hosted Private Cloud 8
3.5 Public Cloud 9
Introduction to vSAN 11
4.1 vSAN Configurations 11
4.2 vSAN Availability 12
Deploying vSAN 13
vSAN Solutions Design 15
6.1 vSAN Cluster Design 15
6.2 vSAN Networking Design 16
6.3 vSAN Storage Design 16
6.4 vSAN Storage Policy Design 17
vSAN and Recoverability 18
7.1 vSAN Stretched Cluster 18
7.2 VMware vSphere Data Protection Advanced 20
Managing vSAN 21
8.1 Management and Operations 21
Conclusion 22
Assumptions and Caveats 23
Reference Documents 24

List of Figures

Figure 1. vSAN Based Management Cluster 7
Figure 2. vRealize Automation with vSAN Logical Architecture 9
Figure 3. vCloud Director with vSAN Logical Architecture 10
Figure 4. vSAN Cluster Example 11
Figure 5. Rainpole Cluster Layout 14
Figure 6. VMware Cloud Provider Program Stretched Cluster Example 19
Figure 7. Network Connectivity for Stretched Cluster 20

Overview

A software-defined storage infrastructure can be a critical piece of the foundation of a hybrid cloud service offering in which infrastructure is provided by a combination of private cloud (on premises) and public cloud (cloud service provider) resources. There are many different workload profiles in the data center, but from the perspective of storage, workloads have traditionally been separated into two profile types: file-based and block-based. While these technologies have had a long track record of success in the enterprise space, their track record has been somewhat limited in the cloud space.

VMware Cloud Providers™ require that their storage infrastructure not only be reliable and scalable, but also cost effective. The consumers of this storage expect that this cost efficiency will be passed on to them. More often than not, these solutions are designed leveraging commodity servers in conjunction with internal storage capabilities.

There are several approaches for architecting a VMware software-designed storage solution and the approach will depend on the use case and the technology deployed. However, the end goal is always the same. This *VMware vCloud Architecture Toolkit for Service Providers* (vCAT-SP) solution architecture document describes a VMware Cloud Provider Program solution for supporting a software-defined storage solution using VMware vSAN™.

Target Audience

This document is targeted towards architects, engineers, application owners, and technology leaders involved in the key decision making process and anyone else interested in guidance on designing a software-designed storage solution leveraging VMware technologies.

Use Cases

Cloud service offerings can utilize either traditional infrastructure (servers with legacy storage), or hyper-converged platforms such as vSAN. A hyper-converged platform is a platform that combines compute, storage, networking, and virtualization resources and management, enabling scale-out consumption. The tight integration of vSAN with VMware vSphere® enables VMware Cloud Providers to host public or private cloud offerings and easily utilize vSAN as the underlying storage infrastructure.

The customers of service providers that leverage vSAN primarily run mission-critical workloads in which stability, reliability, and predictability are extremely important. Many traditional infrastructures, over time, hinder VMware Cloud Providers' ability to remain competitive in such a price-sensitive market. vSAN enables VMware Cloud Providers launching new offerings to avoid the large capital expenditures typically associated with legacy storage arrays.

VMware Cloud Providers who offer consumers hybrid cloud-based virtualized solutions, such as vSphere as a Service, Desktop as a Service, Tier 1 applications, and Database as a Service, are ideal candidates for vSAN. Additionally, VMware Cloud Providers can leverage vSAN as the storage infrastructure for their VMware vCloud Director® cloud offering.

3.1 Management Cluster

Traditionally, providing dedicated storage for management-only clusters increased providers' infrastructure costs. It required the purchase of an additional storage array to provide the expected performance and availability for what is typically a high I/O performance and highly-available virtual environment. These costs, in turn, were inherently passed on to the consumer. With vSAN, these costs are significantly reduced, making it an ideal solution for a dedicated management environment. Additionally, leveraging vSAN for the management cluster eliminates single points of failure and is completely integrated with the VMware vCenter Server® system. See the following figure for a logical diagram of this architecture.

Figure 1. vSAN Based Management Cluster

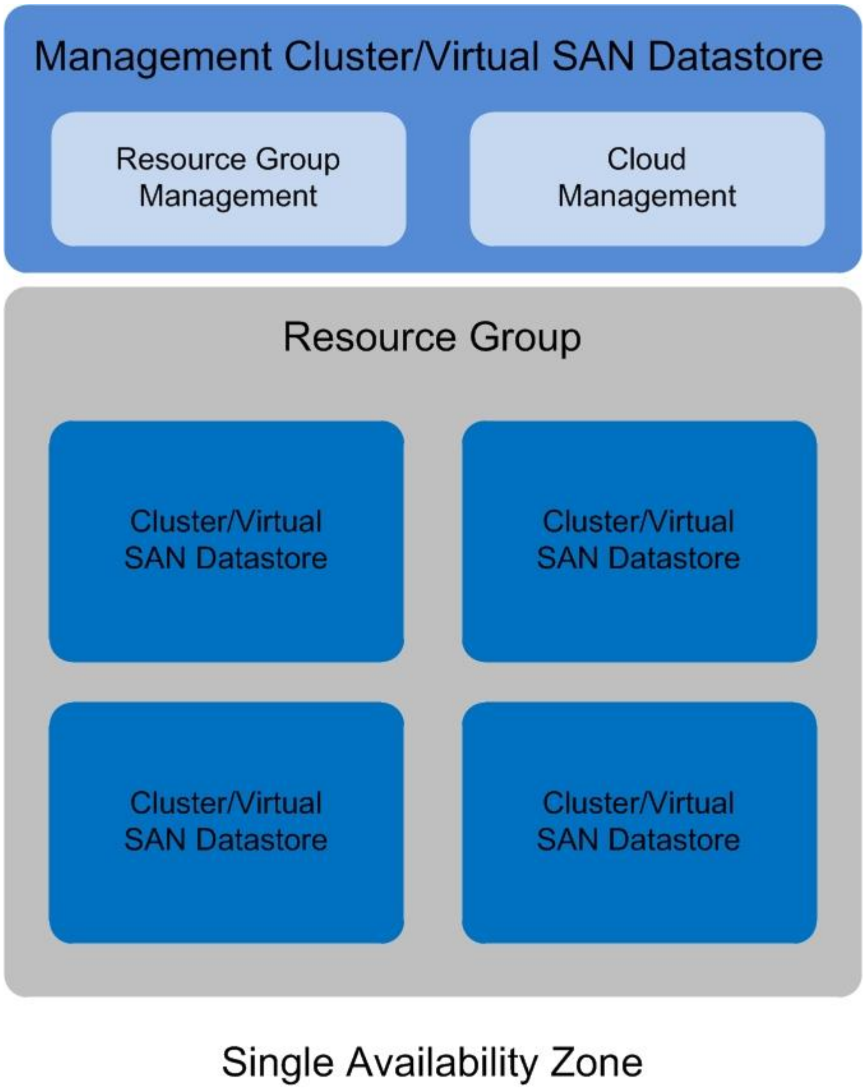

3.2 Tiered Storage

Tiered storage is probably one of the most effective uses of vSAN in a cloud environment. Many service providers deploy multiple tiers of offerings, for example, gold, silver, and bronze. Each provider virtual data center / resource cluster must be configured with a vSAN datastore that meets the specific capability requirements set out by the Service Level Agreement (SLA) for that tier of service.

For example, the number of acceptable failures allowed, cache/capacity ratio, and number of disk stripes per object for the gold offering can be higher than that of the silver or bronze offerings. The combination of storage policy, disk group configuration, and physical hardware can be used to efficiently and easily segment consumer storage resources into pools that align with the provider's SLAs. See the *vCloud Director with Virtual SAN Sample Use Case* article at https://blogs.vmware.com/vcat/2015/09/vcloud-director-and-virtual-san-sample-use-case.html for a detailed description of this use case.

3.3 Hosted vSphere

For customers who want single-tenant, dedicated compute, network, and storage, a hosted vSphere infrastructure is the ideal solution. In this scenario, the entire infrastructure, physical and virtual, is owned and managed by the service provider. Consumers are able to access this infrastructure through a service portal where they have the ability to provision and manage virtual machines, on demand, through an intuitive graphical user interface.

As discussed earlier, vSAN is a hyper-converged offering which makes the hosted vSphere environment an ideal use case. First, it enables providers to stand up an infrastructure in a timely and predictable manner. There is no need to add external arrays or configure LUNs, so the time to deployment is greatly reduced and the process is highly repeatable. Second, vSAN is supported on all major server OEM platforms. Most service providers have a preferred server vendor and receive favorable pricing. Because vSAN is likely supported on their preferred platform, the existing procurement vehicles do not need to be altered.

In summary, service providers can leverage their current VMware skill sets, reduced time to deployment, lower CapEX/OpEX, and the ability to use existing processes and procedures for procuring equipment. End clients get increased agility, reduced risk, and reduced CapEx/OpEx. These benefits make vSAN the best hyper-converged offering for hosted vSphere environments.

3.4 Hosted Private Cloud

VMware vRealize® Automation™ gives providers another option for offering a hosted single-tenant virtualized solution to customers. A hosted vRealize Automation solution allows customers to manage their own resources using custom provisioning workflows, manage virtual machines, create snapshots, and distribute specific resources into resource pools and/or catalogs as needed. With a host vRealize Automation solution, customers can benefit from the agility of a private cloud management platform, without the overhead that is required to deploy and manage the private cloud management infrastructure.

vRealize Automation can be used to provision virtual machines and/or applications onto a vSAN datastore while leveraging the storage-based policy management capabilities of vSAN. This is extremely valuable because storage policies can be configured on a per-application or per-virtual machine basis, allowing providers to assign policies based on a particular use case, such as tiered storage.

When implementing vSAN in a hosted Private Cloud model, vSAN will be managed in a way similar to how a customer would experience an on-premises implementation of VMware vSphere with vSAN. The VMware vSphere Web Client will provide the primary interface for any required maintenance and configuration of vSAN. The hosting solution can either be managed by the service provider or self-managed by the consumer.

Figure 2. vRealize Automation with vSAN Logical Architecture

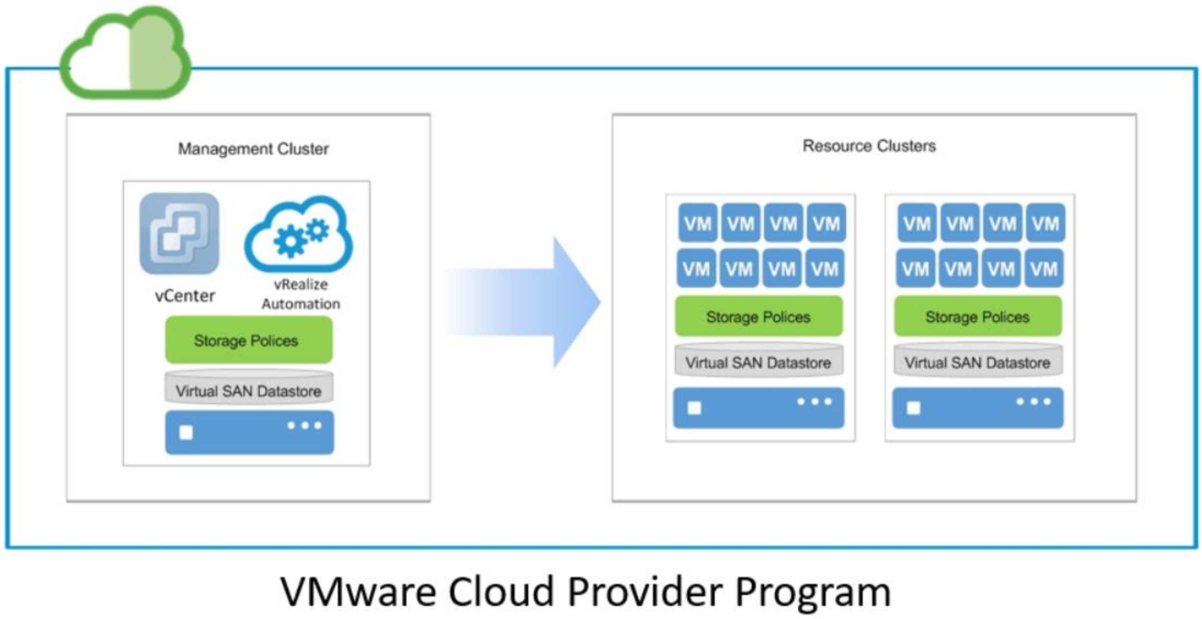

3.5 Public Cloud

When leveraging vSAN with a vCloud Director based cloud infrastructure, the typical service provider architecture consists of two clusters based on vSAN. The first is a management cluster, which hosts all components needed for a vCloud Director environment, in conjunction with a resource cluster. This enables the provider to start out with a relatively low investment in hardware and quickly scale as new customers are added or when existing customers require new hardware in the cloud environment.

This design is in alignment with typical cloud architecture in which the management components are deployed in the management cluster and tenant workloads are hosted in the resource cluster. It enables the provider to offer different SLAs for management components and tenant workloads, provides separation of duties, and allows both clusters to easily scale by adding hosts where needed. The following figure shows this architecture.

Figure 3. vCloud Director with vSAN Logical Architecture

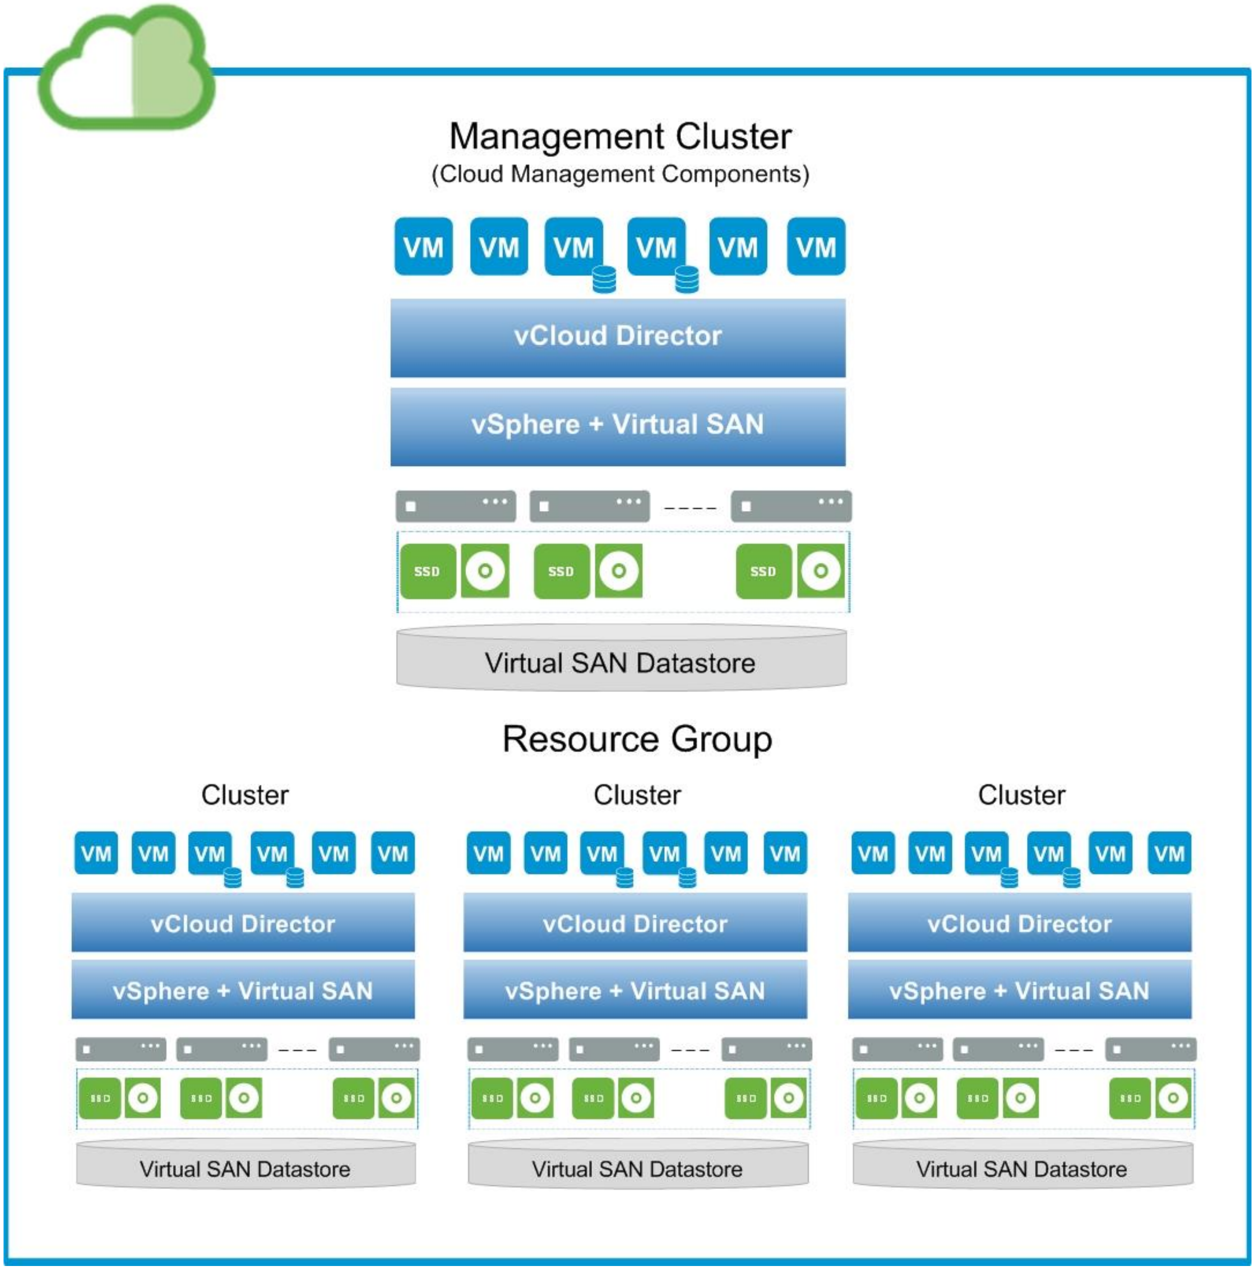

Introduction to vSAN

Object-based storage is considered a leading technology for hybrid cloud deployments because many of its most prominent features, such as massive scalability, geographic independence, and multi-tenancy, have proven ideal for cloud storage.

vSAN is a hyper-converged, software-defined storage solution that stores and manages data in the form of storage objects. Similar to vSphere, vSAN can leverage commodity physical servers and locally-attached flash devices and/or hard disks. vSAN is also completely integrated into the VMware stack, so it works seamlessly with advanced availability features such as VMware vSphere vMotion®, VMware vSphere High Availability (HA), VMware vSphere Distributed Resource Scheduler™ (DRS), and VMware vSphere Fault Tolerance (FT).

4.1 vSAN Configurations

vSAN is available in two configurations: hybrid and all-flash.

See the following figure for a sample vSAN cluster.

Figure 4. vSAN Cluster Example

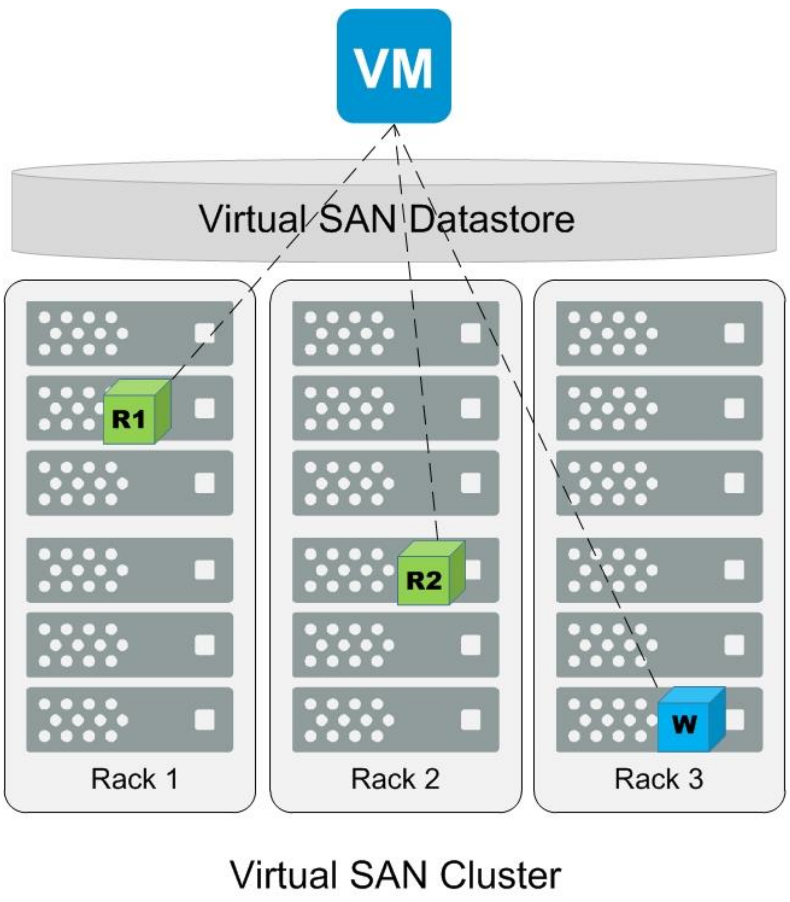

4.1.1 Hybrid Configuration

The hybrid configuration for vSAN uses magnetic disks for the capacity layer and flash memory devices for the cache layer. This option is ideal for workloads that have average storage performance requirements.

4.1.2 All-Flash Configuration

The all-flash configuration for vSAN uses flash devices for both capacity and caching. This option is available for workloads with high storage performance requirements.

4.2 vSAN Availability

4.2.1 Fault Domains

The previous example demonstrates the vSAN concept known as a fault domain. This feature allows providers to group multiple hosts into failure zones across multiple server racks so that replicas of virtual machine objects are not provisioned onto the same logical failure zones or server racks.

4.2.2 Failures to Tolerate Configuration

In this example, the "number of failures to tolerate" configuration is set to 1 in the virtual machine storage policy so that there is another replica of the VMDK created in the capacity layer of another fault domain, providing two copies of the data (R1 and R2 respectively as shown in the previous figure). If the number of failures to tolerate is set to 2, there is an additional replica of the VMDK maintained across the cluster for a total of three copies of the data. The virtual machine's components are distributed through multiple hosts onto multiple racks. Should a rack failure event occur, the virtual machine that is not hosted on the hosts located where the server rack failed, continues to be available.

When architecting a hybrid cloud storage solution, it is also important to consider the characteristics of the application or service, such as IOPS, capacity, availability, and recovery. While the infrastructure, including storage, is important, applications are the key to success of any organization utilizing cloud services.

Deploying vSAN

vSAN allows for per-VM policy-driven storage provisioning. vSAN provides tight integration with vSphere, allowing VMware Cloud Providers offering a hybrid cloud service to easily utilize vSAN in the storage infrastructure.

The remainder of this section provides an architectural overview of a hybrid cloud model comprised of a VMware Cloud Provider data center providing dedicated hosted resources for their business customer, Rainpole, a hypothetical company used to illustrate potential use case.

Rainpole is preparing to stand up new infrastructure to be used to host a hybrid cloud environment. This particular VMware Cloud Provider is launching a new service based on vSAN, and Rainpole's new infrastructure appears to be a good use case.

This infrastructure will host a new application, therefore, Rainpole does not have any historical data, so they are designing from scratch. Rainpole's infrastructure requirements are as follows:

- 400 virtual machines
- Each VM has a 50-GB VMDK and 8 GB of virtual memory
- Tolerate 1 failure and support 1 snapshot
- 70% read, 30% slack space, and 20% future growth

Using this data along with the vSAN sizing tool, the service provider has calculated the infrastructure required to support this inaugural customer as the following:

- 20 nodes (2x Intel Xeon Ivy Bridge 8C E5-2640V3 2.6 G, 8x 16 GB 1866 MHz DDR3 SDRAM, 1x LSI 3008 based controller, 1x 10G support w/X540 chip with 2x RJ45)
- 80 HDDs (1 TB, 7.2K RPM, NL-SAS, 2.5-inch drives)
- 20 SSDs (200 GB, SATA)

The vSAN resource cluster is configured in the following manner:

- Each node is configured with 1 disk group (1x SSD, 4x HDD) for a total raw capacity of 4 TB and total cache capacity of 200 GB.
- 2 resource clusters of 10 nodes for a total raw capacity of 40 TB and total flash capacity of 2 TB per cluster.
- Max IOPS per node is 8K and 80K per cluster.

The infrastructure is divided into two resource clusters for two reasons. The first is that a smaller cluster size, somewhere between 10 to 15 nodes depending on use case, provides a higher level of availability. Having all of the nodes in a single cluster provides a larger fault domain and increases risk. The second reason is that it provides the ability to scale both up (within nodes) and out (adding nodes to the cluster). In this example, each cluster is spread over three fault domains for enhanced reliability. See the following figure for a diagram of the cluster layout.

Figure 5. Rainpole Cluster Layout

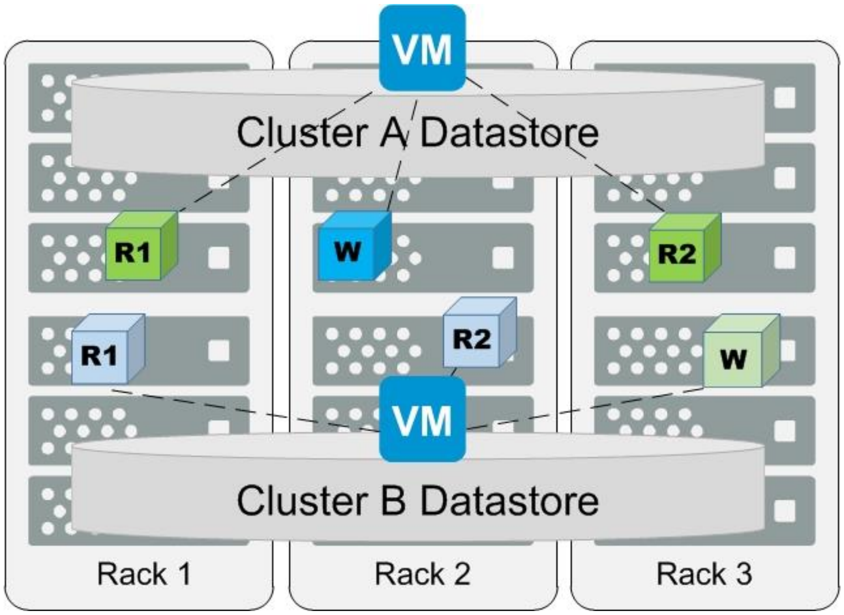

vSAN Solutions Design

Designing storage resources for a cloud model differs from the traditional vSphere approach used in defining storage for non-cloud onsite data center storage. Platform features such as VMware vSphere Storage DRS™ and storage policies assist in balancing workloads across storage resources, allowing providers to offer differentiated storage. This allows the provisioning of flexible pools of storage resources while maintaining consistent performance for consumers. Users can choose the appropriate storage tier for a particular type of workload.

- Historical data is usually the best guide for designing and determining the size and scale of a vSAN infrastructure. For example, if the average virtual machine requirements, in terms of storage, memory, IOPS, and compute can be correlated, this information can be used to efficiently design your vSAN infrastructure.
- Application data is also required in the design. Information, such as whether the application is cache-friendly, requires file services, and has availability requirements, is critical. Once this information is calculated, it must be correlated and aligned to future growth needs and the financial requirements of the provider with respect to operational and capital expenditures as well as profitability goals. The objective is to design the most efficient architecture that will meet current and future demand.

6.1 vSAN Cluster Design

The first and most important design consideration for vSAN is to adhere to the *VMware Compatibility Guide* at http://www.vmware.com/resources/compatibility/search.php?deviceCategory=vsan. When designing a vSAN cluster, use similarly configured and sized hosts. Otherwise, the performance and capacity expectations will be compromised. Most VMware Cloud Providers have a preferred server vendor and, more than likely, they have a preconfigured and validated server (Ready Node) architecture for vSAN. If not, providers can also build their own vSAN nodes using the information provided by the compatibility guide. From an operational standpoint, the Ready Nodes have been the preferred choice of infrastructure for vSAN.

vCloud Director management clusters have different requirements than a resource cluster. Traditionally, providers would allocate dedicated storage for management-only clusters, which would ultimately increase the infrastructure costs. Deployment of a management cluster would normally entail the purchase of an additional storage array to provide the expected performance and availability for what is typically a high I/O throughput and highly available virtual environment. These costs, in turn, would inherently be passed on to the consumer. With vSAN, these costs are significantly reduced making it an ideal solution for a dedicated management environment. Using vSAN for the management cluster provides an entire integrated stack that is automated by the vCenter Server system. The storage is locally attached and backed by SSD-based cache. Take the following best practices into consideration when designing a management cluster backed by vSAN:

- Use a minimum of four nodes with storage
- Use a balanced cluster with identical host configurations
- Do not use a stateless boot image. Use an SD card / USB / SATADOM (Serial ATA Disk on Module)

When designing a vSAN backed vCloud Director infrastructure, VMware recommends not using a vSAN datastore for catalogs. All catalog media images are uploaded as file objects into the same directory structure by vCloud Director. If the same vSAN datastore is used as storage policy for different catalogs, they will share one VM Home Namespace object with maximum size of 256 GB. In this case, use a third-party virtual storage appliance that provides NFS file services to provide catalogs that scale larger than 256 GB.

6.2 vSAN Networking Design

Although VMware supports both 1-Gb and 10-Gb Network Interface Cards (NICs) for vSAN network traffic in hybrid configurations, VMware highly recommends that 10-Gb be used for provider environments. (For all-flash configurations, 10 Gb is required.) Multicast communication is a requirement and one thing to keep in mind is that some 10-Gb switches have better performance than others, even though the feature sets are seemingly identical.

NIC teaming is recommended for availability, but does not provide load balancing. vSAN network traffic has not been designed to load balance across multiple network interfaces when these interfaces are teamed together.

Jumbo frames can reduce CPU utilization and improve throughput, however, both gains are minimal, because vSphere already uses TCP Segmentation Offload (TSO) and Large Receive Offload (LRO) to deliver similar benefits. However, if jumbo frames are already enabled on the network, the recommendation is to move forward with jumbo frames enabled for vSAN deployment.

6.3 vSAN Storage Design

Disk groups can be thought of as storage “containers” on vSAN. They contain a maximum of one flash cache device and up to seven capacity devices: either magnetic disks (hybrid configuration) or flash devices (all-flash configuration) are used as capacity devices. Each disk group assigns a cache device to provide the cache for a given capacity device. The recommendation is to have at least a 10 percent cache-to-capacity ratio. This provides a degree of control over performance as the cache-to-capacity ratio is based on disk group configuration. This also needs to be taken into account when planning future growth. For example, you want to make sure that the flash layer devices are large enough to scale the capacity layer for growth. Otherwise, you will not be able to maintain the minimum flash-to-capacity ratio. Depending on the use case, it might be necessary to design with additional cache up front to allow for future growth of the capacity layer.

To rebuild components after a failure, the design must be sized so that there is a free host’s worth of capacity to tolerate each failure. There must be at least one full host’s worth of capacity free for maintenance. The number of failures to tolerate will determine whether there is a requirement for additional host capacity. For example, to rebuild components after one failure (FTT=1), there must be one full host’s worth of capacity available. To rebuild components after a second failure (FTT=2), there must be two full hosts’ worth of capacity free.

When evaluating hardware for cluster nodes in a hybrid cloud environment, the hardware must be identical, with special attention given to the storage I/O controllers. Queue depth must be as large as possible. At a minimum, the queue depth must be able to accommodate the throughput of current and future devices. In general, SATA drives must the lowest queue depth of the supported magnetic disks, and for this reason, they are not recommended in a cloud environment. Equally important, verify that the storage I/O controller supports pass-through mode. RAID 0 is not recommended in a hybrid cloud environment due to the increased maintenance of setting up and replacing disks. Some additional considerations when designing vSAN storage include the following:

- The number of magnetic disks matter in hybrid configurations due to the eventual de-staging of read cache. Multiple disk spindles can speed up this process. Having more, smaller magnetic disks will often give better performance than fewer, larger disks in hybrid configurations.
- Allow 30 percent slack space when designing capacity.
  - vSAN begins automatic rebalancing when a disk reaches the 80 percent of full threshold.
  - Target configurations must be approximately 10 percent of the 80 percent threshold.
- Multiple disk groups typically provide better performance and smaller fault domains, but might sometimes come at a cost and consume additional disk slots.

6.4 vSAN Storage Policy Design

As discussed earlier, vSAN is an object storage technology. Each virtual machine deployed on vSAN comprises a set of objects. VMDKs, snapshots, VM swap space, and the VM home namespace are all objects. Each of these objects is comprised of a set of components that are determined by capabilities configured in the VM storage policy. For example, if a virtual machine is deployed with a policy to tolerate one failure, objects will be made up of two replica components. If the policy contains a stripe width, the associated object will be striped across multiple devices in the capacity layer. Each stripe is a component of the object.

As you plan the design with respect to component maximums, be aware that vSAN might decide that an object needs to be striped across multiple disks, even with the default policy of one stripe in place. Normally, this is the result of an administrator requesting that a VMDK be created that is too large to fit on a single physical drive. As previously highlighted, the maximum component size is 255 GB. To elaborate, objects that are greater than 255 GB in size will automatically be divided into multiple components. For example, if an administrator deploys a 2 TB VMDK, it will result in eight or more components being created in the same RAID-0 stripe configuration. Finally, all virtual machines created on vSAN will be thin-provisioned, so plan accordingly to avoid overcommitting resources.

vSAN and Recoverability

vSAN enables providers to protect consumer data, even in the event of disk, host, network, or rack failures, with built-in distributed RAID and cache mirroring. Resiliency is enhanced by defining availability on a per-VM basis. The number of host, network, disk, or rack failures to tolerate in a vSAN cluster can be customized based on customer, application, or business requirements.

vSAN supports clustering technologies from both Oracle and Microsoft. With Oracle Real Application Cluster (RAC), customers can run multiple Oracle RDBM instances, accessing the vSAN datastore to deliver better performance, scalability, and resilience. Additionally, Windows Server Failover Clustering (WSFC) can provide protection against application and service failures.

Technically, there is a minimum requirement of three VMware ESXi™ hosts in a vSAN cluster. Although VMware fully supports three-node configurations, providers must be aware of the difference, with respect to configurations, with four or more nodes. In particular, in the event of a failure, there is no way for vSAN to rebuild the components on another host in the cluster to tolerate another failure. There is also an operational limitation because vSAN will not have the ability to migrate all data from a node during maintenance.

Business continuity and disaster recovery are paramount for ensuring that consumers' business critical environment, data, and online presence are available with minimal downtime. vSAN recoverability can be addressed through a range of solutions, from near-line capabilities such as snapshots, offline capabilities provided by VMware vSphere storage integrated backup solutions (using APIs for data protection), or inter-data center solutions, such as stretched cluster.

7.1 vSAN Stretched Cluster

Stretched storage with vSAN allows you to split the vSAN cluster across two sites, so that if a site fails, providers are able to seamlessly fail over to the other site without any loss of data. vSAN in a stretched storage deployment accomplishes this by synchronously mirroring data across the two sites. The failover is initiated by a witness VM that resides in a central location, accessible from both sites. This is a specific implementation for environments where disaster avoidance and unscheduled downtime are key requirements.

A vSAN stretched cluster is a specific deployment where the provider sets up a vSAN cluster with two disparate active/active sites with an identical number of ESXi hosts distributed evenly between the two sites. The witness host resides at a third site and the sites are connected by way of a high-bandwidth, low latency link. The third site, hosting the witness host, is connected to both of the active/active data sites. The sites can be a combination of the VMware Cloud Provider program, customer, and third-party data centers. See the following figure for an example.

Figure 6. VMware Cloud Provider Program Stretched Cluster Example

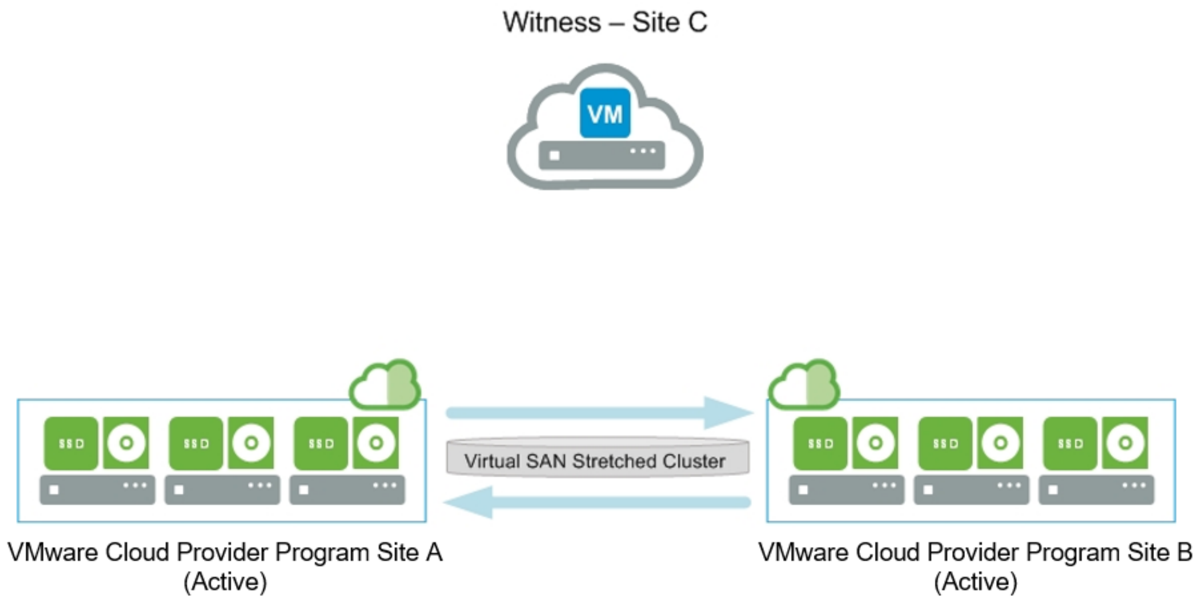

In a vSAN stretched cluster implementation, each site is configured as a vSAN fault domain, and there is only one witness host in any configuration. Each site can be considered a fault domain and a maximum of three sites (two data, one witness) is supported. For deployments that manage multiple stretched clusters, each cluster must have its own unique witness host.

A virtual machine deployed on a vSAN stretched cluster will have one copy of its data on site A, a second copy of its data on site B, while any witness components are placed on the witness host in site C. This configuration is achieved through fault domains and affinity rules. In the event of a complete site failure, there will be a full copy of the virtual machine data available, as well as greater than 50 percent of the components. This will allow the virtual machine to remain available on the vSAN datastore. If the virtual machine needs to be restarted on the other data site, vSphere HA will accommodate this task.

Geographic distance is, in theory, not a concern when designing a vSAN stretched cluster. The key requirement is the latency between the respective sites. VMware requires a maximum latency of no more than 5 ms RTT (Round-Trip Time) between data sites and no more than 200 ms RTT between data sites and the witness host. As long as the latency requirements are met, there is no restriction on geographical distance.

As discussed earlier, vSAN stretched cluster requires three disparate sites, and each site must communicate on the management, vSAN, VM, and vSphere vMotion networks. Detailed networking design is beyond the scope of this document. However, to minimize uncertainty with the implementation, VMware recommends that providers implement a stretched L2 between the data sites, and a L3 configuration between the data sites and the witness site.

In the example illustrated in the following figure, the choice was made to use a virtual witness connected over L3 with static routes. The witness is deployed on a physical ESXi host with two preconfigured networks for the management and vSAN networks, respectively. The data sites are connected by way of stretched L2 which backs the management, vSAN, VM, and vSphere vMotion networks. All hosts in the cluster must be able to successfully communicate, and to facilitate this communication, static routes must be configured (per host) between the data hosts in Site A and B, and the witness host in Site C, for vSAN traffic to flow between the data sites and witness site.

Figure 7. Network Connectivity for Stretched Cluster

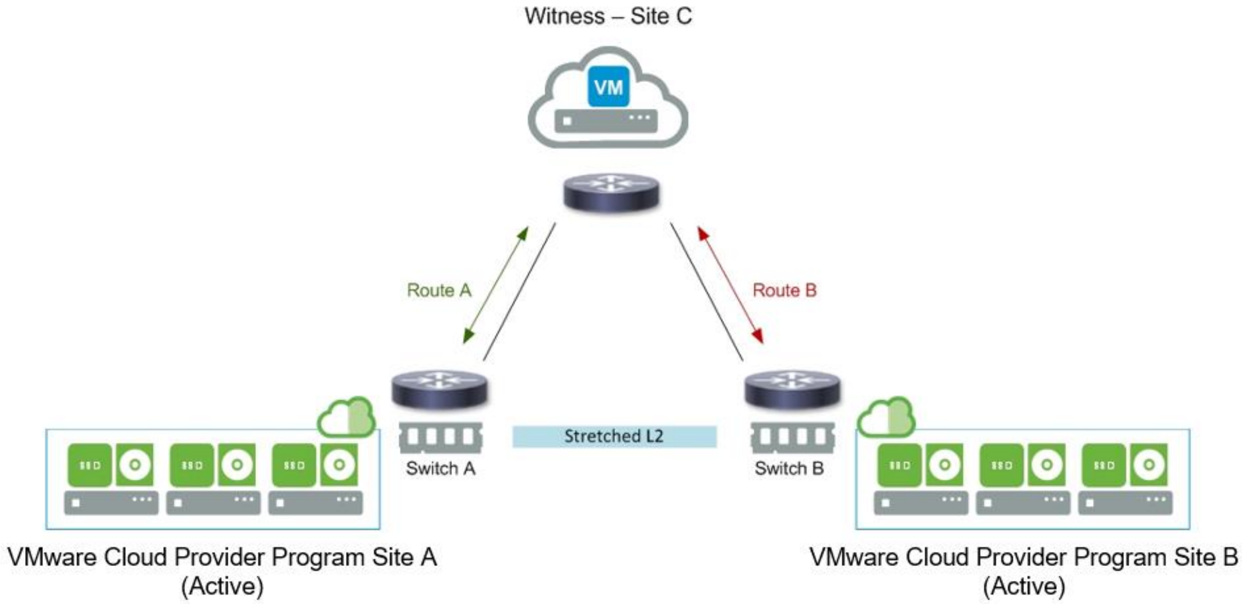

Ultimately, the success and design considerations of a specific vSAN stretched cluster implementation depend upon many factors ranging from choice of topology to the physical capabilities of a providers' networking infrastructure. For detailed guidance on designing a vSAN stretch cluster, consult the *VMware Virtual SAN Stretched Cluster Guide* at http://www.vmware.com/files/pdf/products/vsan/VMware-Virtual-SAN-6.1-Stretched-Cluster-Guide.pdf.

7.2 VMware vSphere Data Protection Advanced

vSAN interoperates with VMware vSphere Data Protection™ Advanced, a backup and recovery solution designed for vSphere environments and powered by EMC Avamar. vSphere Data Protection Advanced provides agentless backup and recovery of virtual machines running on VMware vSphere VMFS, NFS, and vSAN datastores. Backups are deduplicated using a variable-length segment algorithm, resulting in a significant reduction in backup data storage capacity consumption. Backup data can also be moved offsite using reliable, secure, network-efficient replication.

vSphere Data Protection Advanced uses VMware vSphere APIs that enable a product to back up virtual machines from a central backup server or virtual machine, without requiring agents or processing to be done inside each virtual machine.

Managing vSAN

8.1 Management and Operations

Because vSAN is a policy-driven storage solution, provisioning and management are significantly simplified. vSAN automatically and dynamically matches requirements with underlying storage resources. With vSAN, many manual storage tasks are automated to deliver a more efficient and cost-effective operational model. However, from an operational aspect, providers have additional concerns and requirements that go beyond the simplification and streamlining of management. vSAN enables additional capabilities that allow providers to monitor and manage their vSAN infrastructure.

8.1.1 Health Check Plug-In

The vSAN Health Check plug-in checks all aspects of a vSAN configuration. It implements a number of checks on hardware compatibility, networking configuration and operations, advanced vSAN configuration options, storage device health, and virtual machine object health. In addition to simplified troubleshooting, the Health Check plug-in enhances the providers' support experience by enabling uploading logs to an existing support request. In summary, Health Check provides two primary benefits for providers:

- Gives providers peace of mind that their vSAN deployment is fully supported, functional, and operational.
- Provides immediate indications to a root cause in the event of a failure, leading to faster resolution times.

VMware recommends that the Health Check plug-in be utilized to do initial triage of any vSAN problems. A detailed description of the operational aspects of the health check can be found in the *VMware Virtual SAN Health Check Plugin Guide* at http://www.vmware.com/files/pdf/products/vsan/VMW-GDL-VSAN-Health-Check.pdf.

8.1.2 VMware vRealize Operations Management Pack

The VMware vRealize Operations™ Management Pack for Storage Devices offers full support for vSAN and provides an end-to-end global view of vSAN from a centralized location. In typical vRealize Operations fashion, out-of-the-box monitoring is provided across clusters, HDD/SSD devices, network, CPU, and memory components. The vRealize Operations Management Pack for Storage Devices is located in the Cloud Marketplace on the *VMware Solution Exchange* web site at https://solutionexchange.vmware.com/store.

Conclusion

The goal of this document was to describe the considerations for designing a highly available solution for leveraging vSAN in a cloud environment. This document outlined a dedicated IaaS offering for a sample customer of a VMware Cloud Provider using the cost-effective, high availability, scalability, and predictable performance capabilities of vSAN.

Achieving economies of scale means scaling storage resources in a consistent and predictable manner. Scaling the storage capacity in vSAN can be achieved by simply adding more magnetic disks (hybrid configuration) or SSDs (all-flash) to an existing disk group or by creating an entirely new disk group. Scaling the performance layer (flash layer) of vSAN is achieved in a similar way to scaling the storage capacity in the sense that you can quickly select the necessary disk devices and create a disk group.

With the combined resources of the VMware Cloud Provider Program and vSAN, an organization can achieve business-critical levels of performance and availability for a variety of workloads including databases, ERP systems, streaming content, and web services. All this can be achieved while accelerating the time to deployment, reducing costs, and simplifying operations.

Assumptions and Caveats

VMware and other third-party hardware and software information provided in this document is based on the current performance estimates and feature capabilities of the versions of code indicated. These are subject to change.

Reference Documents

For more information, see the following configuration and administration guides, white papers, blogs, and best practices documents.

Document Title	Link or URL
VMware Virtual SAN 6.0 Design and Sizing Guide	https://www.vmware.com/files/pdf/products/vsan/VSAN_Design_and_Sizing_Guide.pdf
vSAN sizing tool	https://vsantco.vmware.com/vsan/SI/SIEV
VMware vSphere 6.0 Storage Guide	https://pubs.vmware.com/vsphere-60/topic/com.vmware.ICbase/PDF/vsphere-esxi-vcenter-server-60-storage-guide.pdf
Overview of vSAN 6.1	http://cormachogan.com/2015/08/31/a-brief-overview-of-new-virtual-san-6-1-features-and-functionality/
VMware Virtual SAN 6.1 Stretched Cluster Guide	https://www.vmware.com/files/pdf/products/vsan/VMware-Virtual-SAN-6.1-Stretched-Cluster-Guide.pdf
VMware Virtual SAN Health Check Plugin Guide	http://www.vmware.com/files/pdf/products/vsan/VMW-GDL-VSAN-Health-Check.pdf