VMware vCloud® Architecture Toolkit™
for Service Providers

Leveraging VMware vSAN™ for Highly Available Management Clusters

Version 2.9

January 2018

Danilo Feroce

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

Contents

List of Tables

List of Figures

Introduction

1.1 Overview

The VMware Cloud Provider™ Program is a global network of thousands of cloud service providers distributed across many regions who are offering or are planning to offer cloud and hosting services founded upon VMware cloud software platform. This platform allows them to provide world-class cloud services to their base of tenants and to build and package tailored offerings to win in their own market or location.

The Cloud Service Provider services are most commonly linked to a service plan that underlines the class of service being offered to the tenants. The flexibility of the software platform allows the integration of many different components. At the same time, complexity in a data center introduces questions about the reliability of each component and its influence on the overall solution.

The Cloud Management Platform is the management foundation for VMware Cloud Providers. It includes a critical set of integrated components to deliver a resilient environment for VMware vCloud® services consumers and to provide a powerful management instrument.

In such a delicate ecosystem, a critical component or a foundation pillar of every architecture that requires top priority care in design is the underlying storage infrastructure. The infrastructure is designed and built to provide the performance and reliability required for designated service levels.

The use case depicted in this paper covers the scenario of a hyper-converged infrastructure built with VMware vSAN™ to address the availability needs of a management cluster.

The key benefits for the use of a hyper-converged infrastructure solution include the following:

- Infrastructure cost reduction by means of replacing dedicated traditional storage devices
- Reduced management complexity thanks to the native integration of vSAN and VMware vSphere®
- Granular policies to tune performance and storage consumption based on a per-workload basis

1.2 Document Purpose and Scope

This document is primarily intended for members of Cloud Service Providers (technical and business leaders, solutions architects, application owners, and operational engineers) who are involved in planning, designing, deploying, and operating VMware Cloud Provider Program solutions, or who are interested in investigating the benefits of cloud solutions powered by VMware Cloud Provider Program platforms and services.

This document assumes the reader is familiar with VMware product suites and has a general understanding of systems and networking concepts and operations. Some references to operational frameworks and their terminologies are also included.

The reference architecture outlined in this paper is meant to represent a building block that can be replicated to scale the solution as workloads require. However, the scope of this document does not include specific or detailed sizing information. Although hardware vendors might be included in the examples, do not consider this as direction or guidance because VMware Cloud solutions are hardware-independent and vendor-agnostic.

1.3 Definitions, Acronyms, and Abbreviations

The following terms, acronyms, and abbreviations are used throughout this document.

Cloud Service Provider (CSP): A business entity that offers computing, storage, or software services powered by VMware cloud platforms to consumers through a private or public network.

Tenant (or client): A business entity using the services offered by the CSP and accepting some sort of service level agreement.

Consumer (or customer or end user): Someone who consumes the services offered by the tenant or directly by the CSP. Also accepts the service level agreement with the CSP through the tenant or has one of his own with the tenant.

Software-Defined Data Center (SDDC): A data center facility where the elements of the infrastructure are virtualized and delivered as a service and where the provisioning and operation is abstracted from the hardware and fully implemented through software.

Total Cost of Ownership (TOC): An estimate of the expenses associated with a product or with a particular equipment during the spanning of its entire lifecycle (purchase, deployment, management, usage, and retirement).

Information Technology Infrastructure Library (ITIL): A set of practices focused to align IT services and business by describing processes and their relationships with people and technologies.

Service: A means of delivering value to customers by facilitating outcomes customers want to achieve without the ownership of specific costs and risks.

Service Level Agreement (SLA): An agreement between a service provider and a customer. The SLA describes what service is to be provided, details how it is provided, and documents the responsibilities of the parties involved (service provider and customer).

Operation Level Agreement (OLA): An agreement among internal groups of the service provider, or between the service provider and its partners, that defines the relationships and responsibilities to support the level of service agreed upon.

Cloud Management Platform (CMP): A suite of integrated products that provide management for public, hybrid, and private cloud environments (includes self-service interfaces, provisioning systems, metering and billing, and workload optimization).

Hyper-Converged Infrastructure (HCI): A type of infrastructure system with a software-centric architecture tightly integrating compute, storage, networking, and virtualization resources and leveraging commodity hardware to support it.

vSAN: A software-defined enterprise-class shared storage solution embedded in the hypervisors.

VMware vCloud Suite®: An integrated suite of VMware technologies that offers cloud management capabilities.

VMware vRealize® Suite: A complement and extension of the vCloud Suite targeting the most complex cloud-based solution with hybrid environments and mixed hypervisors.

Open Virtualization Archive (OVA): An archive file containing a compressed and installable version of a virtual machine.

Wide Area Network (WAN): A telecommunications network that interconnects locations extended over a large geographical distance.

Metropolitan Area Network (MAN): A network that interconnects computer resources in a limited geographic area or region, usually identified by a city.

Round-Trip Time (RTT): The length of time it takes for a signal pulse or packet to travel from a specific source to a specific destination and back again.

Solution Architecture Overview

This section provides a high-level overview of the elements and topics that play a role in the proposed solution: management components, storage infrastructure, and high availability.

Each of these three elements is strictly interconnected with the others to attain an integrated solution that is easily replicated and offers out-of-the box increased value.

2.1 Management Cluster Conceptual Overview

A management-only cluster is a group of host hypervisors joined together to host all the VMware vCloud® management components required for the specific environment under examination. The cluster resources are reserved for the specific and sole purpose of powering and supporting the management components with the desired performance, availability, security, and reliability.

The following products are key Cloud Management Platform (CMP) elements. The list is not exhaustive and might contain additional service provider management services to increase operational manageability and efficiency of the solution.

- VMware vCenter Server® and its database, or VMware vCenter® Server Appliance®
- VMware vCloud Director® cells and database
- VMware NSX® Manager™
- VMware vRealize® Automation™
- VMware vRealize Operations Manager™
- VMware vRealize Log Insight™
- VMware Site Recovery Manager™
- Optional infrastructure services tailored for the specific environment (directory services, name resolution, time keeping, IP address management, logging, and so on)

The advantage of deploying a management cluster separated from the other resource clusters is to provide access to sufficient resources (compute, storage, network) at any time without having to compete with other virtualized workloads, and the freedom to manage different plans for data protection and recoverability tailored to the specific workloads running on the clusters.

The following figure shows a conceptual overview of the management components overlaying a vSphere cluster of hypervisors.

Figure 1. Management Cluster Overview

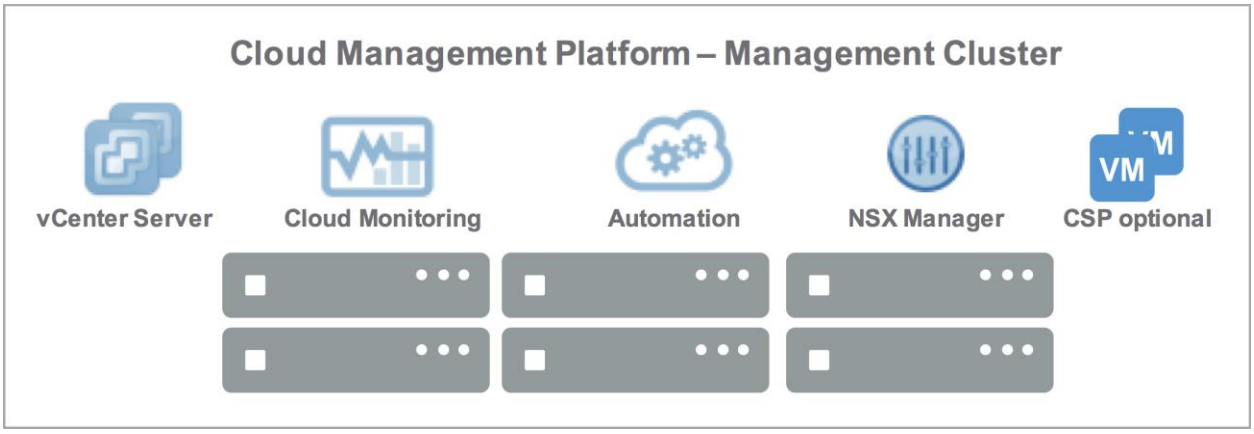

In this scenario, the management cluster and its managed resource groups supporting the vCloud infrastructure reside on the same site to minimize latency issues associated with the speed or reliability of the network.

2.2 vSAN Conceptual Overview

vSAN is a hyper-converged, software-defined storage solution offering enterprise class performance, reliability, and availability. vSAN is designed to take advantage of commodity servers with locally-attached storage (traditional spindle-based disks or high-performance flash devices) to create and maintain its storage abstraction layer (datastore), saving and manipulating data in the form of storage objects.

vSAN architecture is natively embedded in the hypervisor storage stack and provides complete integration with the vSphere management interface while supporting its advanced availability features:

- VMware vSphere vMotion®
- VMware vSphere Fault Tolerance (FT), providing continuous availability for virtual machines up to a limited size of 4 vCPUs and 64 GB RAM
- VMware vSphere High Availability (HA), providing automatic restart of virtual machines on remaining hosts in case of failure of one of the cluster nodes
- VMware vSphere Data Protection™, providing a combination of backup and restore functionalities for both virtual machines and applications
- VMware vSphere Replication™, providing asynchronous virtual machine replication

vSAN supports two main configurations in respect to the underlying storage devices connected to the nodes:

- Hybrid – A combination of magnetic disks used for the capacity tier, and flash devices used for the cache tier. Traditional spindle-based disks are supported in several formats and speeds, and are usually selected based on the capacity requirements of the environment.
- All-flash optimized – Both capacity and cache are supported by flash-based devices (write-intensive/high-endurance and read-intensive/cost-effective SSD devices, depending on the associated tier they are supporting). vSAN using all-flash clusters is available with version 6.0 and later, and requires a 10-Gbps network between the nodes of the cluster.

The following figure shows a typical hybrid vSAN deployment, where the storage resources of the four hypervisors represented are joined together in the vSAN datastore for that cluster and made available to all the virtual machines running on the nodes.

Figure 2. vSphere and vSAN

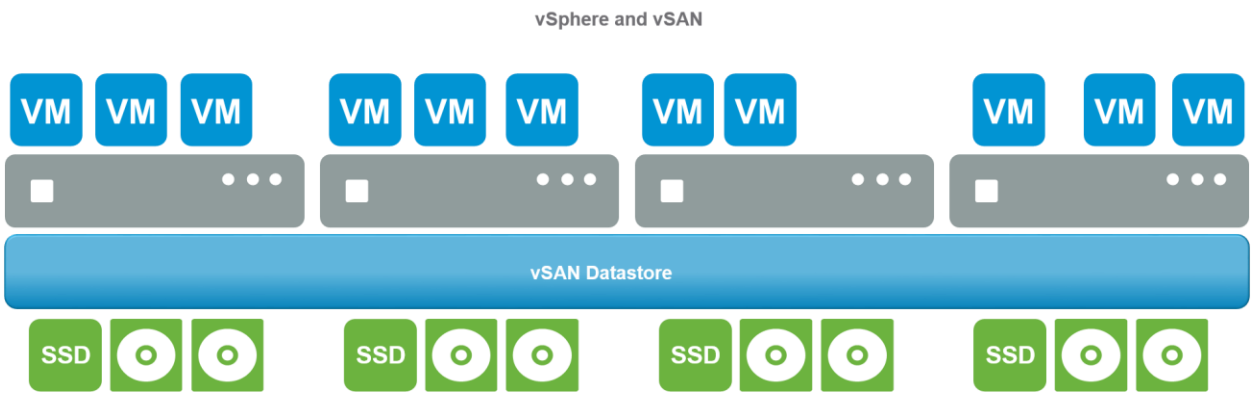

2.2.1 Number of Nodes

vSAN supports clusters with a variable number of nodes depending on the implemented version of the product. The number of nodes in a cluster defines the behavior of the cluster in case of failure of a node.

- Two-node clusters (vSAN 6.1 and later) hold up to two replicas of the data and require an additional witness residing on a third host outside the cluster.
- Three-node clusters hold up to two replicas of the data and a witness, all residing on different hosts of the cluster.
- Four-node or more clusters (up to 64 hosts for vSAN 6.x, or up to 32 hosts for earlier releases) hold three or more copies of the data, all residing on different hosts of the cluster.

2.2.2 Fault Domains Setting

vSAN implements a concept identified as *fault domains*. The fault domains feature allows vSAN to group together multiple hosts (typically within the same chassis or rack) into a logical boundary domain. The fault domains setting makes sure that multiple replica copies of storage objects are not saved within the same boundary, but are instead distributed across the domains. This way, in case of the failure of an entire domain (chassis or rack), only one replica is affected.

The following figure displays how a storage object of a virtual machine is saved within vSAN when the fault domains feature is enabled. In this case, three fault domains are identified and each one contains either one replica copy of the object (R1 or R2) or the witness (W).

Figure 3. vSAN Replica Objects Distribution

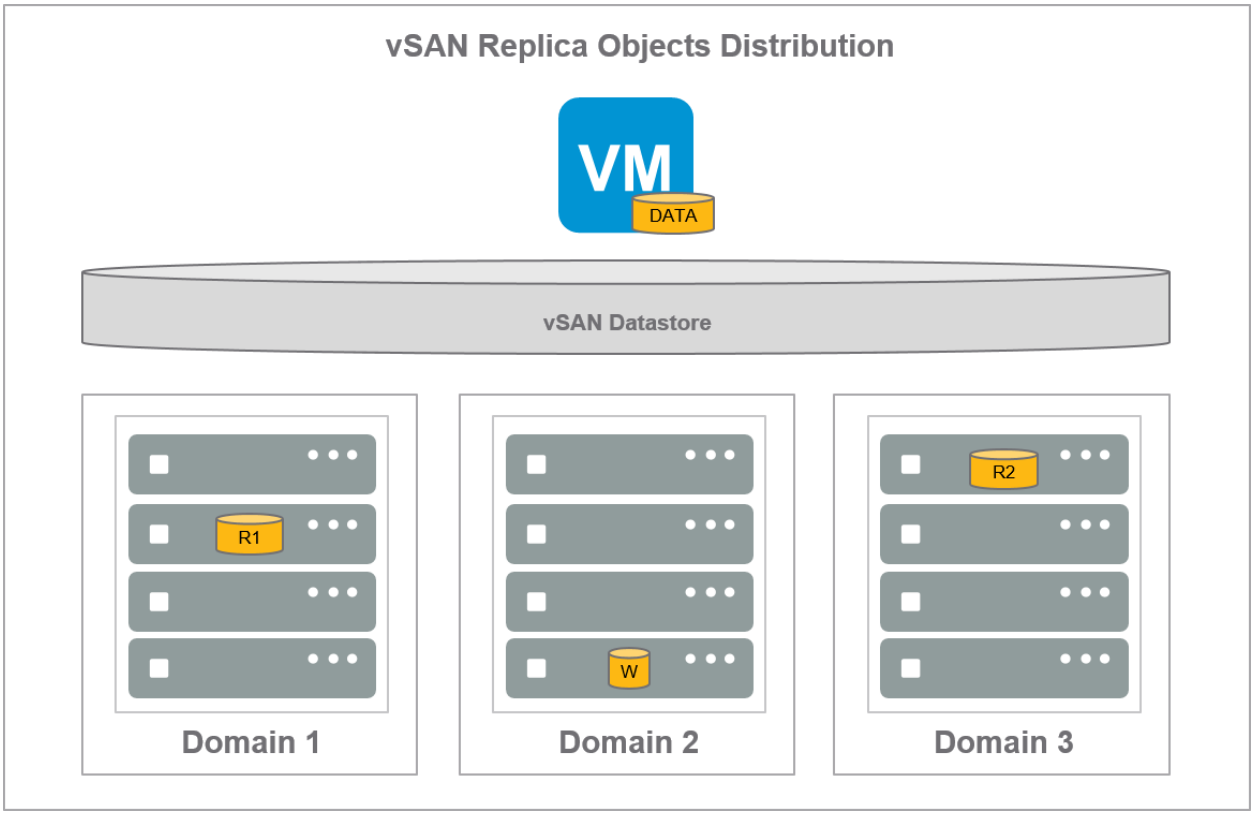

2.2.3 Failures to Tolerate policy

The *failures to tolerate* virtual machine policy (FTT or NumberOfFailuresToTolerate) regulates the number of failures the vSAN underlying the VM is able to sustain while the VM remains available. The fault domains and the FTT policy are tightly related when both are configured.

The following table reports the FTT values with the associated number of fault domains to implement to support the policy, and the common formula to calculate the fault domains (2n+1) when the FTT value (n) is selected[1].

Table 1. FTT Policy and Required Hosts

Number of Failures To Tolerate (FTT)	Number of Fault Domains Required
n	2n + 1
1	3
2	5
3 (maximum)	7

The more virtual machines that are deployed with the FTT policy enabled and the higher the value of the policy setting, the greater the impact on the storage capacity requirements for the vSAN environment.

2.2.4 vSAN and vSphere HA considerations

The combination of vSAN as a highly available storage infrastructure and vSphere HA as a highly available solution for virtual machine compute resources provides a powerful synergy for delicate workloads.

The following are requirements for the interoperability of vSAN and vSphere HA to function:

- vSAN must be configured on the cluster before enabling vSphere HA, or vSphere HA must be temporarily disabled to allow vSAN configuration
- vSphere HA must use the vSAN network for its communication
- vSphere HA must not use the vSAN datastore as the Heartbeat datastore
- vSAN and vSphere HA do not interoperate to pre-allocate storage resources in case of a host failure, contrary to what vSphere HA does by reserving CPU and memory resources. The storage sizing is a task delegated to the administrators.

The following figure illustrates the dynamic of the vSphere HA feature in action at the occurrence of a failure of a single host within the pool of nodes in a cluster. When the failure occurs, the computing resources and the portion of the datastore storage allocated from the lost node become unavailable. vSphere HA then restarts the virtual machines previously running on the failed node, redistributing them across the remaining nodes of the cluster.

[1] Virtual SAN 6.2 RAID Fault Tolerance methods require additional considerations.

Figure 4. vSphere High Availability in Action

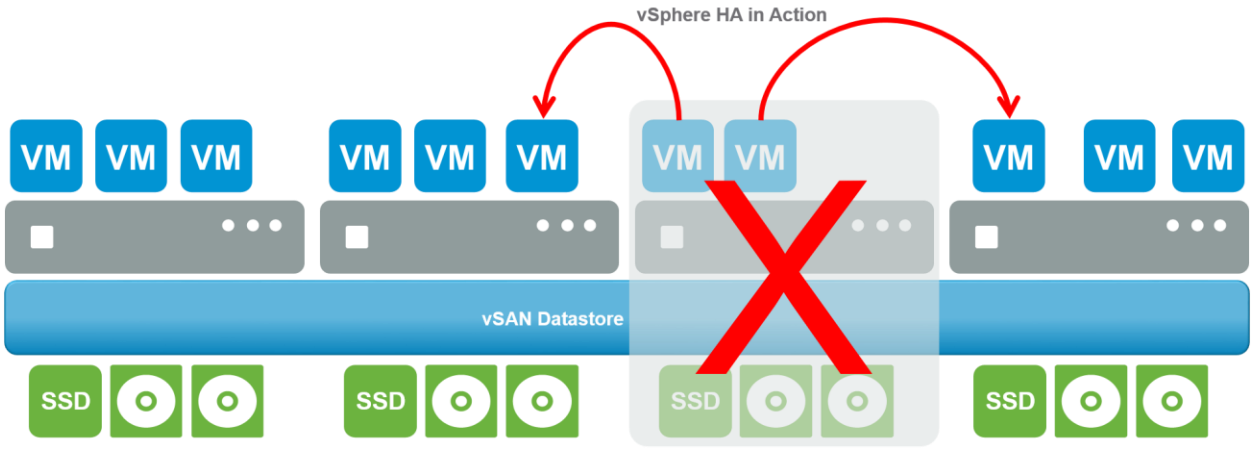

2.3 Availability Concepts Overview

The concept of availability extends across multiple fields, in addition to cloud computing. When dealing with computer systems and networking, as a general term, *availability* can be described as "the degree to which a system or component is operational and accessible when required for use," according to the IEEE 610 definition (see Section 6, References).

The availability of a system or component is commonly calculated as a percentage (of the time the system or component is available when compared to the whole), and it is often measured in terms of number of 9s, for example, 99.99%.

A formula to capture the availability of a simple system or component is calculated as follows, by tracking the uptime and downtime of the system itself:

```
Availability = Uptime/(Uptime + Downtime)
```

The overall availability of a more complex system is calculated as follows, by multiplying the availability in percentage of each component present in the system:

```
Overall Availability = Element#1(availability%)*Element#2(%)*... Element#n(%)
```

A computer server running in a data center represents a complex system when approached as a set of separate devices, each with its own availability (calculated or most likely provided by the manufacturer). To the extent that the rack containing the server and additional devices represents a more evolved complex system to measure in terms of overall availability, the following items must be considered:

- Rack power supplies, cabling, and so on
- Top-of-the-rack network switches
- Server's local storage / hard disks
- Server's power supply (supplies, if redundant)
- Server's internal components (CPUs, memory, controller, and so on)

A reference set of availability percentages, with their corresponding number of 9s and calculated downtime per year, is detailed in the following table.

Table 2. System Availability and Downtime per Year

Number of 9s	Availability %	Downtime/Year	System or Component Inaccessible
1	90%	36.5 days	Over 5 weeks per year
2	99%	3.65 days	Less than 4 days per year
3	99.9%	8.76 hours	Approximately 9 hours per year
4	99.99%	52.56 minutes	Approximately 1 hour per year
5	99.999%	5.26 minutes	Approximately 5 minutes per year
6	99.9999%	31.5 seconds	Approximately half a minute per year

Cloud Service Provider Use Case

3.1 Business Drivers

The key business drivers and benefits for a Cloud Service Provider to implement this solution are as follows:

- Lower storage TCO by deploying standardized hardware components supported by vSAN without elevated upfront financial exposure.
- Single pane of glass management thanks to the native integration of storage, compute, and networking resources.
- Storage performance and consumption allocated on a per-workload basis through granular policies.
- Reduce the cost and complexity of any component of the service provider infrastructure, which translates to a reduced financial burden when the services' costs for the consumers and tenants are defined.

3.2 Cloud Service Provider Considerations

A typical architecture in a Cloud Service Provider environment consists of one or more resource groups dedicated to the tenant workloads supported by the components deployed on the management cluster.

The logical separation between management and workload resources is reflected by the separation of the supporting physical infrastructures for these two areas. The main drivers for this separation are as follows:

- Separation of duties between the management resources (CSP ownership) and the single or multitenant resource groups (tenants or CSP ownership depending on the model implemented).
- Different approaches for data protection, availability, and recoverability between the management resources and the tenant workloads, which might also differentiate the service level of these environments.
- Full independence of the tenant resource clusters from the management layer allowing the freedom to scale the workload resources vertically or horizontally as building blocks according to business needs and without additional constraints.

3.2.1 Multisite Service Provider Considerations

Locate the management cluster and the related resource groups at the same physical site to provide a consistent experience in the level of service provided, because some of the management components might have restrictions when deployed across geographically distributed networks (for example, vCloud Director support for multisite deployments). The main reasons to recommend same-site management is to optimize the connectivity and reliability by avoiding possible latency issues. Nevertheless, the management across distributed topologies or stretched clusters across multiple sites are feasible solutions that can be assessed as viable on a case-by-case basis.

A Cloud Service Provider might operate a more complex topology of data centers with multisite scenarios to manage and supervise. Some of these scenarios might include, but are not limited to:

- MAN and/or WAN topologies with their connectivity speed and latency.
- Network layer available between sites (Layer 2 or Layer 3).
- Resource clusters offering services for workloads that require either stretched or separated configurations across the multisite topologies.

The options and advantages to deploy localized management clusters in each site are related to each characteristic workload:

- Multisite topologies with WAN or MAN connectivity having an RTT over 20 ms are not suited for deployment of single-site style management infrastructures including vCloud Director.
- Multisite topologies with WAN or MAN connectivity having RTT below 20 ms, but with Layer 3 only connectivity, still require separate management clusters.
- Multisite topologies with WAN or MAN connectivity having RTT below 20 ms and Layer 2 connectivity offer the choice between stretched or separate management clusters.
- In the event of a site failure and subsequent recovery, the management functions are the first and most rapid to be recovered and so are preferred to be at the same location of the resources to be managed.
- vSAN recommends communications across data nodes to be over Layer 2, and for witness only communications to occur over Layer 3.
- vSAN and vSphere FT are not supported together in a stretched virtual SAN cluster deployment.

3.3 Tenant Viewpoint

The architecture and infrastructure selected for the management clusters supporting single or multitenant environments are transparent to vCloud consumers.

The definition of the level of service for each offering of the Cloud Service Provider must declare the expected overall availability of the infrastructure to allow the consumer of the vCloud services to correctly plan the positioning of its workloads depending on their criticality and on business needs.

The Cloud Management Platform, its architecture, and its availability are critical for delivering a resilient tenant infrastructure. If a Service Level Agreement is defined between the Cloud Service Provider and the vCloud consumers, the level of availability for the CMP must match or be at least comparable to the highest requirement defined across the SLAs to maintain both the management cluster and the resource groups in the same availability zone.

3.4 vSAN Monitoring

While vSAN is integrated with the vSphere management environment, many details of its configuration and status are not immediately visible, because the architecture can become complex depending on the deployment, the amount of hardware involved, and the amount and kind of workloads assisted.

The vSAN Health Check Plug-In, distributed as an additional software component to add to the vCenter Server installation (either vCenter Server or vCenter Server Appliance), monitors every aspect of the vSAN configuration, including the following:

- Hardware compatibility
- Storage device health
- Virtual machine object health
- Network configuration
- Network operations
- Advanced vSAN configuration options

The main benefits for the use of the plug-in are to gain immediate feedback about the supportability, functionality, and operational status of the components of a vSAN. The Health Check Plug-In also has the ability to upload logs to a support request, streamlining the customer experience and case troubleshooting for any Cloud Service Provider encountering configuration or operation issues.

3.5 vSAN Operations

To achieve an even greater insight to vSAN operations, the VMware vRealize Operations Management Pack™ can be added to a pre-existing vRealize Operations Suite installation.

The management pack offers several out-of-the-box dashboards with relevant information on the vSAN activities:

- Global View to monitor and proactively alert/notify across multiple vSAN clusters, including device and cluster insights, entity usage, heat maps, VMs, hosts, datastores, and disk group relationships.
- Health Monitoring and Availability View to assess device connectivity issues, failures, network congestions, SSD life, and disk failures.
- Performance View to evaluate aggregated data on throughput performance and latencies at the disk group and hard disk level (both SSD and HDD).
- Capacity View to monitor disk usage and perform capacity planning based on simulation scenarios.

Figure 5. vRealize Operations vSAN Management Pack Sample View

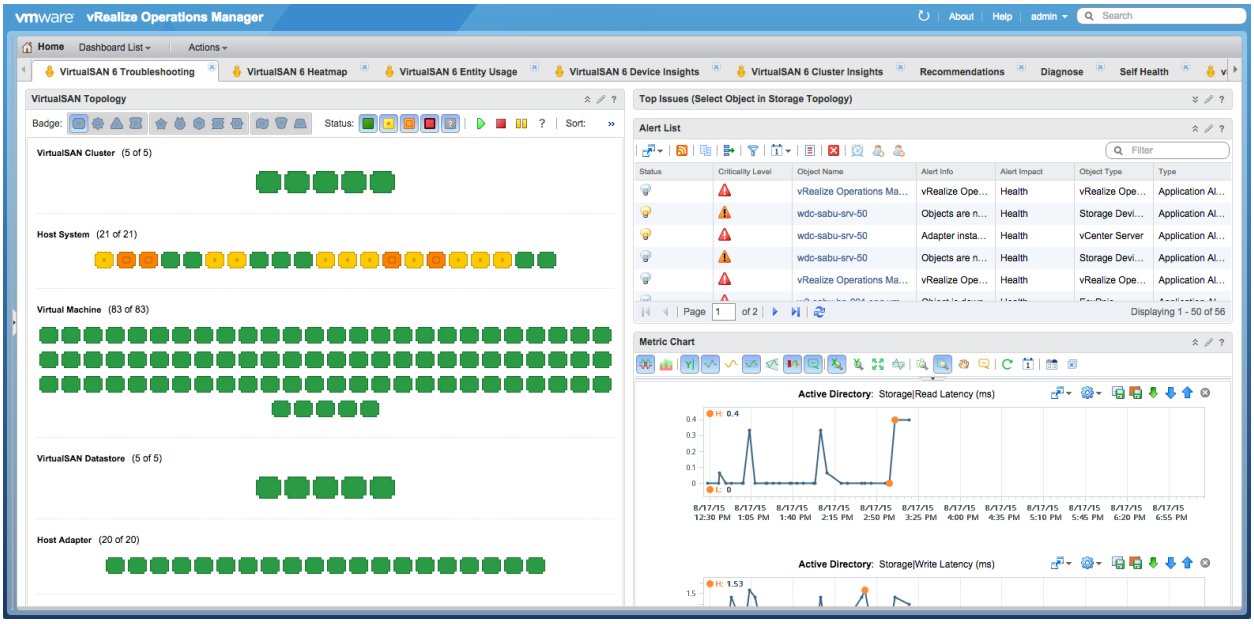

Use Case Architecture

4.1 Two-Node vSAN Management Cluster Overview

The system design example described in this section provides a scenario for a Cloud Management Platform deployment over a cluster with basic resiliency that is supported by vSAN. The design elements for this architecture include the following:

- The management cluster holds dedicated compute, network, and storage resources.
- The management cluster is entirely located at a single site.
- The cluster supports a full CMP suite of products, including optional CSP-related services.
- vSAN powers the storage solution, vSphere HA is enabled for the VMs, and additional high availability features (vSphere FT and vSphere Data Protection) are enabled where required for the workloads.

The following figure illustrates the deployment for this scenario.

Figure 6. Example of Cloud Management Platform Deployment over a vSAN Cluster

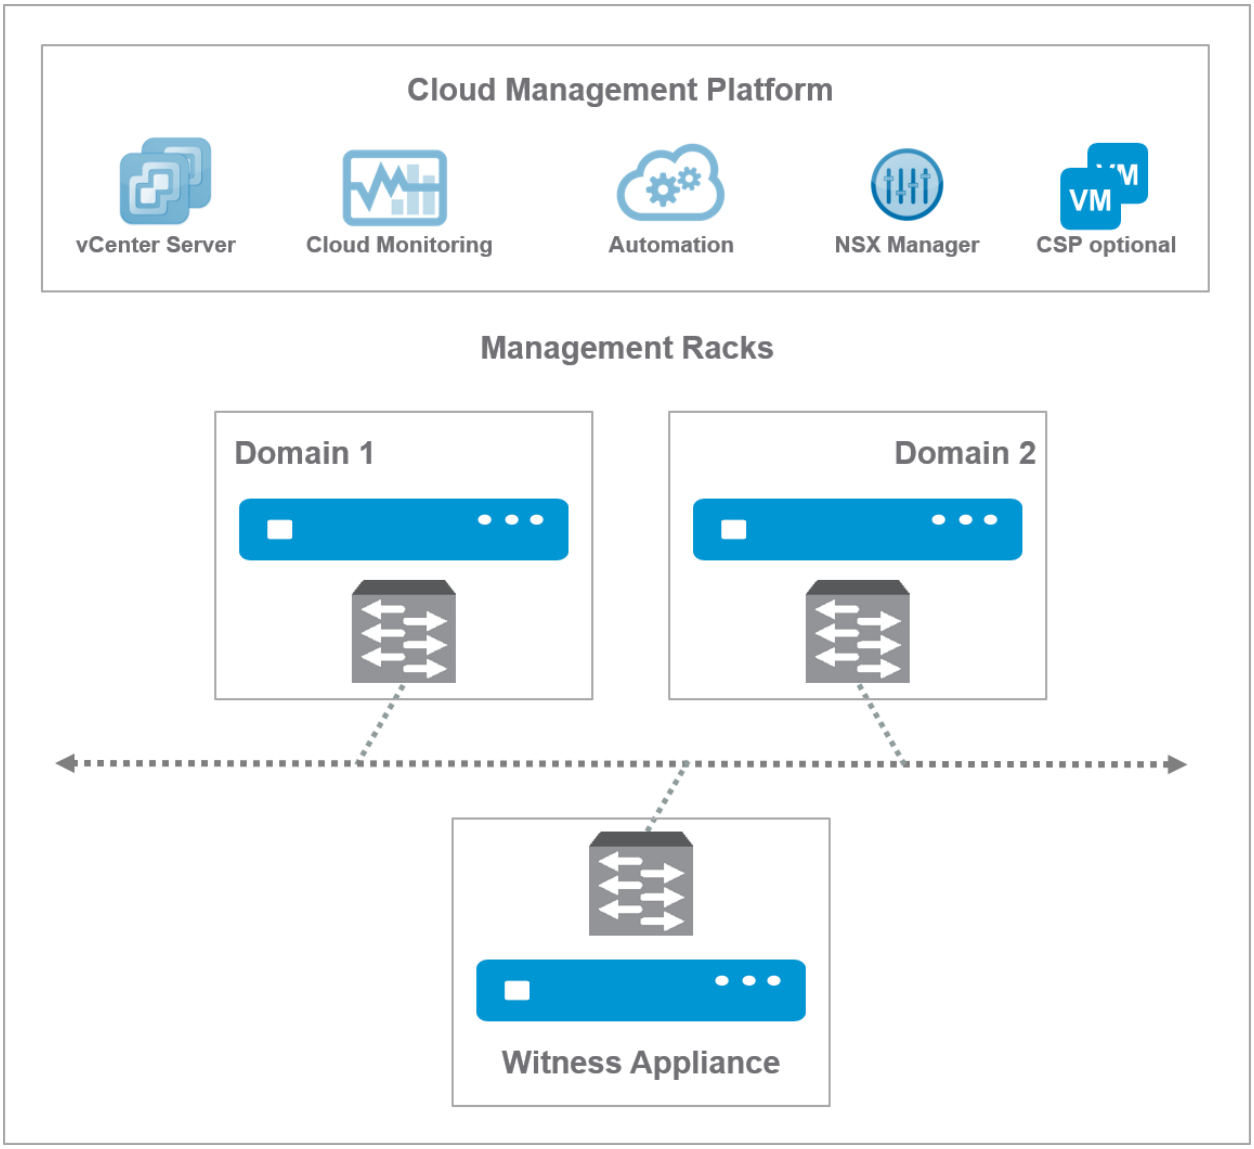

Additional specific configurations, also represented in the preceding figure, include the following:

- A minimal two-node cluster is implemented.
- A witness host appliance for vSAN is deployed on another server in the local data center.
- Two fault domains are identified with the two rack servers hosting the respective nodes of the cluster. One additional domain refers to the witness host.
- Failures to Tolerate policy is set to the value of 1 to enforce two copies per object.

The availability of the vSAN architecture in this example is close to four 9s, or 99.99%. vSAN is able to provide greater availability rates by increasing the number of copies per object (FTT) and the number of fault domains, and therefore, the number of data nodes in the cluster as well.

Some of the availability metrics for computing the overall availability are variable and lie outside the scope of this architecture example because they are primarily linked to hardware resources. These variable metrics include the following:

- Sizing of the resources presented to the virtual machines running the management workloads.
- Rack (power supplies, cabling, and so on).
- Top-of-the-rack network switches.
- Host physical server and hardware components (including CPUs, memory, controller, and so on).
- Hard disk MTBF (traditional spindle-based disks and high-performance flash devices).
- Hard disk capacity and performance, which influences rebuild time.
- FTT setting, which influences the required capacity of the management cluster.

4.2 Management Cluster Operation Restrictions

The vSAN cluster with two data nodes and one witness has the following restrictions during regular and extraordinary operations:

- When a data node fails, the workloads are still operative, but no attempt to rebuild the lost data can be made until the failure is resolved or the failed node is replaced. The drawback of a small two-node configuration is the reduced reliability after the first failure.
- When a data node is placed in maintenance mode, the same behavior applies until the node is back online.
- During regular operation, storage I/O is split equally across the two data nodes. During an outage or maintenance, 100 percent of I/O is sustained by one data node and the storage performance must be sized to sustain that particular load.

Note If any of these constraints are not acceptable for the SLA in place with the tenants or with the end users, a larger deployment with three data nodes must be considered as the reference configuration instead.

4.3 Witnesses and Witness Host

In versions prior to vSAN 6.0, witnesses are an element of every storage object. A witness does not contain data, only metadata (approximately 2 MB per witness object). The purpose of a witness is to serve as a tiebreaker when a failure in the cluster occurs and an availability decision must be made to enforce the failures to tolerate policy.

vSAN 6.0 introduces a slightly different way to manage how the cluster's quorum is computed: each element has a number of votes which contributes to the final decision in the case of failure. It is possible for an object to be distributed in such a way as to not require the witness for the failures to tolerate decision.

4.3.1 Witness Host

With a two-node cluster, as in this example, the minimum configuration still requires a third node as witness, hosting exclusively metadata, but not data.

The witness host for a vSAN is supported using two different deployment models:

- A physical server with ESXi host installed.
- A special appliance containing a virtual ESXi host deployed as virtual machine (nested ESXi).

In either case, the witness host (physical or virtual appliance) is dedicated to one vSAN cluster. It cannot be shared with more than one cluster.

4.3.2 Witness Host Virtual Appliance

The witness host appliance is provided by VMware in the form of an Open Virtualization Archive (OVA) that must reside on a physical ESXi host as a nested hypervisor. Contrary to a physical witness host, the witness host appliance does not consume a vSphere license to run. The witness host appliance must not reside on the same vSAN cluster for which it is providing services.

4.3.3 Witness Host and Virtual Appliance Sizing

For the physical witness host, the resource sizing requirements are met by the minimum ESXi host requirements.

For the witness host appliance, the resource sizing depends on the configuration selected during the OVA deployment (see the following table).

The sizing requirement for this architecture example is small to normal, depending how many virtual machines are loaded into the CMP. A normal configuration for the witness host appliance is recommended for the current scenario of a management cluster.

Table 3. Witness Host Appliance Sizing

Deployment Model	Configurations
Tiny (10 VMs or fewer)	2 vCPUs, 8 GB vRAM, 8 GB ESXi boot disk, one 10-GB SSD, one 15-GB HDD Supports a maximum of 750 witness components
Normal (up to 500 VMs)	2 vCPUs, 16 GB vRAM, 8 GB ESXi boot disk, one 10-GB SSD, one 350-GB HDD Supports a maximum of 22,000 witness components
Large (more than 500 VMs)	2 vCPUs, 32 GB vRAM, 8 GB ESXi boot disk, one 10-GB SSD, one 350-GB HDD Supports a maximum of 45,000 witness components

4.4 Architecture Prerequisites and Constraints

The following table lists the basic prerequisites for the architectural example: software versions, configurations, and operations.

Table 4. Architecture Prerequisites Summary

Area	Specific requirement
VMware vSAN	Version 6.1 or later (to support 2 nodes and a witness host)
Witness Host	Network connectivity from the witness host to the vSAN network
Witness Host	Dedicated deployment, one for each vSAN cluster
Witness Host	Must run same vSphere version as the cluster data nodes
ESXi Supporting Witness Host Appliance	Must run vSphere v.5.5 or later
Witness Host Appliance	Must not reside on the same vSAN cluster resources

Conclusion

As outlined in this document, vSAN provides a significant value for any VMware Cloud Provider business in terms of both CapEx and OpEx. vSAN replaces traditional storage devices that require significant upfront investments with standardized commodities hardware. It also integrates the storage management operations under the same pane of glass used by the vCloud management platform instead of using multiple interfaces built on proprietary devices having variable learning curves.

vSAN enables a wide range of workloads, from the less critical to the enterprise-class workloads, with great flexibility and granular policy control. It becomes the strategic hyper-converged infrastructure solution in an elastic, non-disruptive, scalable, and evolving environment that is a software-defined data center founded upon the VMware Cloud Provider Program.

Additionally, the use case of the management cluster allows a service provider to introduce and evaluate the vSAN product in a scenario that is transparent to the tenants' workloads, providing the opportunity of adoption on a larger scale at a subsequent time for service providers new to the vSAN architecture.

References

The following table lists references for additional information pertinent to this document and its topics.

Document Title	Link or URL
VMware vCloud Architecture Toolkit for Service Providers	https://www.vmware.com/cloud-computing/cloud-architecture/vcat-sp.html
vCloud Architecture Toolkit (vCAT) Blog	https://blogs.vmware.com/vcat/
VMware vCenter Server	https://www.vmware.com/products/vcenter-server/
VMware vSAN	https://www.vmvware.com/products/virtual-san
VMware Virtual SAN 6.2 Design and Sizing Guide	http://www.vmware.com/files/pdf/products/vsan/virtual-san-6.2-design-and-sizing-guide.pdf
VMware Virtual SAN Health Check Guide	http://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/products/vsan/vmware-virtual-san-health-check-guide-6.1.pdf
VMware vRealize Operations Management Pack for Storage Devices	https://solutionexchange.vmware.com/store/products/management-pack-for-storage-devices#.VdKWSnivN5g
IEEE Standard Glossary of Software Engineering Terminology	http://ieeexplore.ieee.org/servlet/opac?punumber=2238