VMware® Virtual SAN™ Design and Sizing Guide for Horizon™ View™ Virtual Desktop Infrastructures

TECHNICAL MARKETING DOCUMENTATION
REV A/JULY 2014

vmware®

Table of Contents

Introduction

VMware Virtual SAN

VMware Virtual SAN™ is a hypervisor-converged, software-defined storage platform that is fully integrated with VMware vSphere®. Virtual SAN aggregates locally attached disks of hosts that are members of a vSphere cluster, to create a distributed shared storage solution. Virtual SAN enables the rapid provisioning of storage within VMware vCenter™ as part of virtual machine creation and deployment operations.

Virtual SAN uses a hybrid disk architecture that leverages both flash-based devices for performance and magnetic disks for capacity and persistent data storage. This hybrid architecture delivers a scale-out storage platform with enterprise performance and resiliency at a compelling price point.

The distributed datastore of Virtual SAN is an object-store file system that leverages the vSphere Storage Policy–Based Management (SPBM) framework to deliver application-centric storage services and capabilities that are centrally managed through vSphere virtual machine storage policies.

This document focuses on the definitions, sizing guidelines, and characteristics of the Virtual SAN distributed datastore for Horizon™ View™ virtual desktop infrastructures.

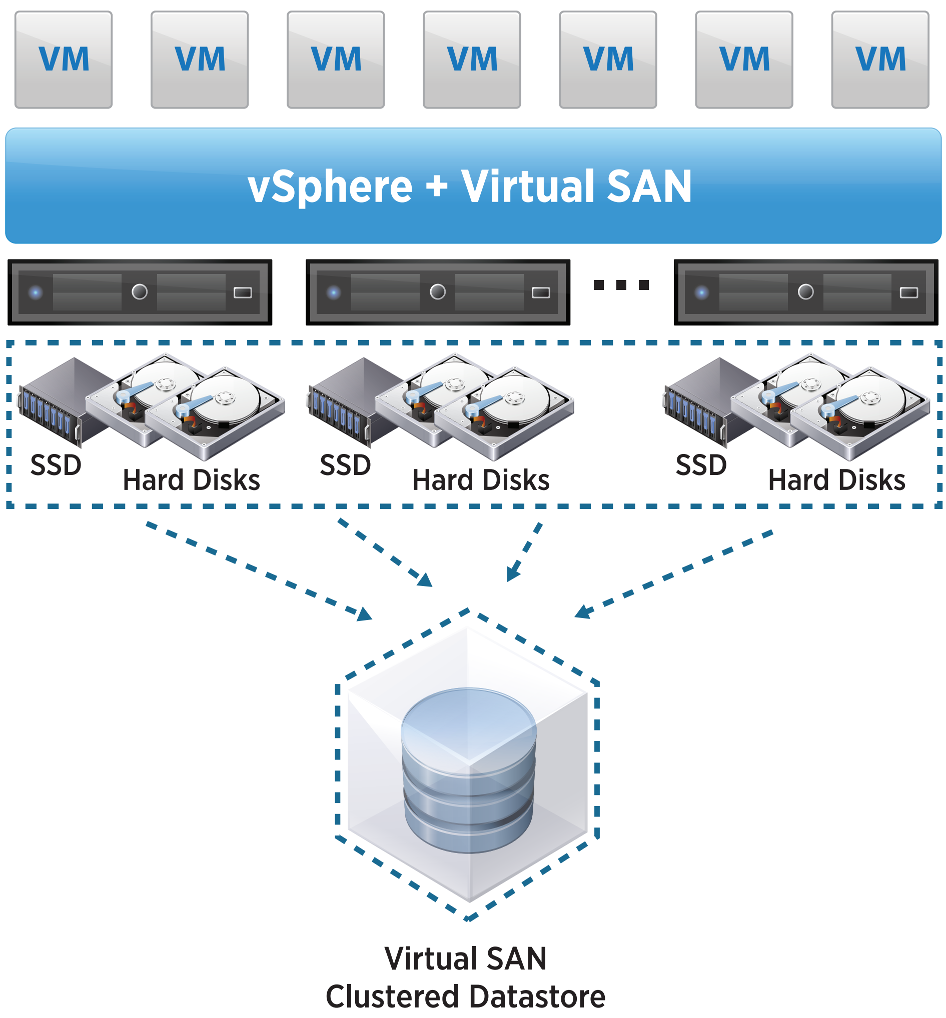

Figure 1. VMware Virtual SAN

Horizon View

Horizon View is a virtual desktop infrastructure solution that simplifies management by placing all of an organization's desktops into one centralized service that can easily customize them to fit user needs.

View delivers the highest-fidelity experience for geographically dispersed users who need applications, unified communications, and 3D graphics as part of their daily workspace.

With access available from a wide variety of device platforms, and performance optimization to accommodate users in even the most remote locations, View delivers maximum workspace productivity and convenience for end users. View is part of VMware Horizon™.

Horizon Suite is designed for today's mobile workforce, connecting end users to their data and applications from anywhere, at any time, and from any device, without sacrificing IT security and control. With Horizon Suite, IT has control, and users have choice.

Visit the VMware Web site for more information on Horizon Suite.

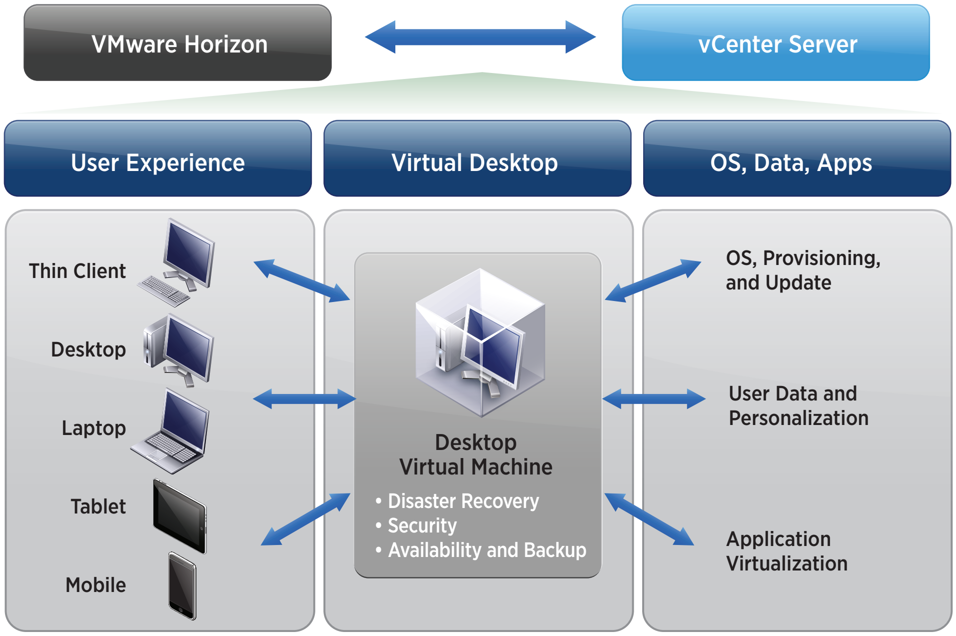

Figure 2. Horizon View

Virtual SAN Characteristics

Virtual SAN is a hypervisor-converged solution that combines compute resources and locally attached storage devices to deliver CPU and memory resources as well as a distributed shared datastore in a vSphere cluster.

From a storage solutions perspective, Virtual SAN introduces a number of new and unique constructs that impact the way in which View designs and sizing considerations are performed. The following are Virtual SAN constructs:

- Disk groups
- Datastore
- Objects
- Components
- Network
- Storage Policy–Based Management (SPBM)

VMware Virtual SAN Design and Sizing Guide provides a detailed description of each of these constructs: https://www.vmware.com/files/pdf/products/vsan/VSAN_Design_and_Sizing_Guide.pdf

System Sizing and Desktop Classification

Virtual SAN provides scale-out storage within a View infrastructure, enabling use of a grow-as-you-go model either by scaling up a storage infrastructure on a per-host basis or scaling out through addition of nodes to the Virtual SAN cluster.

When sizing a hypervisor-converged Virtual SAN node, a number of factors should be considered. Virtual SAN supports as many as 100 desktops per host and as many as 32 hosts per cluster. The higher the quantity of desktops to be deployed per host, the more hardware resources that are necessary to support them.

In addition to the number of desktops per host, the following factors impact the required hardware resources and sizing guidelines for a solution of View on Virtual SAN:

- Choice of desktop operating system (OS), applications, and class of worker (for example, knowledge worker, task worker, and so on)
  - This affects the desktop I/O profile and hardware resource requirements, which in turn drive the underlying server and storage design.
- Type of View pool used
  - Floating linked-clone virtual desktops
  - Dedicated linked-clone virtual desktops
  - Dedicated full-clone virtual desktops
- Method of desktop connectivity (RDP/PCoIP/HTML5) and user type, impacting both network bandwidth consumption and CPU utilization per desktop

Host Sizing Considerations

When classifying suitable hosts for View and Virtual SAN designs, it is important to identify the necessary CPU and memory requirements to support a particular virtual desktop design. These requirements are calculated to provide an estimate on resources a host must provide.

The sizing estimation should include all components that will utilize CPU and memory resources, such as virtual desktops and their specific workload profiles, as well as additional amounts of resources to manage anticipated peaks without negatively impacting the infrastructure. This information is essential to providing an adequate estimated value to support a chosen design.

For View designs that are intended to use Virtual SAN, it is important to account for the CPU and memory estimated overhead introduced by Virtual SAN as part of the sizing estimation calculation.

Host CPU Sizing

The sizing exercise for CPU requirements for a host should account for the average number of vCPUs that will be utilized per virtual desktop, as well as the average CPU processing requirements, measured in megahertz.

VMware recommends increasing the average desktop CPU utilization by as much as 15 percent to provide additional resources to account for virtualization components and processes overhead such as RDP/PCoIP, VMware Tools™, VMware View Agent®, and others.

For View designs that are based on the use of Virtual SAN, VMware recommends adding as much as an additional 10 percent of CPU utilization for Virtual SAN management overhead.

After the average number of vCPUs per virtual desktop and the average CPU utilization have been identified, the total amount of CPU resources required per host can be calculated. Table 1 summarizes the items utilized as part of the host assessment.

CPU SIZING ASSESSMENT	
Average number of CPUs per virtual desktop	Implementation specific
Average CPU utilization per virtual desktop	Implementation specific
Additional CPU percentage to cover View solution overhead (RDP/PCoIP, VMware Tools, and others) per desktop virtual machine	15 percent or less
Additional CPU percentage to cover Virtual SAN overhead per host	10 percent or less

Table 1. Host CPU MHz Requirement Sizing

VMware recommends the use of VMware View Planner as well as third-party tools to obtain overall and average CPU utilization metrics. The ability to perform an accurate assessment that is based on actual requirements will provide the best possible results and overall design.

Memory Sizing

The memory sizing calculations should incorporate the total amount of memory required by the virtual desktops as well as any memory overhead that will be introduced by hosts and other components for management purposes.

The goal is to identify, obtain, and gather all of the usable memory requirements, after which calculating the amount of memory required per host is a straightforward process.

The use of vSphere host Content Based Read Cache (CBRC) capability should be considered particularly for virtual desktop infrastructures based on View. With View, the vSphere host CBRC feature is known as VMware View® Accelerator™. This feature is recommended to address the I/O created during operations such as simultaneous boot-up of a high number of virtual machines (boot storms) in a solution with View on Virtual SAN. The recommendation is to allocate 2GB of memory for View Accelerator, because that is the maximum allowed.

Virtual SAN is designed to consume no more than 10 percent of memory per host, based on the maximum number of devices and capacity the system will manage. Table 2 summarizes the factors to consider in memory sizing calculations.

HOST MEMORY OVERHEAD	Number of concurrent active users	Memory allocated per virtual desktop	View Accelerator (as much as 2GB)	Virtual SAN 10 percent memory overhead
VIDEO MEMORY OVERHEAD	Number of monitors per virtual desktop	Monitor resolution	Video memory overhead per virtual desktop	
vCPU MEMORY OVERHEAD	Number of vCPUs per virtual desktop	vCPU memory overhead per virtual desktop		

Table 2. Host Memory Requirements Categories

Memory sizing should be carefully considered. Granting a virtual desktop too much memory impacts the storage capacity utilized for the virtual machine swap files and Microsoft Windows paging files. Granting a virtual desktop too little memory impacts the user experience. Utilize data gathered from an assessment to understand the memory requirements for various use cases. Table 3 shows the impact of a variety of memory configurations on shared storage. The paging file impacts storage requirements only if a disposable disk is used for the OS temporary files.

MEMORY ALLOCATED	SWAP SIZE	PAGING FILE (DEFAULT WIN OS) – DISPOSABLE DISK ONLY	SUSPENDED MEMORY	MEMORY-RELATED CONTRIBUTION TO STORAGE – PER DESKTOP
512MB	512MB	512MB	512MB	1.5GB
1,024MB	1,024MB	1,024MB	1,024MB	3GB
2,048MB	2,048MB	2,048MB	2,048MB	6GB
4,096MB	4,096MB	4,096MB	4,096MB	12GB

Table 3. Impact of Various Memory Configurations on Shared Storage

NOTE: The storage total does not take into account the additional storage required by the Virtual SAN Number of Failures to Tolerate (FTT) policy, which by default doubles the storage utilization (FTT = 1).

To capture data points required for accurate sizing when modeling designs around existing desktop infrastructures, the ideal approach is to conduct an assessment of the infrastructure. The use of third-party tools such as Liquidware Labs, Lakeside Software, and Windows Performance Monitor can be of great assistance in capturing resource utilization requirements.

If an existing desktop infrastructure is not available for evaluation, the other solution is to build new virtual desktops that run all required software and components and to monitor their resource utilization.

Virtual SAN Datastore Sizing Considerations

Storage considerations are among the most important elements for successful View virtual desktop infrastructure designs. With Virtual SAN, View can be used to leverage a storage system composed of locally attached server-side storage.

The Virtual SAN storage solution is designed to deliver virtual machine–centric storage capabilities through the SPBM framework. It is important to fully understand how differing SPBM capabilities affect the consumption of storage in Virtual SAN with View.

The storage capabilities focus on availability, capacity, and performance characteristics and are delivered in virtual machine storage policies. The following are Virtual SAN 5.5 storage polices that affect sizing:

- Number of Failures to Tolerate (FTT)
  - This defines the number of host, disk, or network failures a storage object can tolerate. For "n" failures tolerated, "n+1" copies of the object are created and "2n+1" hosts contributing storage are required.
  - The Number of Failures to Tolerate capability has the greatest impact on capacity in a Virtual SAN cluster. Based on the availability requirements of a virtual machine, the setting defined in a virtual machine storage policy can lead to the consumption of as much as four times the virtual machine or individual disk capacity.
  - Consider setting FTT = 0 for the OS disk in linked-clone floating pools or for full-clone nonpersistent desktops. If Virtual SAN experiences a failure, only the nonpersistent data will be lost.
- Object space reservation (percentage)
  - This is the percentage of the logical size of the storage object that will be reserved (thick provisioned) upon virtual machine provisioning. The rest of the storage object is thin provisioned.
  - For full-clone desktops, consider setting object space reservation = 100 percent to allow for deterministic placement of desktops by the Virtual SAN algorithm based on the full potential desktop capacity. This will guarantee that components are evenly balanced across the Virtual SAN cluster as increased capacity is utilized by the desktops.
- Flash read cache reservation (percentage)
  - This is the flash capacity reserved as read cache for the storage object, specified as a percentage of the logical size of the object.
  - Consider setting a 10 percent read cache reservation on the replica disk in linked-clone floating pools, because this one object will serve the majority of read operations for the desktop pool.
- Number of disk stripes per object
  - This is the number of HDDs across which each replica of a storage object is distributed. Higher values might result in better performance when there are significant numbers of read cache misses. However, in a datastore with many active workloads, this effect might be negated due to the HDDs' being busy with other tasks.
  - The default stripe = 1 setting is recommended for all View desktop pool types, based upon performance analyses.

NOTE: The recommended policies described in this section are set by default in Horizon 6.0.

Objects

Each object that resides on the Virtual SAN datastore comprises multiple components, which are distributed across hosts that are members of a Virtual SAN cluster. Objects are assigned storage performance and availability services requirements through virtual machine storage policies.

The number of virtual desktop objects is calculated based on the number of devices configured per virtual desktop. Virtual SAN objects include virtual machine home, virtual machine swap file, VMDKs, and snapshots. The default number of objects per virtual desktop type is listed in Table 4. The objects marked “shared” have only a single instantiation, which is then shared across all desktops in the pool; the rest are created on a per-desktop basis.

The default number of objects and components is based on the virtual desktop pool type provisioned; it differs between Horizon View 5.3.1 and VMware Horizon™ 6.0 due to contrasting default SPBM policies, as indicated in the following table:

USER ASSIGNMENT	VIRTUAL MACHINE TYPE	DISPOSABLE DISK	OBJECTS	NUMBER OF OBJECTS PER DESKTOP	NUMBER OF COMPONENTS PER DESKTOP (HORIZON VIEW 5.3.1)	NUMBER OF COMPONENTS PER DESKTOP (HORIZON 6.0)
Floating	Linked clone	N	Namespace, VMDK, snapshot, internal, swap	5	9 replica + 15 per virtual machine	9 replica + 9 per virtual machine
Floating	Linked clone	Y	Namespace, VMDK, snapshot, internal, disposable, swap	6	9 replica + 18 per virtual machine	9 replica + 10 per virtual machine
Dedicated	Linked clone	Y	Namespace, VMDK, snapshot, internal, disposable, persistent, swap	7	9 replica + 21 per virtual machine	9 replica + 21 per virtual machine
Floating	Full clone	-	Namespace, VMDK, swap	3	9	7
Dedicated	Full clone	-	Namespace, VMDK, swap	3	9	9

Table 4. Horizon View Virtual Desktop Types and Objects

NOTE: Adjusting the default SPBM policies will alter the component counts listed in Table 4.

Components

The number of objects per virtual desktop, in addition to their performance and availability policy settings, dictates the number of components that will be created per desktop. Virtual SAN 5.5 currently supports a maximum of 3,000 components per host, so it is important to size the Virtual SAN datastore not only for capacity but also for component count.

The following formula can be used to calculate the number of components per virtual machine. It accounts for the replicas and witnesses created, based on the Number of Failures to Tolerate setting. The resulting number of components is split across all the hosts in the cluster.

- *Formula: Object x [ftt x 2 + 1] = y*

After the number of components per desktop has been calculated, it is multiplied by the number of desktops in a pool. This provides the total number of components in a desktop pool. Because all linked-clone desktops are composed of and share the same replica disk, this does not contribute to the total component count in a desktop pool.

Component Calculation Examples with Number of Failures to Tolerate = 1

- Floating linked clones with data redirection: 4 x (1 x 2 + 1) = 12 components per desktop
- Floating linked clones with no data redirection: 3 x (1 x 2 + 1) = 9 components per desktop
- Floating full clones: 3 x (1 x 2 + 1) = 9 components per desktop
- Dedicated linked clones with data redirection: 5 x (1 x 2 + 1) = 15 components per desktop
- Dedicated linked clones with no data redirection: 3 x (1 x 2 + 1) = 9 components per desktop
- Dedicated full clones: 3 x (1 x 2 + 1) = 9 components per desktop

NOTE: If the Number of Disk Stripes per Object policy is increased beyond the default value of 1, each stripe is a separate component. In this scenario, however, the recommended number of disk stripes for virtual desktops is kept at the default value of 1, so it does not affect the calculation.

Network Sizing Considerations

The following are network recommendations specific to View on Virtual SAN nodes:

- Dedicated redundant 10GbE uplinks for the Virtual SAN VMkernel ports are recommended for hosts within solutions with View on Virtual SAN.
- End-to-end jumbo frames are also recommended for View on Virtual SAN. Jumbo frames decrease the CPU utilization of the network stack, in turn increasing the potential consolidation ratio on hosts in which CPU utilization becomes a bottleneck.

Disk Group Sizing

In Virtual SAN, disk group configurations are limited to a maximum of one flash-based device per disk group. If hosts that participate in a Virtual SAN cluster contain more than one flash-based device, multiple disk groups are created to utilize the additional devices. Disk group sizing recommendations differ depending on the type of View pool deployed.

When creating disk groups, users decide on the ratio of flash-based devices to magnetic disks whenever performance is a requirement. The higher the ratio of flash-based–device capacity to magnetic-disk capacity, the greater the size of the cache layer.

Having the ability to define and reduce storage failure domains is one reason to consider a configuration with multiple disk groups. If a flash-based device fails, all the magnetic disks within that particular disk group become unusable, and the total storage capacity provided by the affected disk group becomes unavailable to the cluster. When working with multiple disk groups, the failure domain is limited to only the magnetic disks in that particular disk group.

Another reason for configuring more than one disk group per cluster is to accelerate rebalancing and resynchronization of components. These operations occur after a device or host goes offline or fails, in addition to when the utilization of the HDDs exceeds 80 percent on a host. Having multiple disk groups per host enables greater parallelism in these operations.

Magnetic Disk Sizing

The following are sizing recommendations based on the type of View pool deployed in a Virtual SAN cluster:

- Floating, linked-clone virtual desktops
  - Recommended to have at least three 10,000RPM or 15,000RPM SAS magnetic disks with a disk group
- Dedicated linked-clone virtual desktops
  - Recommended to have at least three 10,000RPM or 15,000RPM SAS magnetic disks with a disk group
- Dedicated full-clone virtual desktops
  - Recommended to have at least four 10,000RPM SAS or 7,200RPM NL-SAS magnetic disks within each disk group. Multiple disk groups will scale both performance and capacity.
- All desktop types: The recommended total capacity of all magnetic disks should be greater than 130 percent of the total size of data to be stored on Virtual SAN. This ensures approximately 30 percent of free space for future growth. When an individual magnetic disk is 80 percent utilized, Virtual SAN automatically rebalances components to other disks, which incurs a performance overhead on the cluster.

Flash Capacity Sizing

In Virtual SAN, 30 percent of each flash-based device is used as a write-back buffer. Each write first goes onto the flash-based device; it is persisted in as many flash-based devices as there are replicas for the corresponding objects—VMDKs, for example. This rule is strictly enforced to ensure the required availability of data.

Virtual SAN uses 70 percent of each flash-based device as read cache. A block is never placed in the read cache of more than one flash-based device. Unlike with the write cache, Virtual SAN maintains the same percentage of the flash-based device as read cache for an object, irrespective of the number of replicas. In other words, increasing the availability of an object does not increase the usage of the read cache of the flash-based device.

The general rule for sizing flash capacity for Virtual SAN with View depends on the type of View pool deployed. The optimal value can be determined by performing an accurate assessment of the particular environment.

- Floating, linked-clone virtual desktops
  - Flash capacity should be greater than 10 percent of the projected virtual machine storage usage, not including replicas created by Virtual SAN for tolerating failures
- Dedicated linked-clone virtual desktops
  - Flash capacity should be greater than 10 percent of the projected virtual machine storage usage, not including replicas created by Virtual SAN for tolerating failures
- Dedicated full-clone virtual desktops
  - Flash capacity should be greater than or equal to 10 percent of the projected virtual machine storage usage, not including replicas created by Virtual SAN for tolerating failures (taking into account object space reservation = 100 percent)

Sizing Example

A user plans to provision 1,000 dedicated full-clone virtual desktops, each with 40GB of logical address space and 4GB of memory. Table 5 shows a simple sizing scenario based on the general recommendation for flash capacity.

MEASUREMENT REQUIREMENTS	VALUES
Virtual machine space usage	40GB
Projected number of virtual machines	1,000
Total space consumption per virtual machine	40GB x 1,000 = approximately 40,000GB = approximately 40TB
Target flash capacity percentage	10 percent
Total flash capacity required	40TB x .10 = 4TB

Table 5. Flash Capacity Sizing

The aggregate anticipated consumed storage, before replication, is 1,000 x 40GB = 40TB. If the virtual machine's availability factor is defined to support *Number of Failures to Tolerate* equals 1 (FTT = 1), this configuration results in creating two replicas for each virtual machine—that is, a little more than 80TB of consumed capacity, including replicated data. However, the flash sizing for this case is 10 percent x 40TB = 4TB of aggregate flash capacity in the cluster where the virtual machines are provisioned.

Given the supported maximum of 100 desktops per host, a minimum of 10 hosts is necessary to support 1,000 desktops, not taking into account sizing for VMware vSphere High Availability (vSphere HA). This means that each host requires a minimum of 400GB of flash.

- *Formula: [total flash capacity required] / [total desktops required / maximum number of desktops per host]*
- *Example: 4TB / (1,000/100) = 400GB of* flash per host

With 10 hosts and 40TB of virtual machine space usage, each host requires 4TB of usable capacity. Because this is a dedicated full-clone View pool, a FTT = 1 setting is utilized, meaning an additional 4TB of Virtual SAN replicas per host will be created to provide storage availability.

To achieve the minimum recommended 130 percent of used capacity, 12 x 900GB HDDs are required for each host. Given a recommended minimum flash capacity of 400GB per host, two 400GB SSDs are used, both to support two disk groups of one SSD and six HDDs and to provide additional flash capacity for future growth through a scaling up of disk group size.

To calculate the usable capacity of each host, we must also consider Virtual SAN overhead and the virtual machine swap space. Virtual SAN utilizes approximately 1GB of overhead per magnetic disk for metadata. Virtual SAN virtual machine swap space is equal to the memory allocated to the virtual machine, minus memory reservation. A Virtual SAN virtual machine swap is always stored in two replicas, regardless of the Number of Failures to Tolerate setting.

NOTE: The maximum number of vSphere HA protected desktops in a datastore is 2,048; the maximum number of total desktops in a cluster is 3,200.

Horizon View and Virtual SAN Node Design Example

View designs utilize the concept of pods as a building block model for View solutions. Pods are described as a set of management tools or virtual desktops that service the number of users defined as part of a virtual desktop design. Each building block is configured to support a specific purpose.

The two types of building blocks are management and desktop.

- A management building block consists of a VMware cluster that supports server workloads for View. This includes all vSphere and View management components and tools infrastructure management.
- A virtual desktop building block consists of several vSphere clusters that support the View virtual desktop workload. This building block also supports View transfer servers for local mode users.

This section provides examples of nodes that can be used to construct a virtual desktop building block based on View and Virtual SAN.

VIEW AND VIRTUAL SAN – NODE PROFILES		
	Linked Clones	**Full Clones**
NUMBER OF VIRTUAL MACHINES PER NODE* (AVERAGE VIRTUAL MACHINE SIZE)	100 or fewer (Two vCPUs; 1.5GB vRAM; 20GB VMDK)	100 or fewer (Two vCPUs; 1.5GB vRAM; 40GB VMDK)
IOPS PER NODE (30% READ 70% WRITE)	10K or fewer	10K or fewer
RAW STORAGE CAPACITY PER NODE	1.2TB	10.8TB
CPU**	2x 10-core	2x 10-core
MEMORY	256GB	256GB
HDD	4x 300GB SAS 15,000RPM	12x 900GB SAS 10,000RPM
SSD	1x 400GB SSD (Class E)	2x 400GB SSD (Class E)
I/O CONTROLLER***	Queue depth: 512 or greater (example: LSI SAS 2208-based controller)	Queue depth: 512 or greater (example: LSI SAS 2208-based controller)
NETWORK ADAPTER	10GbE (jumbo frames enabled)	10GbE (jumbo frames enabled)

DESKTOP TYPE[1]		
	Linked Clones	**Full Clones**
INFRASTRUCTURE SIZING ASSUMPTION	• Average virtual machine instance size: two vCPUs; 1.5GB vRAM; 20GB VMDK • Memory utilization: 70% • SSD-to-HDD ratio: 10% or less anticipated used capacity • Storage utilization: 90% • Disk group ratio: one SSD, three to seven HDDs • Five or fewer disk groups per node • VMware ESXi™ boot: At least 4GB SD card or one dedicated HDD • VMDK size assumes floating pool	• Average virtual machine instance size: two vCPUs; 1.5GB vRAM; 40GB VMDK • Memory utilization: 70% • SSD-to-HDD ratio: 10% or less anticipated used capacity • Storage utilization: 70% • Disk group ratio: one SSD, four to seven HDDs • Five or fewer disk groups per node • VMware ESXi boot: At least 4GB SD card or one dedicated HDD • VMDK size assumes persistent desktops

Table 6. Example View Node Configuration

As seen in Table 6, infrastructure sizing depends on the type of VDI workload. For desktops that are more CPU intensive, consider using additional cores per host. For desktops that require more RAM, consider using greater RAM per host. Table 7 provides some average resource requirements for typical desktop profiles. Use this as a guide to adjust the baseline sizing previously provided.

	CPU	MEMORY	IOPS
Task User (Light)	One vCPU	1GB RAM	3–7
Knowledge Worker (Medium)	One vCPU	1GB RAM	8–16
Power User (Standard)	One vCPU	2GB RAM	17–25
Power User (Heavy)	Two vCPUs	3GB RAM	25+

Table 7. Average Resource Requirements for Typical Desktop Profiles

1. For more details on storage sizing for View, consult *Storage Considerations for VMware Horizon View 5.2:* http://www.vmware.com/files/pdf/view_storage_considerations.pdf.

Conclusion

VMware Virtual SAN is a hypervisor-converged platform that delivers a shared datastore by combining compute and storage resources of VMware vSphere hosts in a vSphere cluster while providing a much simpler storage management experience for the user. The combination of peformance, scalability, and operational simplicity makes Virtual SAN an ideal storage platorm for Horizon View.

Acknowledgments

We would like to thank Charu Chaubal, group manager of the Storage and Availability Technical Marketing team, and Donal Geary, reference architect engineer, for their contributions and for reviewing this paper.

About the Authors

Rawlinson Rivera is a senior architect in the Cloud Infrastructure Technical Marketing group at VMware. His focus is on storage virtualization, software-defined storage technologies, and integration aspects of VMware products and solutions with the OpenStack framework. Previously he was an architect in the VMware Cloud Infrastructure and Management Professional Services organization, focused on vSphere and cloud enterprise architectures for VMware Fortune 100 and 500 customers.

Rawlinson is among the first VMware Certified Design Experts (VCDX #86) and is the author of multiple books based on VMware and other technologies.

Follow Rawlinson's blogs:

- http://blogs.vmware.com/vsphere/storage
- http://www.punchingclouds.com

Follow Rawlinson on Twitter:

- @PunchingClouds

Wade Holmes, CISSP, CCSK, is a senior technical marketing architect at VMware, currently focusing on software-defined storage reference architectures. Wade has more than 17 years of industry experience in the design and implementation of complex computing environments of all scopes and sizes. He has presented at many industry conferences and is a co-author of *VMware vCloud Architecture Toolkit*. Wade is among the first VMware Certified Design Experts, VCDX #15. He holds a bachelor's degree in information technology and a master's degree in information assurance.

Follow Wade's blogs:

- http://blogs.vmware.com/vsphere/storage
- http://www.vwade.com

Follow Wade on Twitter:

- @wholmes

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com