VMware Horizon View 6.0.2 and VMware Virtual SAN 6.0 Hybrid

REFERENCE ARCHITECTURE

vmware®

Table of Contents

Executive Summary

This is a reference architecture using VMware Horizon® View™ 6.0.2 running on VMware Virtual SAN™ 6.0 in a hybrid configuration and is based on realistic test scenarios, user workloads, and infrastructure system configurations. The architecture is comprised of SuperMicro rack mount servers with local storage to support a scalable and cost-effective VMware Horizon View linked-clone desktop deployment on VMware vSphere® 6.0.

Extensive user experience and operations testing, including use of Login VSI desktop performance testing of up-to 1,600 desktops, desktop provisioning operations of up-to 2,400 desktops, revealed world-class performance at an extremely low cost. VMware Virtual SAN technology allows easy scalability while maintaining superior performance at a competitive price point.

Figure 1 shows the overall test results.

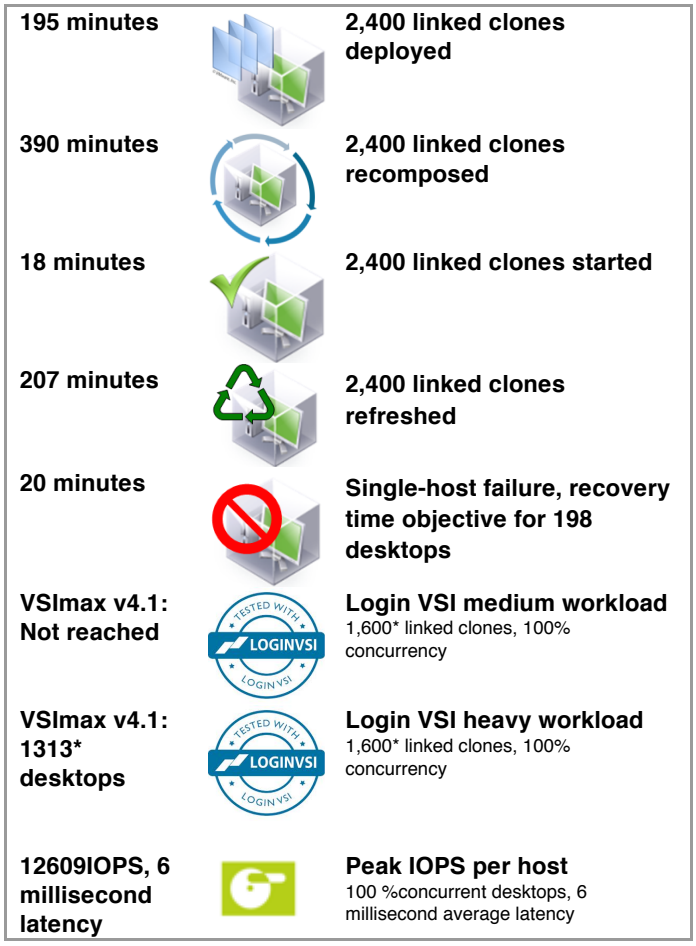

Figure 1. Test Results

Outstanding Desktop Performance

- Performance similar to high-end physical PCs and all-flash arrays
- Excellent user experience characteristics, even with diverse use cases and burst- heavy workloads

Fast Desktop Operations

- User desktop pools provisioned quickly
- Enterprise use cases up and running the same day as requested
- New desktops, applications, and hot fixes rolled out within a single maintenance window

Ease of VDI Management

- Simplified storage configuration and provisioning
- Storage Policy Based Management (SPBM) integration
- Reduced complexity of scaling for VDI deployment
- Tight integration with other vSphere features (high availability, DRS, and others)

*** Workloads push the limits of cost-effective hardware, particularly CPU**

VMware Reference Architecture Overview

VMware reference architectures are built and validated by VMware and supporting partners. They are designed to address common use cases; examples include enterprise desktop replacement, remote access, business process outsourcing, and disaster recovery. A reference architecture describes the environment and workload used to simulate realistic usage, and draws conclusions based on that particular deployment.

This guide is intended to help customers—IT architects, consultants, and administrators—involved in the early phases of planning, design and deployment of Horizon View–based solutions. The purpose is to provide a standard, repeatable, and highly scalable design that can be easily adapted to specific environments and customer requirements.

The reference architecture “building block” approach uses common components to minimize support costs and deployment risks during the planning of large-scale, Horizon View–based deployments. The building block approach is based on information and experiences from some of the largest VMware deployments in production today. While drawing on existing best practices and deployment guides pertinent to many of the individual specific components, the reference architectures are tested and validated in the field and described in detail.

Some key features that can help an organization get started quickly with a solution that integrates easily into existing IT processes and procedures include:

- Standardized, validated, readily available components
- Scalable designs that allow room for future growth
- Validated and tested designs that reduce implementation and operational risks
- Quick implementation, reduced costs, and minimized risk

Technology Introduction

This reference architecture uses common components to minimize support costs and deployment risks.

The desktop virtualization solution, which combines the best of breed of data center, virtualization, and network technologies, uses SuperMicro rack mount servers with local solid state drives (SSD) and hard disk drives (HDD) running the vSphere 6.0 software suite for desktop workloads. In addition, standard SuperMicro servers running vSphere are used for server workloads. The View 6.0.2 environment runs Windows 7 virtual desktops provisioned by VMware View Composer™.

The Virtual SAN storage platform for desktop workloads allows the solution to scale linearly, with each host capable of supporting approximately 200 users per host. This reference architecture shows 2,400 desktops running on 12 VMware ESXi™ hosts. Login VSI workloads are very resource intensive, easily pushing the limits of cost-effective hardware, particularly CPU that is used in this solution. Therefore, 1,600 desktops are tested under Login VSI workloads.

Hardware Components

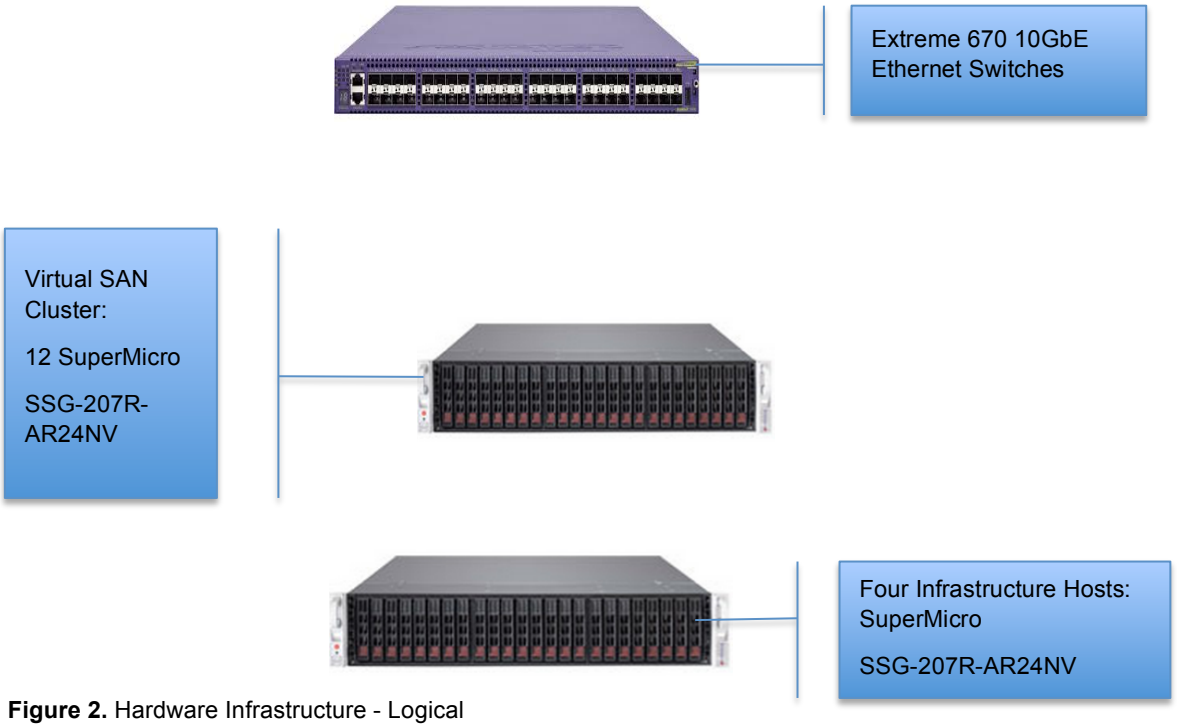

Figure 2. Hardware Infrastructure - Logical

For management workloads, the solution uses four standard SuperMicro SSG-207R-AR24NV rack mount servers.

Desktop workloads use 12 SuperMicro SSG-207R-AR24NV rack mount servers, which offer high-density memory, balanced I/O, and the latest processors for enterprise virtualization and business-processing environments. The system is optimized for running in virtualized and cloud-computing environments.

The Extreme Networks Summit X670-G2 product family provides high density 10 Gigabit Ethernet and 40

Gigabit Ethernet switching. The switches provide 10GbE network connectivity for management, Virtual SAN, and desktop traffic.

Local SSDs and HDDs are used in conjunction with Virtual SAN technology to provide a scalable and enterprise-class storage solution. Each ESXi host has two disk groups each consisting of one SSD and six HDDs. The disk groups are combined to form a Virtual SAN datastore. This next-generation storage platform combines powerful and flexible hardware components with advanced efficiency, management, and software-defined storage.

VMware vSphere

vSphere is the industry-leading virtualization platform for building cloud infrastructures. It enables users to run business-critical applications with confidence and respond quickly to business needs. vSphere accelerates the shift to cloud computing for existing data centers and underpins compatible public cloud offerings, forming the foundation for the industry's best hybrid cloud model.

VMware Virtual SAN

Virtual SAN is a hypervisor-converged, software-defined storage platform that is fully integrated with vSphere. Virtual SAN aggregates locally attached disks of hosts that are members of a vSphere cluster to create a distributed shared storage solution. Because Virtual SAN sits directly in the I/O data path, it can deliver the highest levels of performance, scalability, and resilience without taxing the CPU with additional overhead. Virtual SAN enables the rapid provisioning of storage within VMware vCenter™ during virtual machine creation and deployment operations.

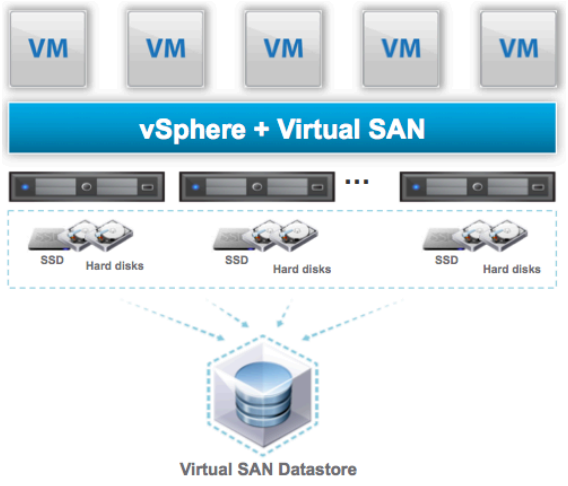

Figure 3. Virtual SAN Clustered Datastore

Virtual SAN uses a hybrid disk architecture that leverages flash-based devices for performance and magnetic disks for capacity and persistent data storage. Its distributed datastore is an object-store file system that leverages the vSphere Storage Policy-Based Management feature to deliver centrally managed, application-centric storage services and capabilities. Administrators can specify storage attributes, such as capacity, performance, and availability, as a policy on a per virtual machine basis. The policies dynamically self-tune and load-balance the system so that each virtual machine has the right level of resources.

Virtual SAN simplifies and streamlines storage provisioning and management for vSphere environments. Use virtual machine-centric storage policies to provide finely granular control and automation of storage service levels. Self-tuning capabilities automatically rebuild and rebalance storage resources to align with the service levels assigned to each VM. Full integration with vSphere and the entire VMware stack delivers an efficient and cost-effective operational model.

VMware Horizon View

Horizon View brings the agility of cloud computing to the desktop by transforming desktops into highly available and agile services delivered from your cloud. View delivers virtual sessions that follow end users across devices and locations. It enables fast and secure access to corporate data across a wide range of devices, including Mac operating system, Windows, and Linux machines and iOS and Android tablets.

You can use View with VMware vCenter Server™ to create desktops from virtual machines that are running on ESXi hosts and to deploy these desktops to end users. After you create a desktop, authorized end users can use Web-based or locally installed client software to connect securely to centralized virtual desktops, back-end physical systems, or terminal servers. View uses your existing Active Directory infrastructure for user authentication and management.

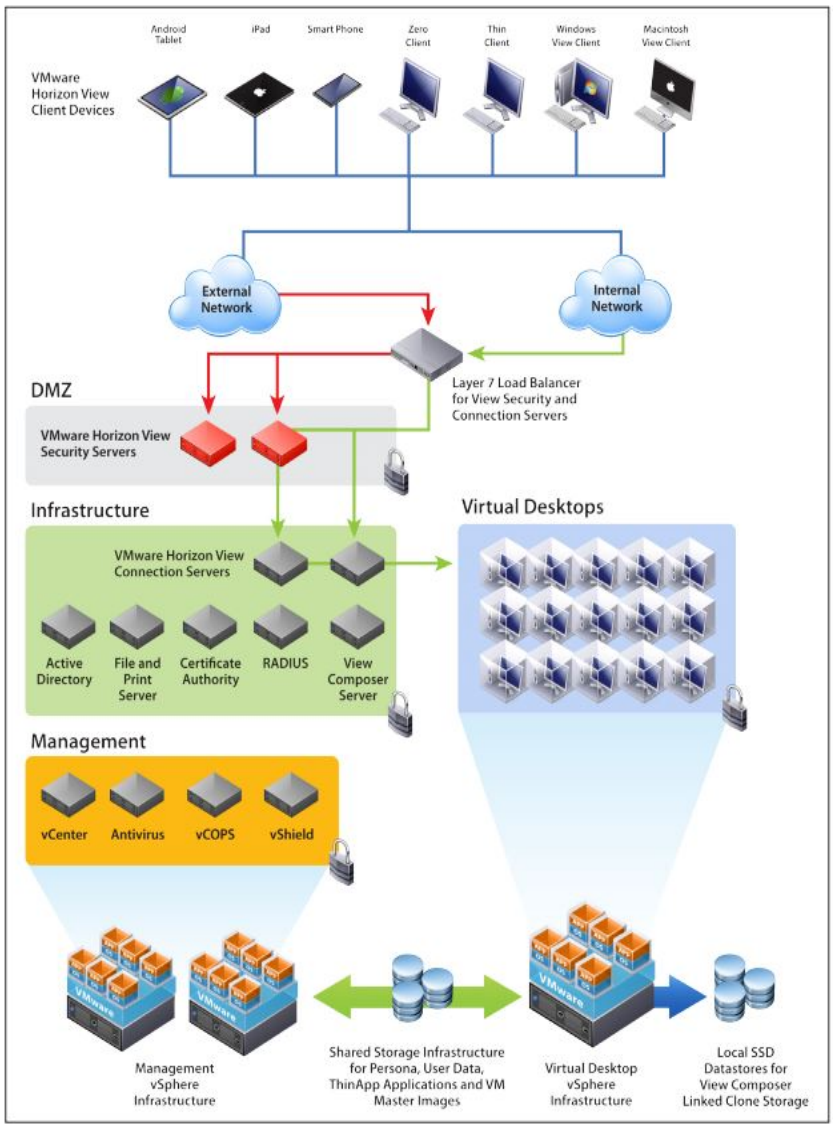

Figure 4. Horizon View Components

VMware View Storage Accelerator

All tests performed in this reference architecture used View Storage Accelerator, an in-memory host caching capability that uses the Content-Based Read Cache (CBRC) feature in ESXi hosts. CBRC provides a per host RAM-based solution for View desktops, considerably reducing the read I/O requests that are issued to the storage layer. It also addresses boot storms when multiple virtual desktops are booted at once—which causes a large number of reads. CBRC is beneficial when administrators or users load applications or data frequently.

View Storage Accelerator minimizes total cost of ownership (TCO) in View deployments by reducing peak input/output operations per second (IOPS) by 80 percent and peak throughput up to 65 percent.

Test Results

This section summarizes the test results of the solution.

Login VSI 4.1 Workload Testing

The test used Login VSI 4.1 to load the system with simulated desktop workloads using common applications like Microsoft Office, Internet Explorer, and Adobe Reader.

The VDI workload in general can be CPU intensive. Virtual SAN can support up to 200 desktops per host from storage perspective if host CPU is sized properly. During the LoginVSI testing, we uncovered that our servers were CPU bound under specific workloads. Therefore, we focused our tests on 1,600 desktops to observe Virtual SAN performance: One test applied a medium workload and the other a heavy workload, both with 100 percent concurrency.

Note: *VMware does not recommend host CPU utilization to exceed 80 percent.*

Test 1: 1,600 Medium-Workload Linked-Clone Desktops

In Test 1, the average host CPU usage reached above 95 percent, as shown in Figure 5, on all ESXi hosts at 1,600 desktops under medium workload with 100 percent concurrency. Despite high CPU usage, VSImax v4.1 was not reached.

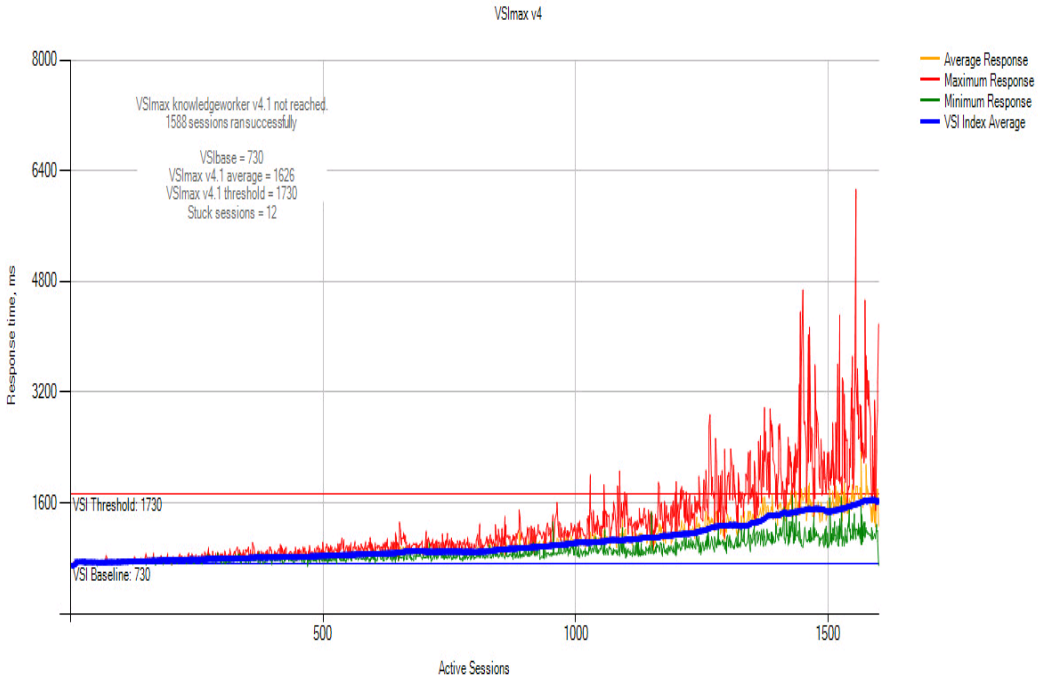

Figure 5. VSImax Not Reached on Login VSI Medium Workload, 1,600 Desktops

See VSImax introduction for more information.

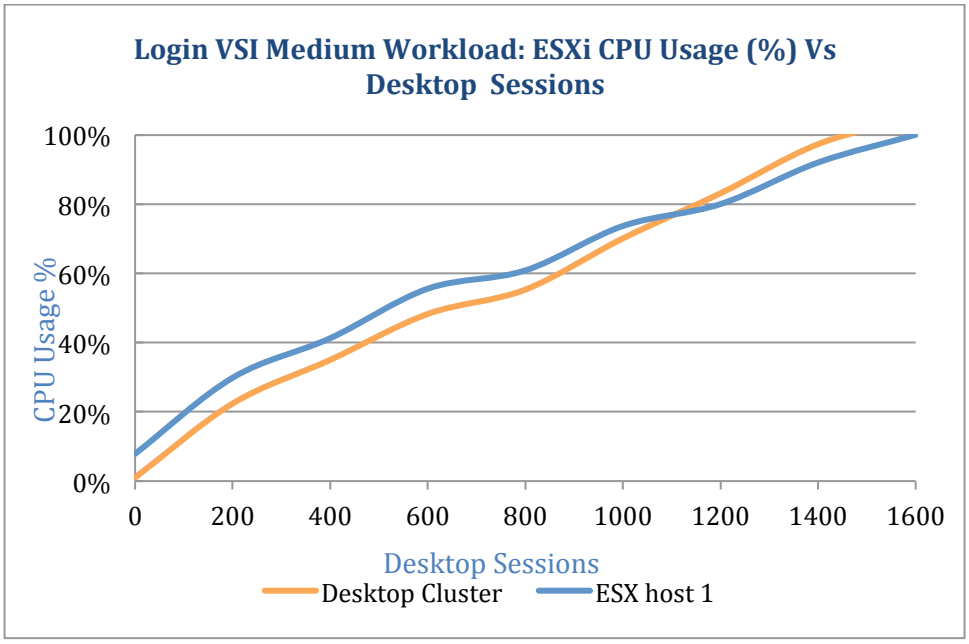

Figure 6. Desktop Cluster CPU Usage during Login VSI Medium Workload

The ESXi average latency is 3.3 milliseconds (ms) and peaked at just over 8 ms during an I/O-intensive phase that reached over 8,700 IOPS.

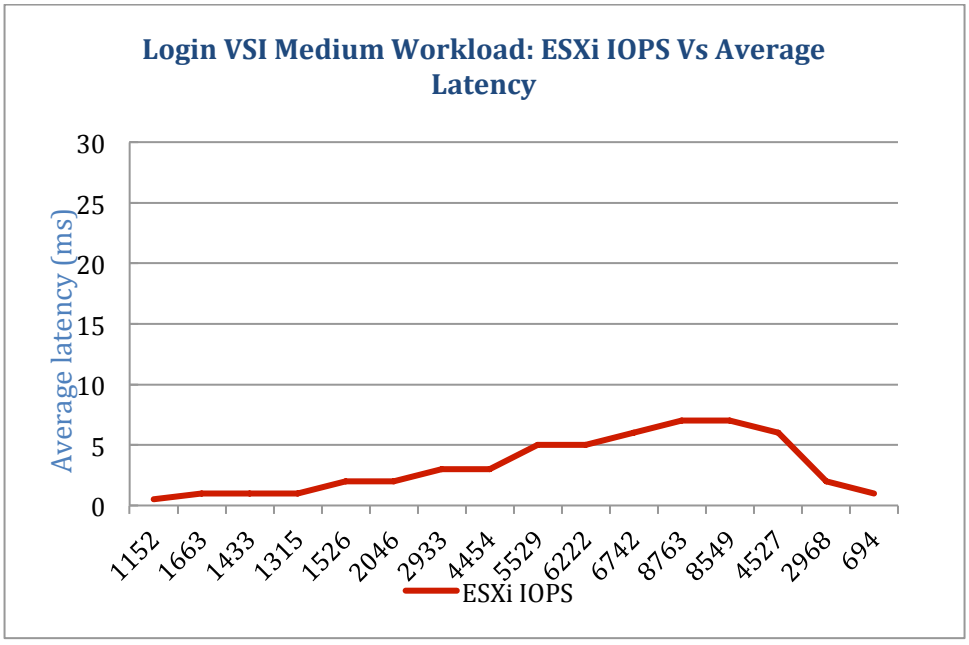

Figure 7. ESXi IOPS versus ESXi Latency during Login VSI Medium Workload

Figure 8 shows ESXi CBRC hit rate during Login VSI medium workload testing. The average hit rate is above 80 percent.

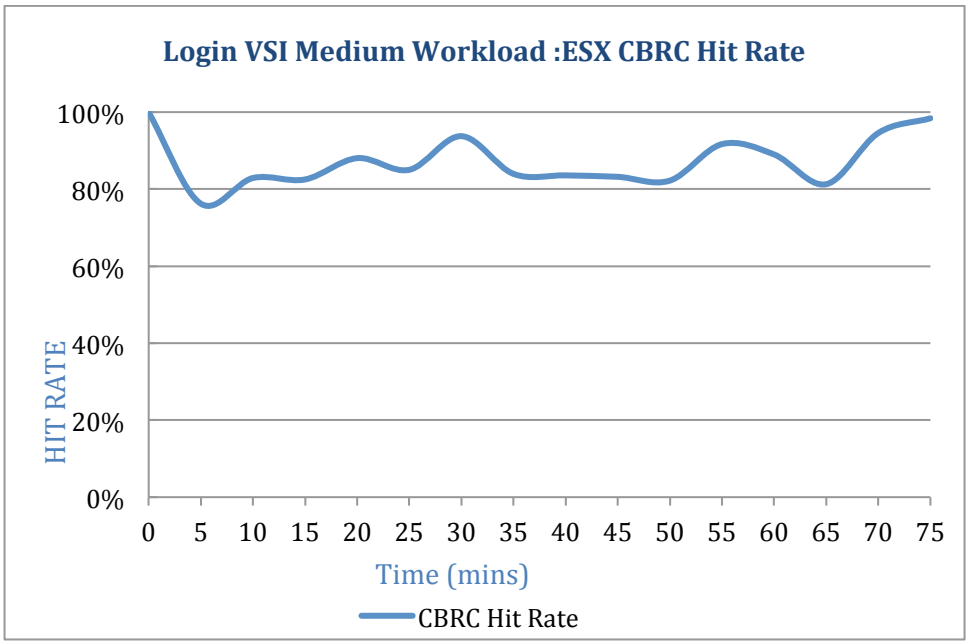

Figure 8. ESXi CBRC Hit Rate during Login VSI Medium Workload

The highlights of test 1 are:

- VSImax v4.1 did not reach the baseline of 730
- CPU usage was high but the memory usage remained under 60 percent
- Excellent average latency on ESXi, even at high load (average 3.3 ms latency, peak 8ms at over 8700 IOPS per ESXi host)
- Peak of 104,800 IOPS on Virtual SAN datastore (51percent writes and 49 percent reads)

CPU Usage		63,299MHz
Memory Usage		260.92GB
Network Adapter		XMIT 56385 / RCV 78039KBps
Storage Adapter		8763 IOPS

Figure 9. ESXi Host Metrics – Test 1

Test 2: 1,600 Heavy-Workload Linked-Clone Desktops

In Test 2, the CPU was saturated at 100 percent usage across all ESXi hosts under heavy workload with 100 percent concurrency. VSImax v4.1 is 1,313 at baseline of 733.

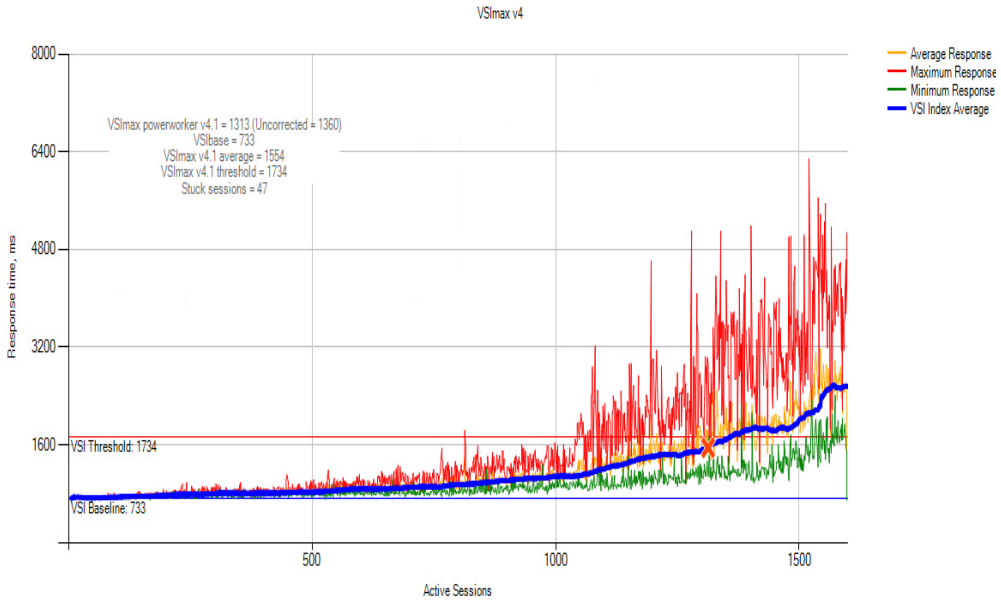

Figure 10. VSImax 1313 during Login VSI Heavy Workload, 1,600 Desktops

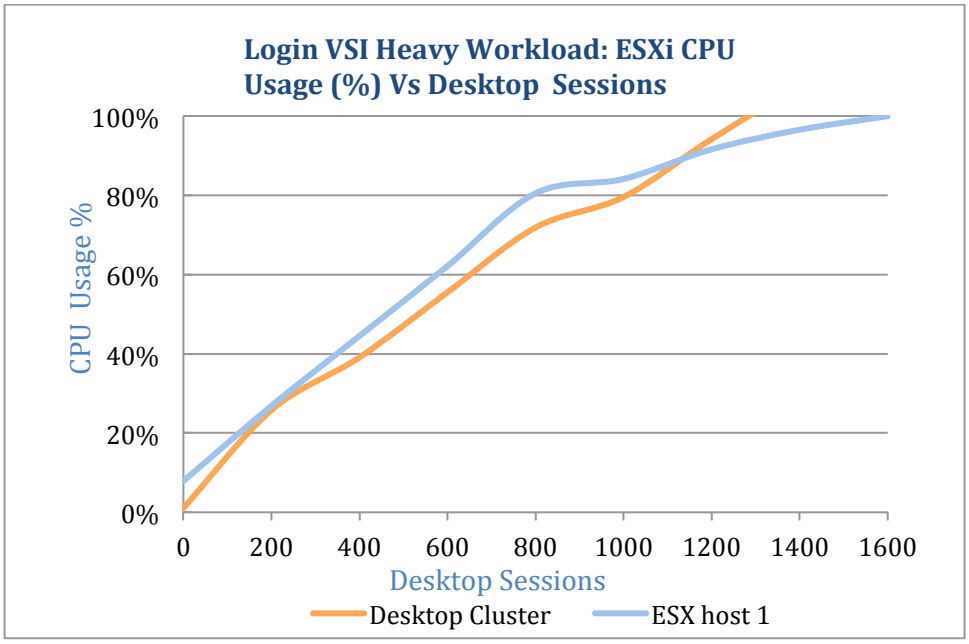

Figure 11. ESXi CPU Usage during Login VSI Heavy Workload

The ESXi latency averaged 4.3 ms and peaked over 9 ms during an I/O-intensive phase that reached over 12,609 IOPS.

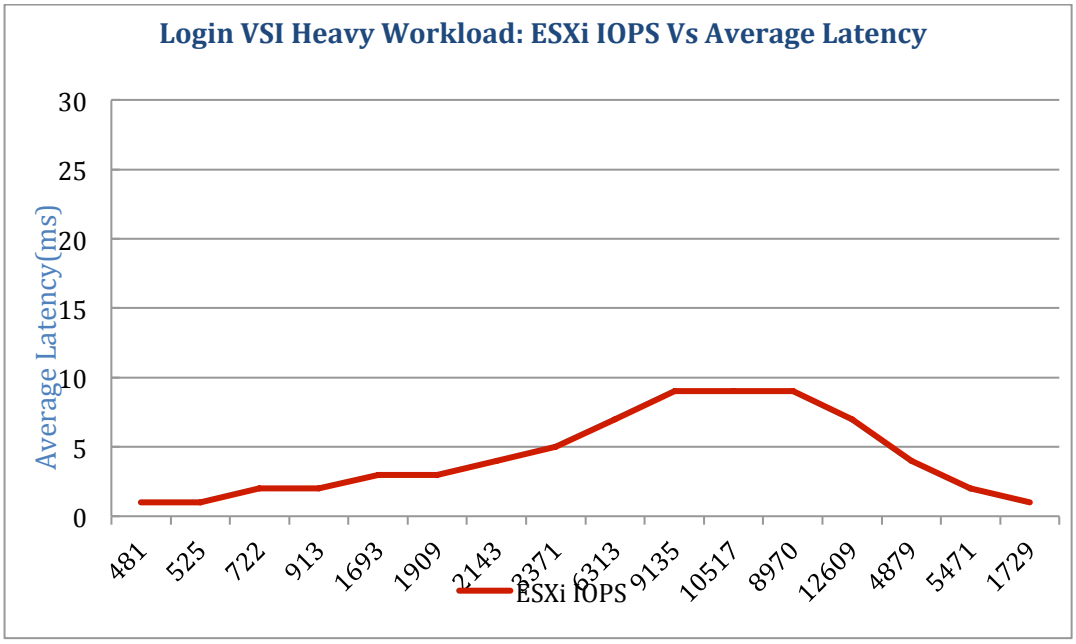

Figure 12. ESXi Average IOPS versus Average Latency during Login VSI Heavy Workload

Figure 13 shows ESXi CBRC hit rate during Login VSI heavy workload testing. The average hit rate is above 70 percent.

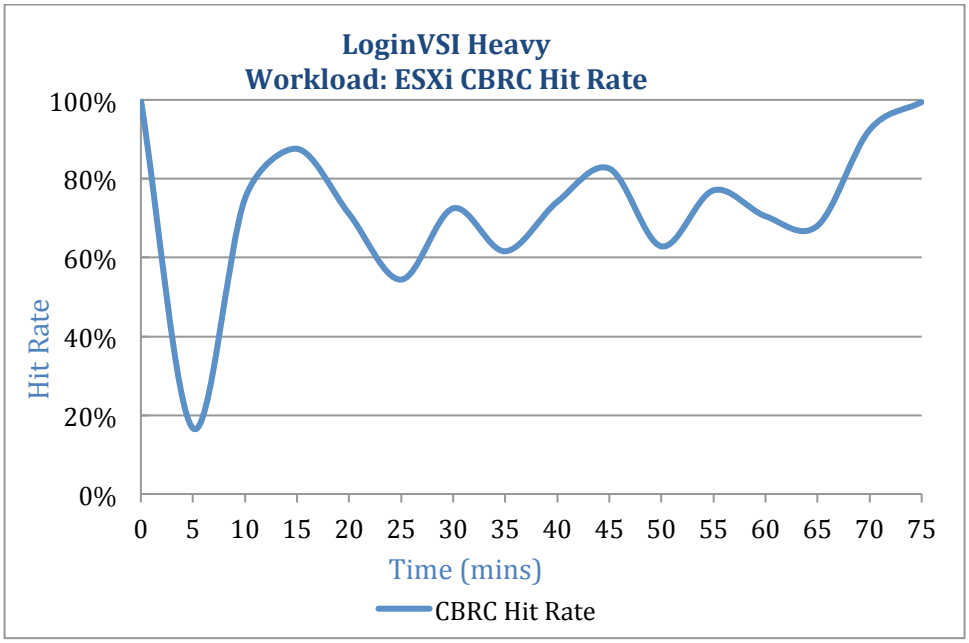

Figure 13. ESXi CBRC hit rate during Login VSI Heavy Workload

The highlights of test 2 are:

- VSImax v4.1 1313 did not reach the baseline of 733
- CPU usage was very high, but memory usage remained under 60 percent
- Excellent average latency on ESXi, even at high load (average 4.3ms, peak 9ms latency at over 12,609 IOPS per ESXi host)
- Peak of 151,308 IOPS on Virtual SAN datastore (51percent writes and 49 percent reads)

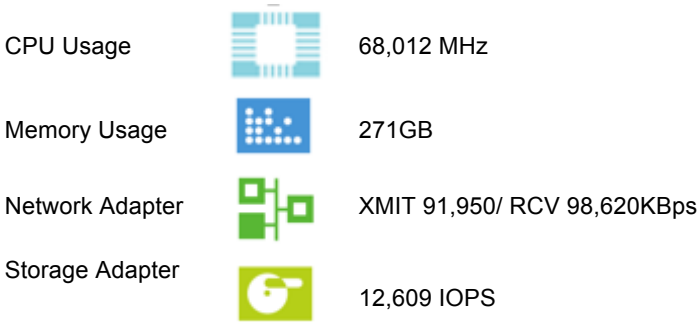

Figure 14. ESXi Host Metrics – Test 2

View Operations Tests

Provisioning 2,400 Linked-Clone Desktops

In this test, a new pool of 2,400 linked-clone virtual desktops was provisioned on the Virtual SAN datastore, with about 200 desktops per ESXi host. To complete this task, View Composer created a replica copy of the 24GB base image on the Virtual SAN datastore. View Composer created and customized the desktops and joined them to the Active Directory domain. It then took a snapshot of the virtual desktop, and the desktop went into an available state.

It took less than 195 minutes to provision 2,400 Windows 7 linked-clone virtual desktops and for them to appear in the available state in the View Administrator console.

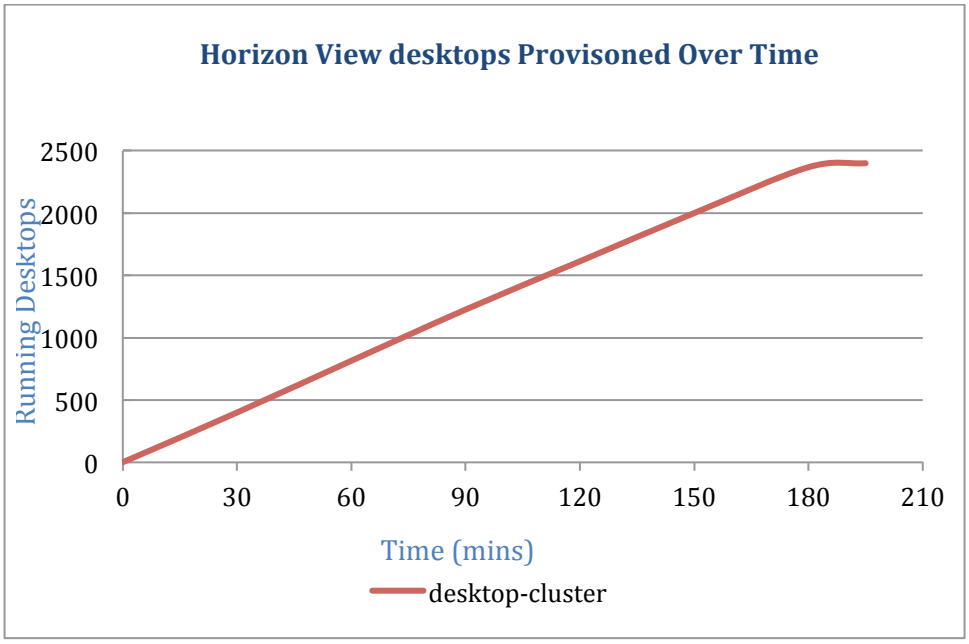

Figure 15. View Provisioning Operation – 2400 Linked-Clone Desktops

Refreshing 2,400 Linked-Clone Desktops

In a refresh operation, a virtual desktop reverted to its snapshot. The operating system disk of each virtual desktop was restored to its original state and disk size.

It took 207 minutes to refresh 2,400 Windows 7 linked-clone virtual desktops to their original-base image.

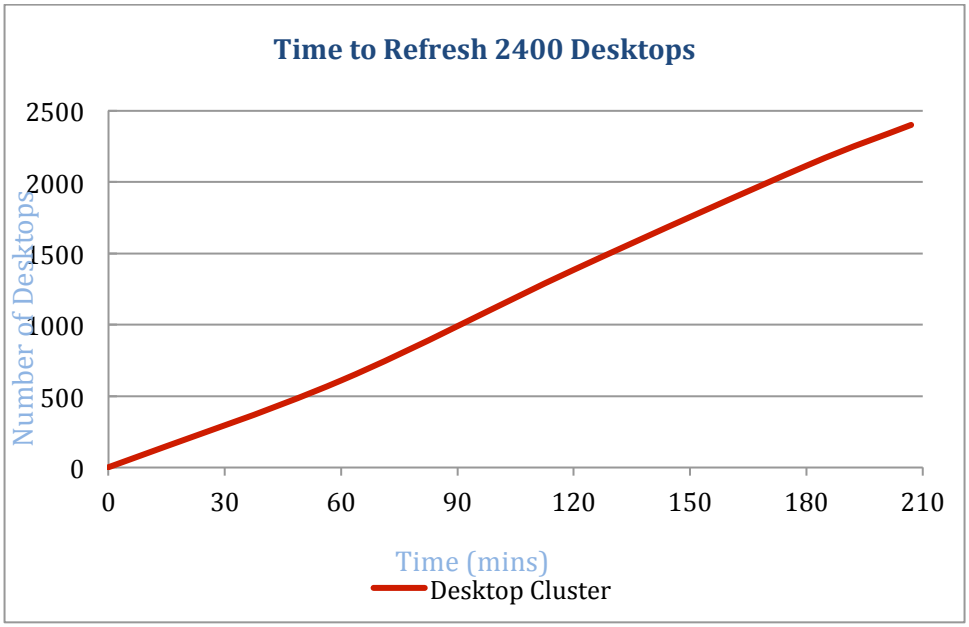

Figure 16. View Refresh Operation – 2,400 Linked-Clone Desktops

Recomposing 2,400 Linked-Clone Desktops

In a recompose operation, a virtual desktop operating system disk was changed to a new base image and snapshot. This feature allows administrators to push out patches and software updates with ease. In this operation, View Composer created a replica of the new base image on the Virtual SAN datastore, created a new operating system disk for each virtual desktop, and deleted the old one. The new desktop was then customized and a new snapshot was created.

It took approximately 390 minutes to recompose 2,400 Windows 7 linked-clone virtual desktops to a fresh-base image.

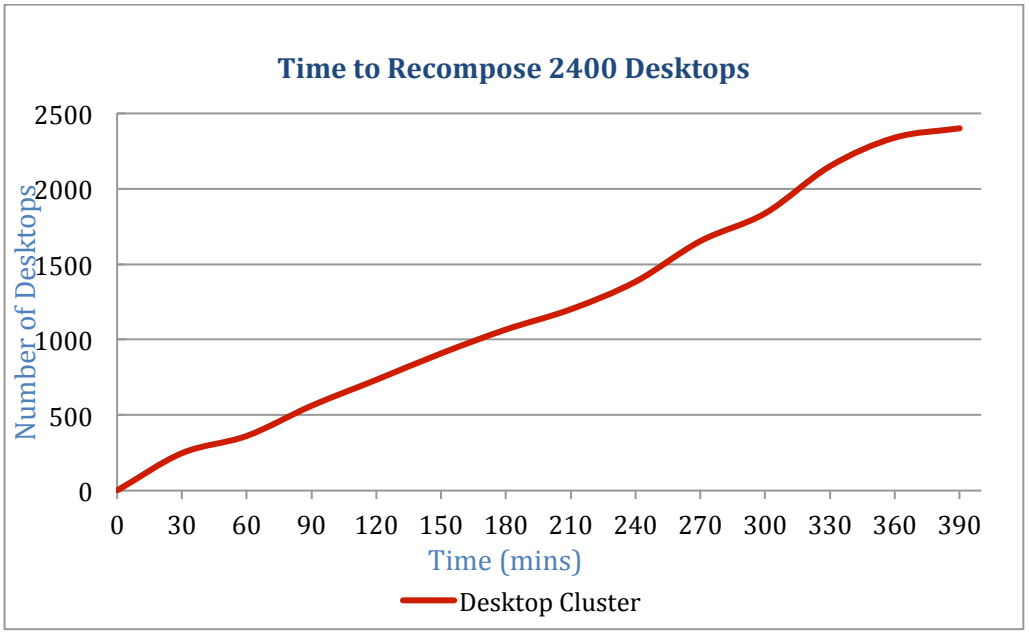

Figure 17. View Recompose Operation – 2,400 Linked-Clone Desktops

Deleting a Pool of 2,400 Linked-Clone Desktops

This test deleted a desktop pool, destroying the associated virtual desktops and replicas. Deleting a pool of 2,400 linked-clone virtual desktops took 195 minutes.

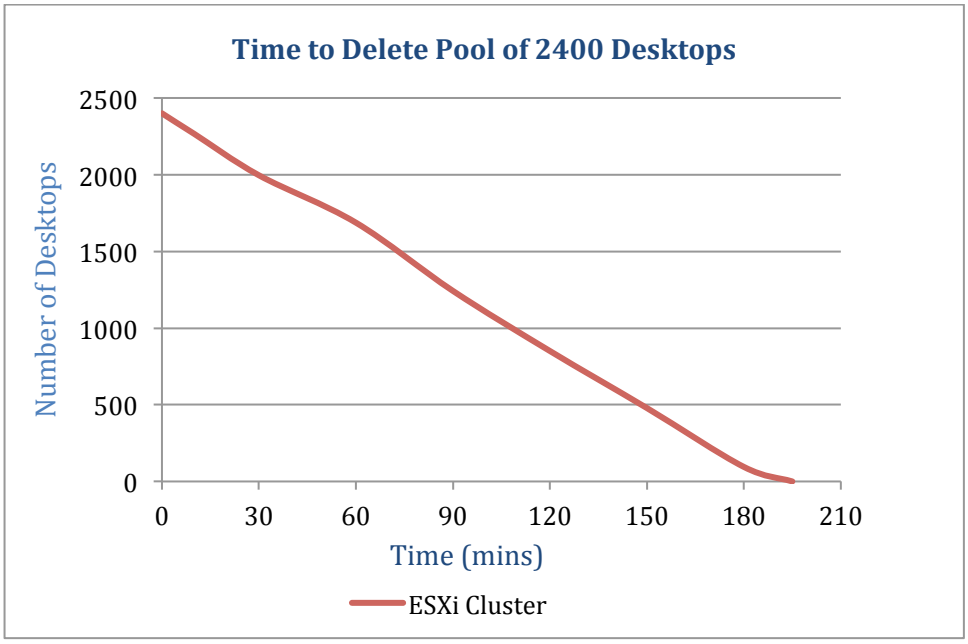

Figure 18. View Pool Deletion Operation –2400 Linked-Clone Desktops

Powering on 2,400 Desktops

The power-on test was carried out on a 12-node Virtual SAN cluster in vCenter. It took just under 18 minutes for all the virtual desktops to be ready for Windows user login. CPU usage on each host was nearly full for about a half hour.

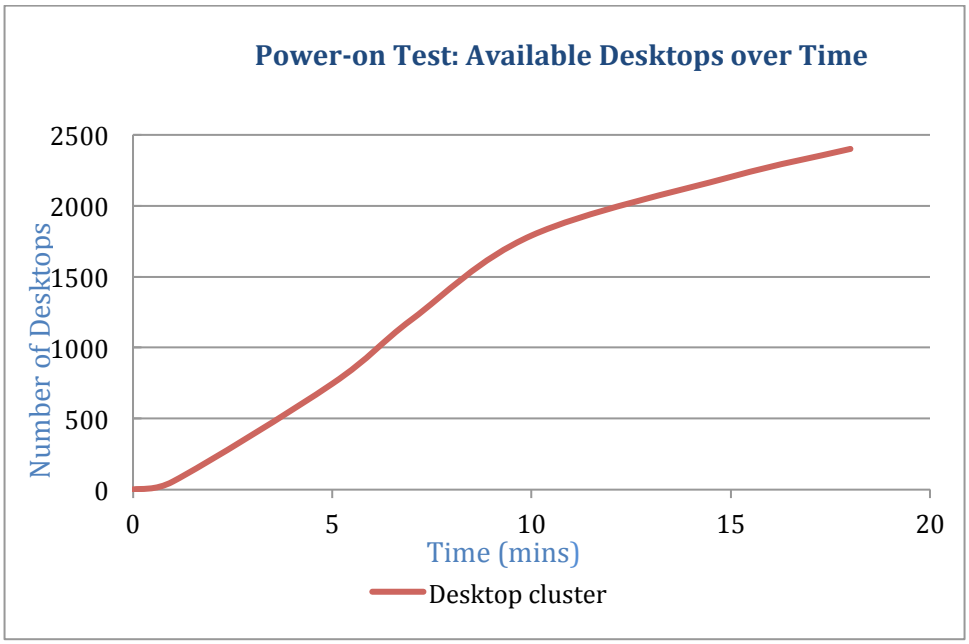

Figure 19. View Power-On Operation – 2,400 Linked-Clone Desktops

Resiliency Test: One-Node Failure

A single Virtual SAN node hardware failure was simulated for a cluster with twelve hosts and 2,400 running virtual desktops, all under simulated workload.

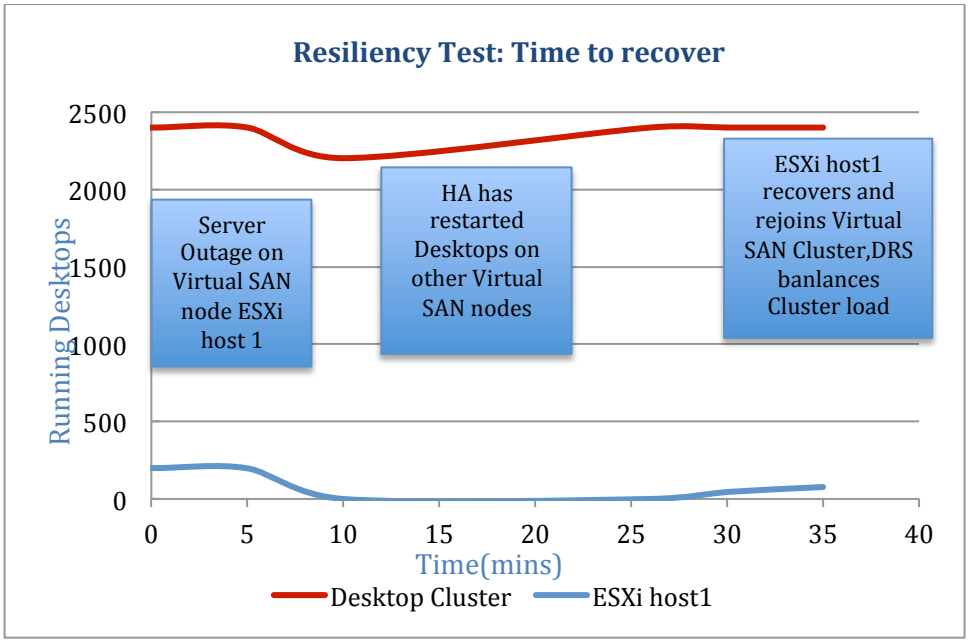

Figure 20. Resiliency Testing – Recovery Time for a Single-Node Failure

An ESXi host with 198 running virtual desktops was reset, and all the virtual machines became unavailable. VMware vSphere High Availability restarted the virtual desktops on the other Virtual SAN cluster nodes, and all were ready for user login within 16 minutes of the simulated failure.

The power was restored to the Virtual SAN node some minutes later. The node rejoined the Virtual SAN cluster, and VMware vSphere Storage Distributed Resource Scheduler™ rebalanced the load across all ESXi hosts in the cluster.

***Note**: A single-node failure does not trigger an immediate rebuild after a host failure is detected. If a failure that returns an I/O error is detected, such as a magnetic disk or SSD, Virtual SAN immediately responds by rebuilding the disk object. However, for host failures that do not return an I/O error, Virtual SAN has a configurable repair delay time (60 minutes by default) after which components are rebuilt across the cluster. Virtual SAN prioritizes the current workload over rebuilding to minimize the impact on the cluster performance.*

System Configurations

This section describes how the reference architecture components were configured.

Architecture

VMware Virtual SAN integrates with the View pod and block design methodology, which comprises the following components:

- **View Connection Server** – A View Connection server supports up to 2,000 concurrent connections. Our testing consisted of two View Connection Server, operating in active/active mode. The two View Connection Servers actively broker and possibly tunnel connections.
- **View block** – View provisions and manages desktops through vCenter. Each vCenter instance supports up to 10,000 virtual desktops. The tests used one vCenter and one Virtual SAN cluster with four hosts. Virtual SAN supports 200 virtual machines per host and clusters of up to 64 hosts.

  Note: *The maximum HA-protected virtual machine in a vSphere cluster is 6,000 per datastore.*

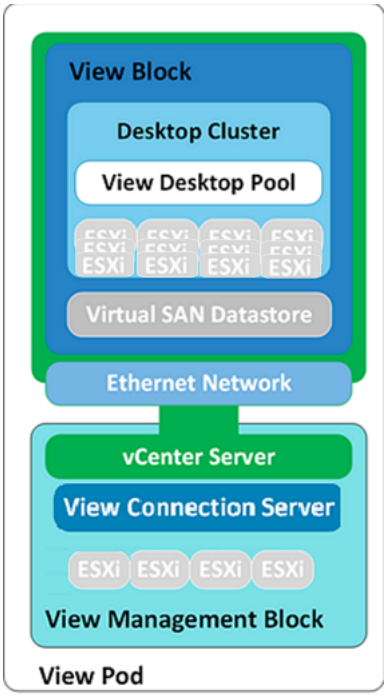

Figure 21. View Pod Configuration

- **View management block** – A separate vSphere cluster was used for management servers to isolate the volatile desktop workload from static server workload. Our testing is for large deployment and we have a dedicated vCenter for the management and View blocks.

vSphere Clusters

A 12-node Virtual SAN cluster was deployed to support 2,400 virtual desktops. Each Supermicro server had an identical configuration and ESXi booted from the local HDD.

For the mmanagement virtual machines, a 4-node cluster was deployed on identical servers.

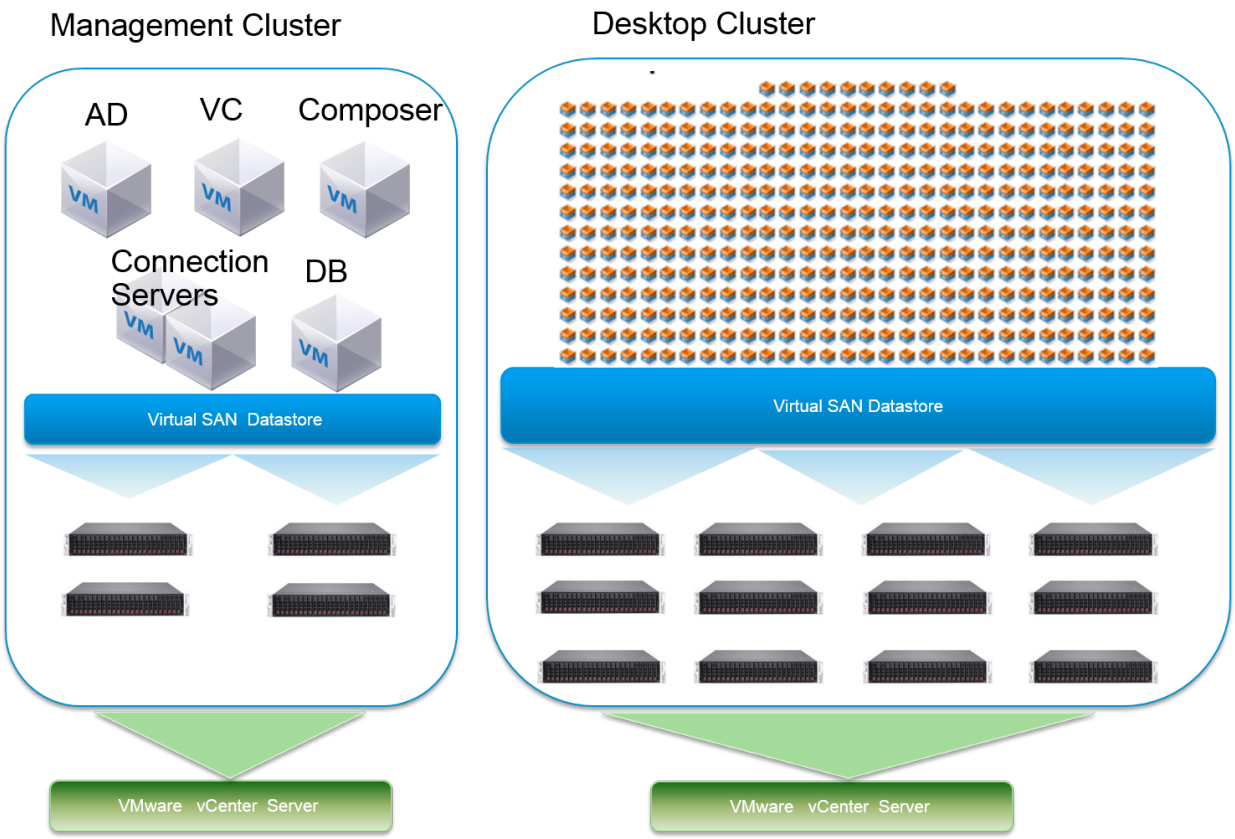

Figure 22. vSphere Cluster Design

Table 1 lists the settings of vSphere Cluster configuration.

Table 1. vSphere Cluster Configuration

PROPERTY	SETTING	DEFAULT	REVISED
Cluster Features	HA	–	Enabled
	DRS	–	Enabled
vSphere HA	Host Monitoring Status	Enabled	–
	Admission Control	Enabled	–
	Admission Control Policy	Host failures the cluster tolerates = 1	–
	Virtual Machine Options > VM restart priority	Medium	–
	Virtual Machine Options > Host Isolation Response	Leave Powered On	–
	Virtual Machine Monitoring	Disabled	–
vSphere DRS	Automation Level	Fully automated (apply 1,2,3 priority recommendations)	–
	DRS Groups Manager	–	–

	Rules	–	–
	Virtual Machine Options	–	–
	Power Management	Off	–
	Host Options	Default (disabled)	–
Enhanced vMotion Capability		Disabled	–
Swapfile Location		Store in the same directory as the virtual machine	–

ESXi Servers

Each Virtual SAN ESXi server in the Virtual SAN cluster had the following configuration settings:

Figure 23. ESXi Host Components

Table 2 lists the settings of ESXi host configuration.

Table 2. ESXi Host Configuration

PROPERTY	SPECIFICATION
ESX server model	12 x SuperMicro Server 2027R-AR24NV
ESX host CPU	2 x Intel Xeon CPU E5-2690 v2 @ 3.00GHz 10-core (60GHz)
ESX host RAM	512GB
ESX version	ESXi 6.0.0, 2494585
Network adapter	Intel Ethernet 82599EB 10-Gigabit SFI/SFP+ Firmware version: 0x8000030d Driver version: ixgbe 3.7.13.7.14iov

Storage adapter	2x LSI00194 SAS 6GB/S 9211-8I HBA W/LSI2008 CONTROLLER Firmware version: 18.00 Driver version: MPT2SAS 19.00.00.00.1vmw Queue Depth: 600
Power management	High Performance (set in BIOS)
Disks	SSD: 2 Intel 800GB S3700 Series SATA 6Gb/s HDD: 12 Seagate 1.2TB 10K RPM 2.5in 6Gbps SAS 64M

Storage Controller Mode

The storage controller LSI9211-8I supports both pass-through and RAID mode. We use pass-through mode in the testing. This mode is the preferred mode for Virtual SAN that gives Virtual SAN complete control of the local SSDs and HDDs attached to the storage controller.

Virtual SAN

The floating linked clones and replicas used Virtual SAN for storage. Each ESXi host had the same uniform configuration of two disk groups, each consisting of one 800GB SSD and six 1.2TB 10K SAS disks. The SSD devices form the caching layer (fixed at a 70 percent read cache and 30 percent write buffer). Only the spinning magnetic disks contribute toward the usable storage capacity of the datastore. The 12 hosts yielded 172.8TB RAW.

Figure 24. Virtual SAN Datastore Components

The virtual desktop replica did not need to be stored on a dedicated tier of flash storage because read I/O operations are served from the flash layer on Virtual SAN.

Virtual SAN Storage Policy

Virtual SAN can set availability, capacity, and performance policies per virtual machine if the virtual machines are deployed on the Virtual SAN datastore. The tests used the default storage policy settings that are created by Horizon View automatically. For Horizon 6.0.2 with View, specific storage-policy recommendations are based on pool type as described in table 16.

Table 3. Virtual SAN Storage Default Settings for View

STORAGE CAPABILITY	SETTING
Number of Failures to Tolerate	1
Number of Disk Stripes per Object	1
Flash Read Cache Reservation	0%
Object Space Reservation	0%

Number of Failures to Tolerate (FTT) – This Virtual SAN storage protection policy is applied to each virtual machine. The FTT policy defines how many concurrent host, network, or disk failures can occur in the cluster and still ensure the availability of the object. The configuration contains at least FTT+1 copies of the virtual machine and a witness copy to ensure that the object's data is available even when the number of tolerated failures occurs.

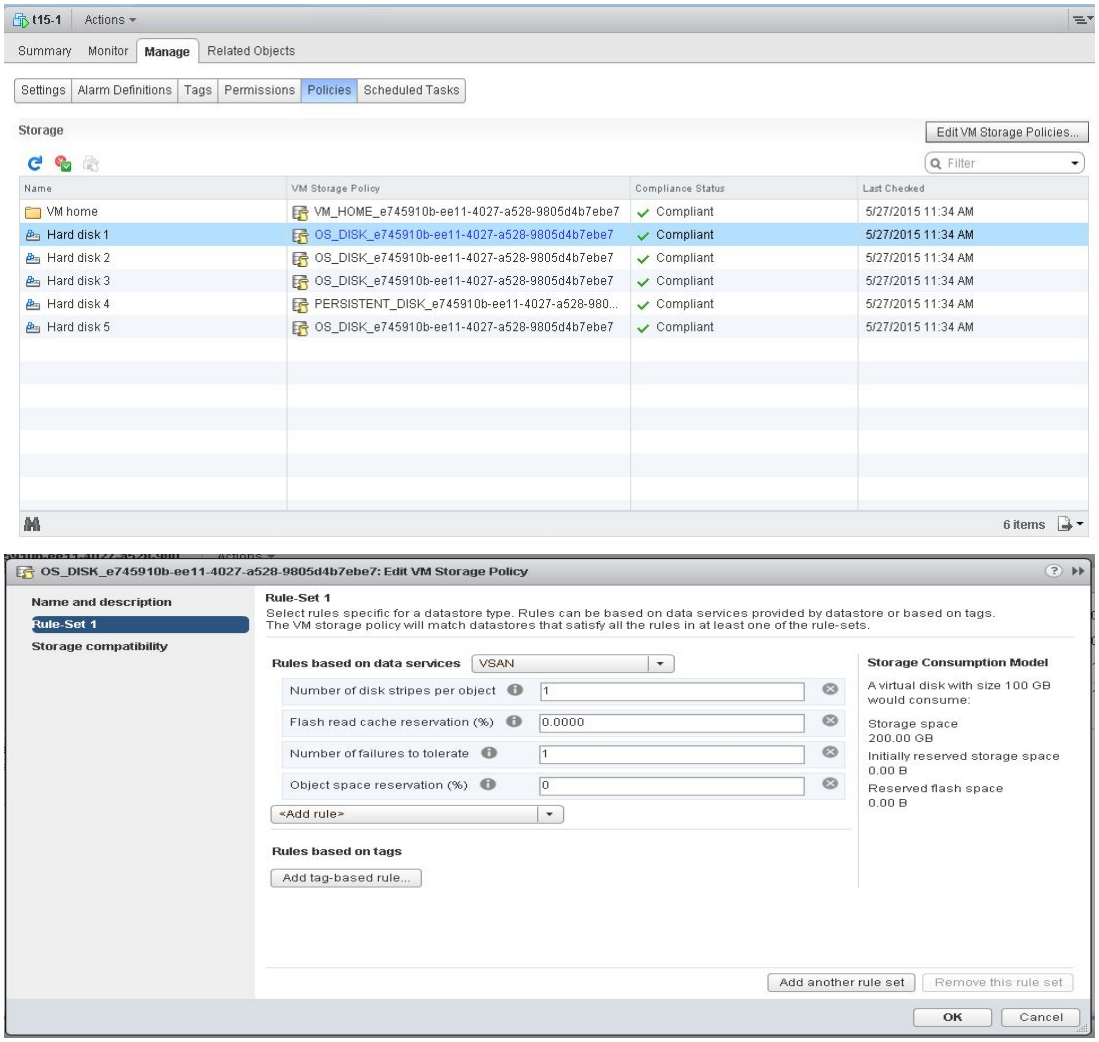

Figure 25. View Auto Created Storage Policy

Object Space Reservation – By default, a virtual machine created on Virtual SAN is thin-provisioned, so it has no capability for object space reservation. It does not consume any capacity until data is written. You can change this setting between 0–100% of the virtual disk size. The virtual machine consumes this capacity of the Virtual SAN datastore when it is created.

The combination of the object space reservation percentage and the FTT settings applied to the virtual machines on the Virtual SAN datastore determines the usable capacity of the datastore.

Number of Disk Stripes per Object – This policy defines how many physical disks across each copy of a storage object are striped. The default value (recommended) of 1 was sufficient for our tested workloads.

Flash Read Cache Reservation – This is the amount of flash capacity reserved on the SSD as a read cache for the storage object. By default, all virtual machines share the read cache of an SSD equally.

Virtual SAN Fault Domains

Fault domains introduce an even higher level of availability in Virtual SAN 6.0. Virtual SAN Fault Domains provide the ability to group multiple hosts within a cluster to define failure domains to ensure replicas of virtual machines data is spread across the defined failure domains (racks). With this new feature, customers

can continue to run their virtual machines on Virtual SAN, even in the event of something catastrophic like a rack failure.

The guideline in tolerating ‘n’ host failures was to have ‘2n + 1′ hosts in the cluster. Similarly, to tolerate ‘n’ domain failures, ‘2n + 1′ fault domains are required.

The fault domain configuration for 12-node Virtual SAN cluster in the testing is shown in table 4:

Table 4. Fault Domain Configuration

FAULT DOMAIN	A	B	C
Host	1, 5, 7,10	2, 4,8, 11	3, 6, 9,12

Networking

A vNetwork distributed switch (dvSwitch) acted as a single vSwitch across all associated hosts in the data center. This setup allows virtual machines to maintain a consistent network configuration as they migrate across multiple hosts. The dvSwitch uses two 10GbE adapters per host.

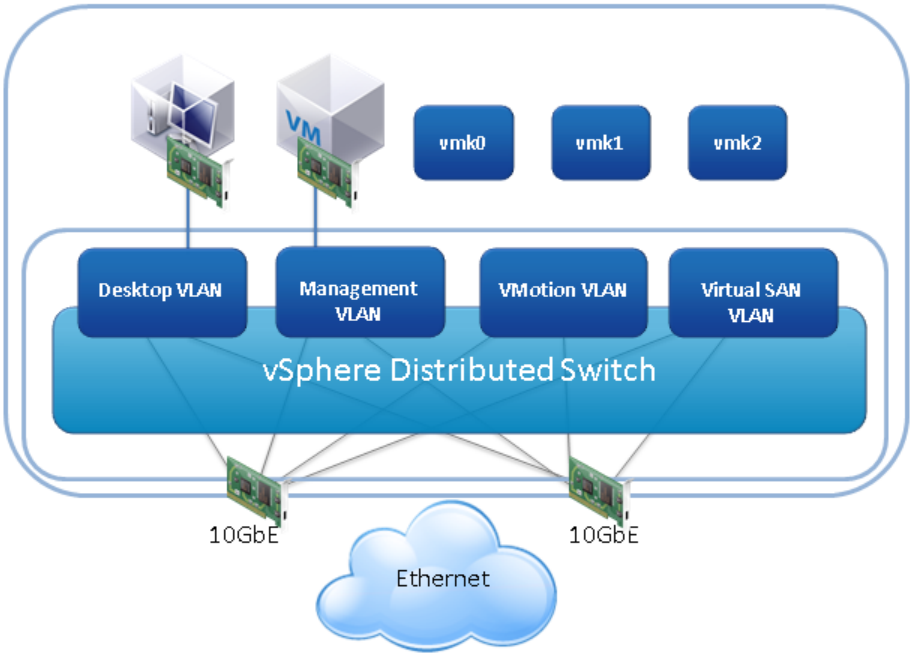

Figure 26. dvSwitch Configuration

Properties regarding security, traffic shaping, and NIC teaming can be defined on a port group. Table 5 shows the settings used with this design.

Table 5. Port Group Properties – dvSwitch v6.0

PROPERTY	SETTING	DEFAULT	REVISED
General	Port Binding	Static	–
Policies: Security	Promiscuous mode	Reject	–
	MAC address changes	Accept	Reject
	Forged transmits	Accept	Reject

Policies: Traffic Shaping	Status	Disabled	–
Policies: Teaming and Failover	Load balancing	Route based on the originating virtual port ID	Route based on physical NIC load
	Failover detection	Caution Link Status only	–
	Notify switches	Yes	–
Policies: Resource Allocation	Network I/O Control	Disabled	Enabled
Advanced	Maximum MTU	1500	–

Network I/O control was enabled for the distributed switch. The settings and share values in table 6 were applied on the resource allocation.

Table 6. Resource Allocations for Network Resources in dvSwitch

NETWORK RESOURCE POOL	HOST LIMIT (MBPS)	PNIC SHARES	SHARES
vMotion	8000Mbit/s	Low	25
Management	Unlimited	Normal	50
Virtual machines	Unlimited	High	100
Virtual SAN Traffic	Unlimited	Normal	50

Horizon View

The Horizon View installation included the following core systems:

- Two connection server (N+1 is recommended for production)
- One vCenter Server (vCenter Appliance) with the following roles:
  - vCenter
  - vCenter single sign-on (SSO)
  - vCenter Inventory Service
- View Composer

Note: *Security servers were not used during this testing.*

View Global Policies

The global policies in table 7 were in place for all system testing.

Table 7. View Global Policies

POLICY FEATURE	SETTING
Multimedia redirection (MMR)	Allow
Remote Mode	Allow
PCoIP hardware acceleration	Allow – medium priority

VMware View Manager Global Settings

The following VMware View Manager™ global policies were used.

ATTRIBUTE	SPECIFICATION
View Administrator session timeout	600 minutes
Forcibly disconnect users	9,999 minutes
Single sign-on (SSO)	Enabled
For Clients that support applications If the user stops using the keyboard and mouse disconnect their applications and discard SSO credentials	Never
Other clients Discard SSO credentials	After 15 minutes
Auto Update	Disabled
Pre-login message	No
Display warning before forced logoff	Yes
Enable Windows Server 2008 R2 desktops	No
Mirage Server configuration	

Table 8. View Manager Global Settings

vCenter Server Settings

View Connection Server uses vCenter Server to provision and manage View desktops. vCenter Server is configured in View Manager.

Table 9. View Manager – vCenter Server Configuration

ATTRIBUTE	SPECIFICATION
Description	View vCenter Server
Connect using SSL	Yes
vCenter Port	443
View Composer Port	18443
Enable View Composer	Yes
Advanced Settings: Maximum Concurrent vCenter Provisioning Operations Maximum Concurrent Power Operations Maximum Concurrent View Composer Maintenance Operations Maximum Concurrent View Composer Provisioning Operations	 20 50 12 12
Storage Settings: Enable View Storage Accelerator Default Host Cache Size	 √ 2048MB

View Manager Pool Settings

Table 10 lists the View Manager pool settings.

Table 10. View Manager – Test Pool Configuration

ATTRIBUTE	SPECIFICATION
Pool Type	Automated Pool
User Assignment	Floating
Pool Definition – vCenter Server	Linked Clones
Pool ID	Desktops
Display Name	Desktops
View folder	/

ATTRIBUTE	SPECIFICATION
Remote Desktop Power Policy	Take no power action
Auto Logoff Time	Never
User Reset Allowed	False
Multi-Session Allowed	False
Delete on logoff	Never
Display Protocol	PCoIP
Allow Protocol Override	False
Maximum Number of Monitors	1
Max resolution	1920 x 1200
HTML Access	Not selected
Flash Quality Level	Do not control
Flash Throttling Level	Disabled
Enable Provisioning	Enabled
Stop Provisioning on error	Enabled
Provision all desktops up-front	Enabled
Disposable File Redirection	Do not redirect
Select separate Datastores for replica and OS	Not selected
Datastores – Storage Overcommit	Conservative
Use View Storage Accelerator	Selected
Reclaim VM disk space*	N/A
Disk Types	OS disks
Regenerate Storage Accelerator after	7 days
Reclaim VM Disk Space	N/A
Use Quickprep	Enabled

Test Methodology

This reference architecture used Login VSI to test performance and operational latency.

Login VSI 4.1 Workload Testing

Login Virtual Session Indexer (Login VSI) is the industry-standard benchmarking tool for measuring the performance and scalability of centralized desktop environments. Login VSI gradually increases the number of simulated users until saturation. When the system is saturated, the response time of the applications increases significantly. This latency indicates that the system is almost overloaded. Nearly overloading a system makes it possible to find out what its true maximum user capacity is.

After a test is performed, the response times can be analyzed to calculate the maximum active session per desktop capacity. This metric is called VSImax. When the system is approaching its saturation point, response times rise. By reviewing the average response time, you can see that the response time escalates at the saturation point.

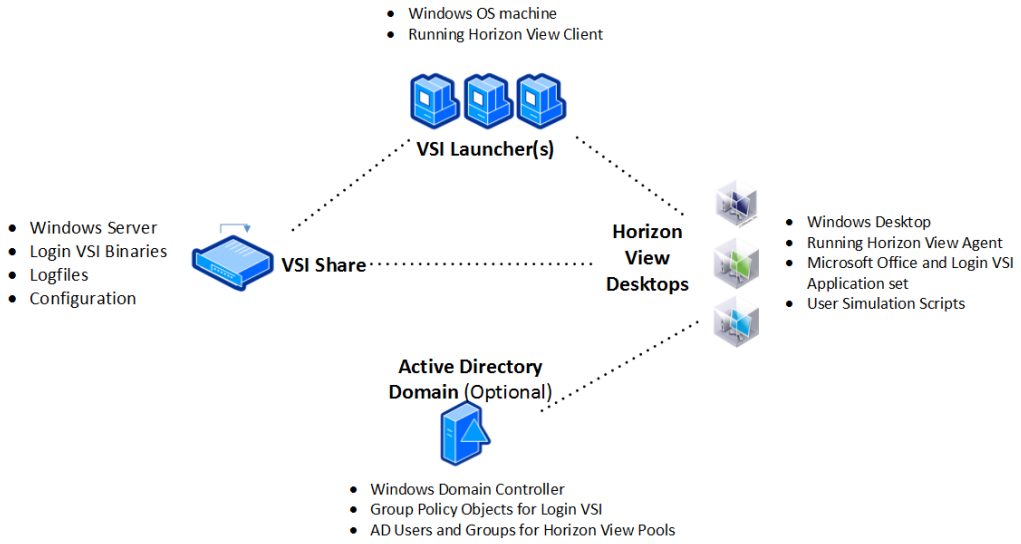

Figure 27. Login VSI System Components

The simulated desktop workload is scripted in a 48-minute loop when a simulated Login VSI user is logged on performing generic office worker activities. After the loop is finished, it restarts. Within each loop, the response times of 12 operations are measured in a regular interval: 12 times within each loop. The response times of these operations are used to determine VSImax.

The operations from which the response times are measured are listed in table 11.

Table 11. Login VSI Operations

ID	ACTION	DESCRIPTION	RELATED RESOURCES
WSLD	Start Microsoft Word and load a random document	Word Start/Load a local random document file from content pool	CPU, RAM, and I/O
NSLD	Start VSI-Notepad and load a document	VSI-Notepad Start/Load a local random text file from content pool	CPU and I/O
WFO	Press file open in VSI-Notepad	VSI-Notepad file open [Ctrl+O]	CPU, RAM, and I/O
NFP	Press print open in VSI-Notepad	VSI-Notepad print open [Ctrl+P]	CPU
ZHC	Compress files with high compression	Compress a local random .pst file from content pool (5MB)	CPU
ZNC	Compress files with no compression	Compress a local random .pst file from content pool (5MB)	I/O

Login VSI has built-in workloads, so you can immediately start testing. See the following brief descriptions of the medium and heavy built-in workloads.

Medium Workload

The Login VSI default workload is medium. It emulates a medium knowledge worker using Microsoft Office, Internet Explorer, PDF files, Java, and FreeMind. The medium workload is designed to run on two vCPUs per desktop virtual machine. It has these characteristics:

- After a session starts, the workload loop repeats every 48 minutes.
- The loop is divided in four segments. Each consecutive Login VSI user login starts at a different segment to ensure that all elements in the workload are equally used throughout the test.
- During each loop, the response time is measured every three to four minutes.
- Five applications are opened simultaneously.
- The keyboard type rate is 160ms for each character.
- Approximately two minutes of idle time is included to simulate real-world users.

Each loop opens and uses:

- Outlook to browse messages
- Internet Explorer, browsing different Web pages; a YouTube style video (480p movie trailer) is opened three times in every loop
- Word, one instance to measure response time, one instance to review and edit a document
- Doro PDF Printer and Acrobat Reader, printing and reviewing PDF files
- Excel, opening a large randomized sheet
- PowerPoint to review and edit a presentation
- FreeMind, a Java-based mind-mapping application

Heavy Workload

The heavy workload requires the Login VSI PRO Content library, which includes 720p and 1080p videos.

In addition to what is performed in the medium workload, the heavy workload includes the following characteristics:

- Begins by opening four instances of Internet Explorer and the instances open throughout the workload loop.
- Begins by opening two instances of Adobe Reader and the instance remain open throughout the workload loop.
- More PDF printer actions.
- A 720p video and a 1080p video are watched.
- Increases the time for playing a Flash game.
- Idle time is reduced to two minutes.

Virtual Machine Test Image Build

Table 12 lists the based image configuration. The configuration is conformed to the testing tool standards and is optimized in accordance with the VMware Horizon with View Optimization Guide for Windows 7 and Windows 8. The VMware OS Optimization Tool was used to make the changes.

Table 12. Virtual Machine Test Images

ATTRIBUTE	LOGIN VSI IMAGE
Desktop OS	Windows 7 Enterprise SP1 (32-bit)
Hardware	VMware Virtual Hardware version 8
CPU	2
Memory	1536MB
Memory reserved	0MB
Video RAM	35MB
3D graphics	Off
NICs	1
Virtual network adapter 1	VMXNet3 Adapter
Virtual SCSI controller 0	Paravirtual
Virtual disk – VMDK 1	24GB
Virtual disk – VMDK 2	1GB
Virtual disk – VMDK 3	
Virtual floppy drive 1	Removed
Virtual CD/DVD drive 1	Removed
Applications	Adobe Acrobat 11 Adobe Flash Player 16 Doro PDF 1.82 FreeMind Internet Explorer 9 MS Office 2010
VMware Tools™	9.0.10 build-2445092
VMware View Agent	6.0.2.-2331487

System Sizing

This reference architecture used the following sizing specifications.

Hosts

As part of the sizing calculations, it is important to factor in the CPU and memory overhead of Virtual SAN. Virtual SAN is designed to introduce no more than ten percent of CPU overhead per host.

Table 13. Host Sizing – Desktop CPU Requirements

DESKTOP PERFORMANCE METRIC	RECORDED VALUE
Average number of CPUs per physical desktop system	1
Average CPU utilization per physical desktop system	350MHz
vCPU overhead	10%

Table 14. Host Sizing – CPU

ATTRIBUTE	SPECIFICATION
Number of CPUs (sockets) per host	2
Number of cores per CPU	10
GHz per CPU core	3.0GHz
Total CPU GHz per CPU	30GHz
Total CPU GHz per host	60GHz
Virtual SAN CPU usage	10%
Available CPU GHz per host	54GHz
Desktops per host	140

***Note**: 140 is the calculated maximum number.*

It is recommended to allow some headroom for CPU spikes, host failures, and maintenance within your vSphere clusters.

The host memory requirement depends on various factors, including memory allocated per virtual desktop, virtual desktop graphics requirements, vCPU memory overhead, Virtual SAN overhead of 10 percent.

Table 15. Host Sizing – Memory

ATTRIBUTE	SPECIFICATION
Virtual SAN memory usage	10%
Total amount of RAM per virtual machine	1536MB
VM CPU No's	2
VM Memory Reservation	0%
VM Resolution	1 x [1920x1600]
VM Video Memory	8.79MB
3D	-
VM Memory Overhead	63MB
Total amount of RAM per host	248GB

Virtual SAN

Virtual SAN introduces some new constructs that have sizing requirements.

Disk Groups

A disk group is a container for magnetic disks and an SSD that acts as a read cache and write buffer. Each disk group must have one SSD and at least one magnetic disk, with a maximum of seven. Each ESXi host can have up to five disk groups.

The type of magnetic disk depends on the level of performance and capacity required. For linked clones, it is

recommended to have at least three magnetic 10K or 15K disks per disk group.

When an individual magnetic disk is 80 percent utilized, Virtual SAN rebalances the components to other magnetic disks. This action incurs a performance overhead on the cluster.

The recommended sizing for the SSD is 10 percent of the total consumed storage capacity (excluding FTT). For example, a pool of 2400 linked-clone desktops is expected to use 5GB of disk space per virtual desktop:

10% (2400 x 5) = 1,200GB

With a Virtual SAN cluster of twelve hosts, the minimum recommended SSD size is 100GB per host.

Only magnetic disks count toward cluster capacity. Cluster capacity equals:

Num_hosts x Num_disk_groups x Num_disks_per_group x disk_size

For our test configuration, the capacity is:

12 x 2 x 6 x 1.2TB = 172.8 TB

Note: For the 2400 linked clone pool, we did not utilize the total capacity of both disk groups. The required capacity should be based on sizing.

Objects and Components

Virtual SAN objects include the virtual machine home (namespace), virtual machine swap files, virtual machine disk format (VMDK) files, and snapshots. The namespace includes VSA files, log files, and virtual machine configuration files. In Virtual SAN 6.0 we will now account for an additional MEM object when the VMs are suspended. In 5.5 this was created within the namespace component.

The number of objects per virtual machine, in addition to their performance and availability requirements, dictates the number of components that will be created. Virtual SAN supports a maximum of 9,000 components per host.

The default storage policies for Virtual SAN and Horizon View are as follows. VMware recommends using the default policies.

- FTT = 1
- Number of disk stripes per object =1

FTT defines the number of hosts, disk, or network failures that a storage object can tolerate. For *n* failures tolerated, *n*+1 copies of the object are created, and 2*n*+1 host contributing storage is required.

FTT 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 several multiples of the virtual desktop consumed space.

The number of disk stripes per object is the number of HDDs across which each replica of a storage object is distributed.

You can use the following formula to calculate the number of components per desktop. It accounts for the replicas and witnesses created based on the FTT setting. The resulting number of components is split across all the hosts in the cluster.

Number of components = Objects x [FTT x 2 + 1]

In Horizon View 6.0, Horizon View specific Virtual SAN Storage Policies leveraged during provisioning; it uses FTT=0 for Floating Desktops utilize less components per host/cluster than Horizon View 5.3.2, for other types, Default of FTT=1, Disk Striping=1 for other objects.

The default number of objects per virtual desktop type is listed in table 16.

Table 16. Default Number of Objects per View Desktop Disk Type

STORAGE POLICIES	NUMBER OF DISK STRIPES PER OBJECT	FLASH READ CACHE RESERVATION (%)	NUMBER OF FAILURES TO TOLERATE (FTT)	OBJECT SPACE RESERVATION (%)
VM_Home	1	0	1	0
OS_Disk	1	0	1	0
OS_Disk_Floating	1	0	0	0
Persistent_Disk	1	0	1	100
Replica_Disk	1	10	1	0
Full_Clone_Disk	1	0	1	100
Full_Clone_Disk_Floating	1	0	0	100

Table 17. Default Number of Objects per View Desktop Type

USER ASSIGNMENT	VIRTUAL MACHINE TYPE	DISPOSABLE DISK	PERSISTENT DISK	NUMBER COMPONENTS PER DESKTOP
Floating	Linked clone	N	N	9 replica plus 12 per virtual machine
Floating	Linked clone	Y	N	9 replica plus 14 per virtual machine
Dedicated	Linked clone	Y	Y	9 replica plus 27 per virtual machine
Dedicated	Linked clone	N	Y	9 replica plus 21 per virtual machine
Dedicated	Linked clone	Y	N	9 replica plus 24 per virtual machine
Floating	Full clone	–	–	10
Dedicated	Full clone	–	–	12

Note: *Adjusting the default Space Policy-Based Management policies alters the component counts listed here.*

For our test configuration, 2,400 floating, linked-clone virtual machines with disposable disk come to components per desktop. The total number of components for our pool of 2,400 virtual machines is:

2400 x 14 = 33,600

In addition, three objects (namespace, swap, and VMDK) are created for the replica virtual disk. With FTT = 1, that is 9 additional components, so 33,609 in total.

Management Blocks

Table 18. Management Block Sizing

SERVER ROLE	VCPU	RAM (GB)	STORAGE (GB)	OS
Domain Controller	2	6	40	Server 2008 64-bit R2
SQL Server	4	8	140	Server 2008 64-bit R2
vCenter Server	16	32	415	SUSE Linux Enterprise 11
View Connection Server	4	10	60	Server 2008 64-bit R2
View Composer	4	10	100	Server 2008 64-bit R2

Table 19 shows the peak resource usage by management block components throughout workload, operations, and resiliency testing.

Table 19. Management Peak Resource Usage

RESOURCE	MAX CPU (MHZ)	MAX MEM USAGE (MB)	MAX NETWORK RX/TX (KBPS)	MAX DISK I/O (KBPS)
View Manager	2887MHz	12569MB	239 / 135KBps	8861KBps
VMware View Composer	1138MHz	12237MB	6046/ 9025KBps	1436KBps
VMware vCenter	13169MHz	16850MB	4270/3701KBps	12474KBps
SQL Server	331MHz	12237MB	120/ 421KBps	1790KBps
Domain Controller	759MHz	6009MB	394/ 297KBps	3558KBps

Bill of Materials

Table 20 summarizes the bill of materials for this reference architecture.

Table 20. Bill of Materials

AREA	COMPONENT	QUANTITY
Host hardware	SuperMicro Server 2027R-AR24NV	16
	Intel(R) Xeon(R) CPU E5-2690 v2 @ 3.00GHz 10-core	2
	512 GB RAM	1
	LSI9211-8I	2
	Intel Ethernet 82599EB 10-Gigabit SFI/SFP+	2
	Intel 800GB S3700 Series SATA 6Gb/s	2
	Seagate 1.2TB 2.5” SAS 6Gb/s 10K RPM 64M	12
Software	VMware ESXi6.0 2494585	16
	VMware vCenter Server 6.0.0, 2562627	2
	View 6.0.2	1
	Microsoft Windows 2008 R2	6
	Microsoft SQL Server 2008 R2	1

Conclusion

VMware Virtual SAN is a low-cost, high-performance storage platform that is rapidly deployed, easy to manage, and fully integrated into the industry-leading VMware vSphere cloud suite.

Virtual SAN scales as your VDI user base does, keeping CapEx costs down and eliminating the need for the large upfront investment that traditional storage arrays often require.

Extensive workload, operations, and resiliency testing shows that Horizon View on Virtual SAN delivers high levels of performance, a great end-user experience, and solid system resiliency, all at a low price point.

References

For additional information, see the following product documentation:

- VMware Horizon 6 Documentation
- VMware Horizon 6 Resources
- Optimization Guide for Windows 7 and Windows 8 Virtual Desktops in Horizon with View
- View Storage Accelerator in VMware View 5.1
- VMware Virtual SAN Network Design Guide
- VMware Virtual SAN 6.0 Design and Sizing Guide
- VMware Virtual SAN Design and Sizing Guide for Horizon View Virtual Desktop Infrastructures
- VMware Virtual SAN Documentation
- VMware Virtual SAN Hardware Compatibility Guide

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