



White Paper

Cisco UCS C240 M4 Rack Server with Horizon 6 and VMware Virtual SAN 6.0 All-Flash Configuration

Reference Architecture

December 2015

Contents

Solution Overview	3
Cisco Unified Computing System.....	4
Cisco UCS C-Series Rack Servers	5
Service Profiles	7
VMware vSphere	8
VMware Virtual SAN 6.0	8
VMware Horizon 6 with View	10
VMware View Storage Accelerator	11
System Configuration: Design	12
Cisco UCS Configuration	13
Service Profile Configuration.....	14
BIOS Policy.....	14
Boot Policy.....	16
Networking	16
QoS Policy	17
Multicast Policy	18
VLANs.....	19
VMware ESXi Network Configuration	19
VMware Virtual SAN Configuration	21
VMware Horizon 6 Configuration	21
VMware vSphere Clusters	22
VMware Horizon with View Configuration	23
VMware View Global Policies	23
VMware View Manager Global Settings.....	23
VMware vCenter Server Settings.....	24
VMware View Manager Pool Settings.....	24
Test Results	26
Test Summary.....	26
Test 1: Benchmarking with 660 Linked-Clones on Four Cisco UCS C240 M4 Servers in a VMware Virtual SAN Cluster	26
Test 2: VMware View Operations with 660 Linked Clones on Four Cisco UCS C240 M4 Servers in an All-Flash VMware Virtual SAN Cluster	30
VMware Virtual SAN Failure Simulation.....	33
VMware Virtual SAN Observer.....	34
System Sizing	35
Virtual Machine Test Image Build	35
Management Blocks	36
Host Configuration	36
Bill of Materials	37
Conclusion	38
For More Information	39

Executive Summary

The reference architecture in this document describes VMware Horizon 6 hosted on the Cisco Unified Computing System™ (Cisco UCS®) with a VMware Virtual SAN 6 All-Flash configuration as a hyper-converged storage solution. The purpose of this reference architecture is to provide guidance on the following aspects of deployment of this joint solution:

- Scalability and performance results for hosting 660 Horizon 6 virtual desktops on four Cisco UCS C240 M4 Rack Servers using industry-standardized benchmarking of real-world workloads
- Design and implementation best practices for Cisco UCS configurations, Virtual SAN 6.0–based storage policies, and data-store sizing guidance for hosting Horizon 6–based virtual desktops

This reference architecture uses LoginVSI as the standardized benchmarking tool for the application-centric benchmarking of real-world workloads. This tool helps ensure that the end-user experience and the performance of the solution components are taken into account. Horizons 6 virtual desktops are hosted on Cisco UCS C240 M4 rack-mounted servers on the VMware Virtual SAN compatibility list.

The test results summarized here and described in more detail later in this document demonstrate that VMware Virtual SAN together with Cisco UCS Rack Servers provides an easily scalable Horizon 6 based virtual desktop environment that also provides superior performance and manageability.

Table 1. Table 1 summarizes the test results.

Table 1. Test Results Summary

Test Results	Test Summary		
Deployed 660 desktops		660 linked-clones (100% concurrency)	<ul style="list-style-type: none"> • Excellent application response times hosting linked-clones: Helps ensure excellent end-user performance for practical workloads • Proven resiliency and availability: Provides greater uptime for applications. • Faster desktop operations: Improves IT efficiency.
69 minutes		660 linked-clones deployed	
6 minutes		660 linked-clones started	
51minutes		660 linked-clones refreshed	
102minutes		660 linked-clones recomposed	
Average application latency of less than 3 milliseconds (ms)		LoginVSI 4.1.4.2 workload	
Average disk latency of less than 13ms		Virtual SAN disk latency	

Solution Overview

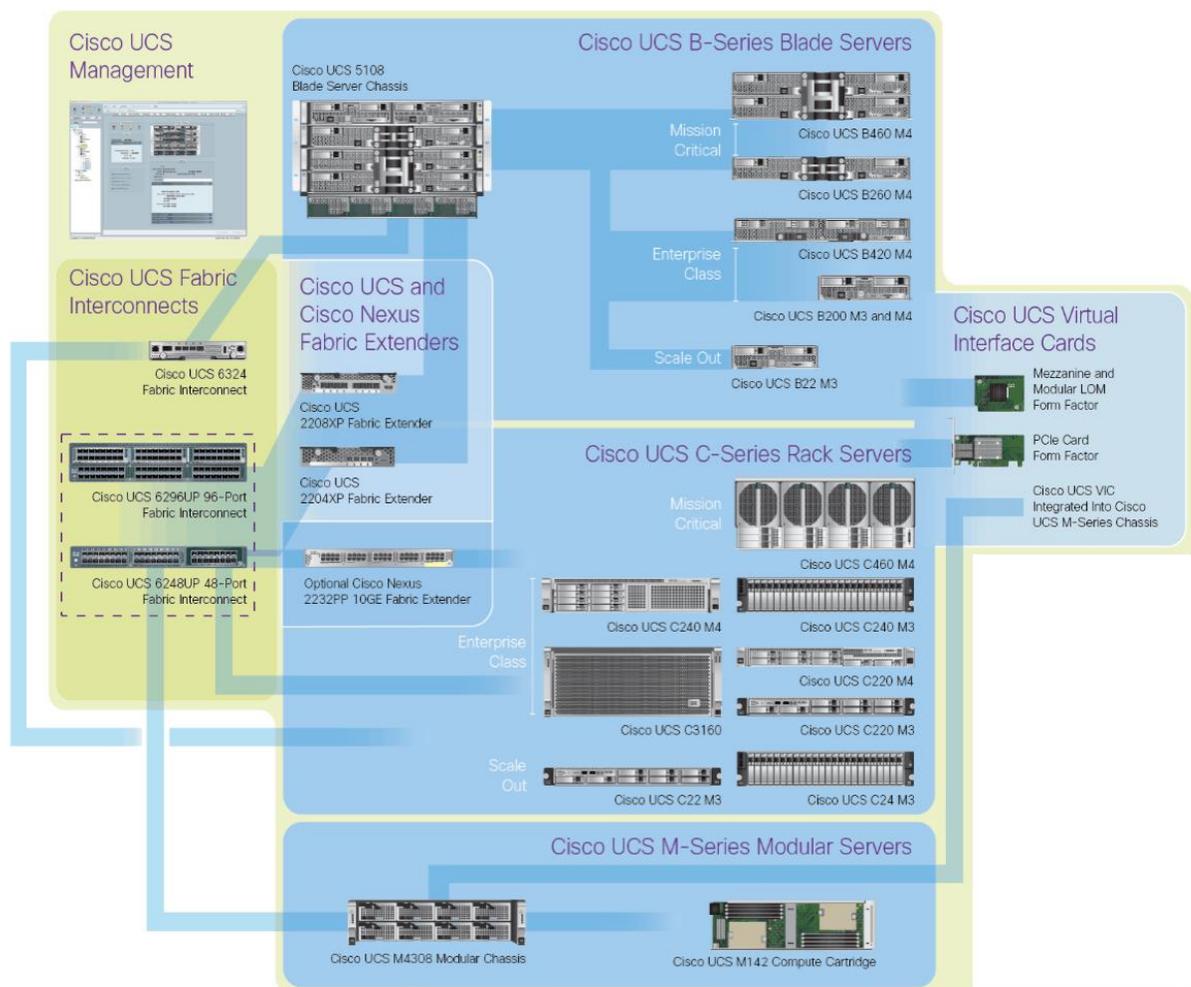
The VMware Horizon 6 with View hosted on Cisco UCS with VMware Virtual SAN 6.0 reference architecture supports and matches the Cisco UCS based [Virtual SAN Ready Nodes](#) for View virtual desktops. This document provides scalability and performance guidance for this solution. This section provides an overview of the individual components used in this solution. For more information about each product, refer to the respective product documentation.

Cisco Unified Computing System

Cisco UCS is a next-generation data center platform (Figure 1). It unites computing, networking, storage access, and virtualization resources into a single cohesive, integrated architecture that is designed to reduce total cost of ownership (TCO) and increase business agility.

The system integrates a low-latency, lossless 10 Gigabit Ethernet unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform that enables all resources to participate in a unified management domain.

Figure 1. Cisco Unified Computing System



Cisco UCS represents a radical simplification of traditional architectures that dramatically reduces the number of servers that are required to enable the platform. Cisco UCS helps reduce the TCO by automating element management tasks through the use of service profiles that enable just-in-time provisioning. Service profiles increase business agility by quickly aligning computing resources with rapidly changing business and workload requirements.

In addition, Cisco UCS delivers end-to-end optimization for virtualized environments, while retaining the capability to support traditional operating systems and application stacks in physical environments.

Cisco UCS offers these main advantages:

- **Less infrastructure and more intelligent servers:** The Cisco UCS architecture enables end-to-end server visibility, management, and control in both virtual and bare-metal environments. The Cisco UCS platform facilitates the move to cloud computing and IT-as-a-service (ITaaS) platforms with fabric-based Infrastructure.
- **Resource consolidation with Cisco UCS servers:** Cisco UCS servers simplify traditional architectures and optimize virtualized environments across the entire system. With Cisco® servers, bare-metal and virtualized applications can be supported in the same environment.
- **Faster server deployment:** The smart, programmable infrastructure of Cisco UCS simplifies and accelerates enterprise-class application and service deployment in bare-metal, virtualized, and cloud computing environments. With unified model-based management, hundreds of servers can be configured as quickly as just one server, resulting in lower cost of ownership and improved business continuity.
- **Simplified management:** Cisco UCS offers simplified and open management with a large partner ecosystem. Cisco UCS Manager provides embedded management of all software and hardware components in Cisco UCS. With Cisco UCS Central Software, management can be extended globally to thousands of servers across multiple Cisco UCS domains and geographies. In addition, Cisco UCS Director unifies management across computing, networking, and storage components in converged infrastructure solutions.

For more information about the capabilities and features of Cisco UCS technologies, see the [For More Information](#) section at the end of this document.

Cisco UCS C-Series Rack Servers

Cisco UCS C-Series Rack Servers deliver unified computing in an industry-standard form factor to reduce TCO and increase agility. Each product addresses different workload challenges through a balance of processing, memory, I/O, and internal storage resources.

Cisco UCS C-Series servers provide the following benefits:

- Form-factor independent entry point to Cisco UCS
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice with unique benefits in a familiar rack package
- Reduced TCO and increased business agility

Several Cisco UCS C-Series server models are available, each optimized for particular types of deployments. For VMware Virtual SAN deployments, disk density is an important factor in model selection. Computing power is also an important consideration.

These servers can be managed by the built-in, standalone software provided by the Cisco Integrated Management Controller (IMC), or by Cisco UCS Manager when connected through Cisco UCS 6200 Series Fabric Interconnects. Cisco UCS Manager supplies a totally integrated management process for both rack and blade servers through a single tool. All Cisco UCS servers use industry-leading Intel® Xeon® processors.

For Cisco UCS with VMware Virtual SAN, the Cisco UCS C240 M4 Rack Server (Figure 2) was considered the optimal choice for the development of this solution. The Cisco UCS C240 M4 is a high-density, enterprise-class, dual-socket, 2-rack-unit (2RU) rack server designed for computing, I/O, storage, and memory-intensive standalone and virtualized applications. The addition of the Intel Xeon processor E5-2600 product family delivers an optimal combination of performance, flexibility, and efficiency gains.

Figure 2. Cisco UCS C240 M4 Small Form-Factor (SFF) Rack Server



The Cisco UCS C240 M4 supports:

- Up to two Intel Xeon processor E5-2600v3 CPUs
- Up to 1.5 terabytes (TB) of RAM with 24 dual in-line memory module (DIMM) slots
- Capacity for up to 24 serial-attached SCSI (SAS), serial ATA (SATA), and solid-state disk (SSD) drives for workloads that demand large amounts of internal storage
- Six PCI Express (PCIe) Generation 3 slots and two 1 Gigabit Ethernet LAN interfaces on the motherboard
- Trusted platform module (TPM) for authentication and tool-free access

The Cisco UCS C240 M4 2RU server comes in two varieties. It can support either large form-factor (LFF; 3.5-inch) or small form-factor (SFF; 2.5-inch) hard drives. The VMware Virtual SAN technology uses SSD drives, so typically the SFF variety is required.

Computing performance is important because the virtual machines that use the Virtual SAN data store reside on the same hosts that contribute disk capacity to the data store.

The Cisco UCS C240 M4:

- Delivers outstanding internal memory and storage expandability and exceptional performance
- Is suitable for nearly all memory-intensive, storage-intensive, 2-socket applications
- Uses the unique Cisco UCS Virtual Interface Card (VIC) 1227: a dual-port 10 Gigabit Ethernet and Fibre Channel over Ethernet (FCoE)-capable PCIe modular LAN-on-motherboard (mLOM) adapter designed exclusively for Cisco UCS C-Series Rack Servers that can support up to 256 PCIe standards-compliant virtual interfaces (Figure 3)

Figure 3. Cisco UCS Virtual Interface Card 1227



- Is an exceptional building block and entry point for Cisco UCS
- Creates a highly flexible I/O environment that requires only one card for all I/O configurations
- Supports continual innovation in Cisco server technology at all levels of Cisco UCS

For more information about the capabilities and features of Cisco UCS technologies, see the [For More Information](#) section at the end of this document.

Service Profiles

In Cisco UCS, a service profile adds a layer of abstraction to the actual physical hardware. The server is defined in a configuration file that is stored on the Cisco UCS 6200 Series Fabric Interconnects. It can be associated with the physical hardware by using a simple operation from Cisco UCS Manager. When the service profile is applied, the manager configures the server, adaptors, fabric extenders, and fabric interconnects to match the specified service profile. The service profile makes the physical hardware transparent to the operating systems and virtual machines running on it, enabling stateless computing and making the most of data center resources.

A number of parameters can be defined in the service profile, according to the environment requirements. Administrators can create policies to define specific rules and operating characteristics. These policies can be referenced in the service profiles to help ensure consistent configuration across many servers. Updates to a policy can be propagated immediately to all servers that reference that policy in their service profile, or in the case of firmware updates, they can be propagated at the next power-cycling event.

In addition, the advantages of the service profile can be extended when server-specific parameters, such as the universal user ID (UUID), MAC address, and worldwide name (WWN), are themselves parameterized, and the service profile is converted to a template. The template can be used to rapidly deploy new servers with consistent general parameters and unique server-specific parameters.

When combined with templates, service profiles enable the rapid provisioning of servers with consistent operation parameters and high availability. Service profiles can be configured in advance and used to move servers to a new blade, chassis, or rack in the event of a failure.

VMware vSphere

VMware vSphere is the industry-leading virtualization platform for building cloud infrastructure. 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 it enables compatible public cloud offerings, forming the foundation for the industry's best hybrid cloud model.

For more information about the capabilities and features of VMware vSphere, see the [For More Information](#) section at the end of this document.

VMware Virtual SAN 6.0

VMware Virtual SAN is a hypervisor-converged storage solution that is fully integrated with VMware vSphere. Virtual SAN combines storage and computing for virtual machines into a single device, with storage provided within the hypervisor, instead of using a storage virtual machine that runs alongside the other virtual machines. Virtual SAN aggregates locally attached disks in a vSphere cluster to create a storage solution, called a shared datastore that can be rapidly provisioned from VMware vCenter Server during virtual machine provisioning operations.

Virtual SAN is an object-based storage system that is designed to provide virtual machine–centric storage services and capabilities through a storage policy–based management (SPBM) platform. The SPBM platform and virtual machine storage policies are designed to simplify virtual machine storage placement decisions for vSphere administrators.

Virtual SAN is fully integrated with core vSphere enterprise features such as vSphere vMotion, vSphere High Availability, and vSphere Distributed Resource Scheduler (DRS). Its goal is to provide both high availability and scale-out storage capabilities. In the context of quality of service (QoS), virtual machine storage policies can be created to define the level of performance and availability required on a per-virtual machine basis.

A Virtual SAN shared data store is constructed with a minimum of three vSphere ESXi hosts, each containing at least one disk group with at least one SSD drive and one magnetic drive, as shown in Figure 4. It can also support up to seven magnetic drives per disk group and up to five disk groups per host, or up to 35 drives per node. The VMware virtual machine files are stored on the magnetic drive, and the SSD drive handles read caching and write buffering. The disk group on each host is joined to a single network partition group, which is shared and controlled by the hosts.

Virtual SAN 6.0 provides a new all-flash architecture on flash-memory devices that delivers high, predictable performance and response times of less than a millisecond for some of the most demanding enterprise applications.

Virtual SAN 6.0 supports double the scalability, with up to 64 nodes per cluster and up to 200 virtual machines per host and up to 6400 virtual machines per Virtual SAN cluster, along with performance enhancements and highly efficient snapshot and clone technology.

Figure 4. VMware Virtual SAN Cluster: All-Flash Data Store

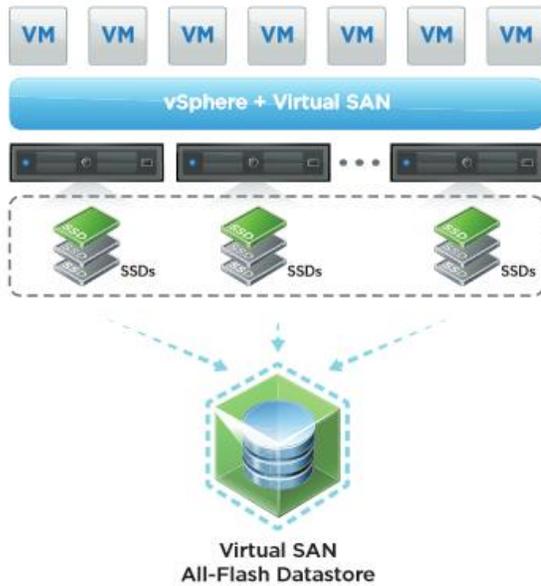
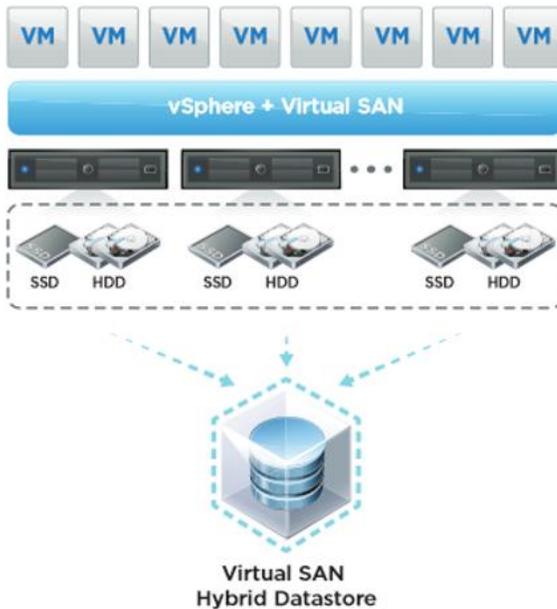


Figure 5. VMware Virtual SAN Cluster: Hybrid Data Store



Note: This document focuses on the All-Flash solution. For a Hybrid solution follow this link: <http://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/unified-computing/whitepaper-c11-735680.html>

The size and capacity of the Virtual SAN shared data store are dictated by the number of magnetic disks per disk group in a vSphere host and by the number of vSphere hosts in the cluster. Virtual SAN is a scale-out solution, in

which more capacity and performance can be obtained by adding more disks to a disk group, more disk groups to a host, and more hosts to the cluster.

With Virtual SAN, the SPBM platform plays a major role in the way that administrators can use virtual machine storage policies to specify a set of required storage capabilities for a virtual machine or, more specifically, to specify a set of requirements for the application running in the virtual machine.

The following Virtual SAN data store capabilities are available in vCenter Server, configurable in virtual machine storage policy:

- Number of failures to tolerate
- Number of disk stripes per object
- Flash read-cache reservation
- Object-space reservation
- Force provisioning

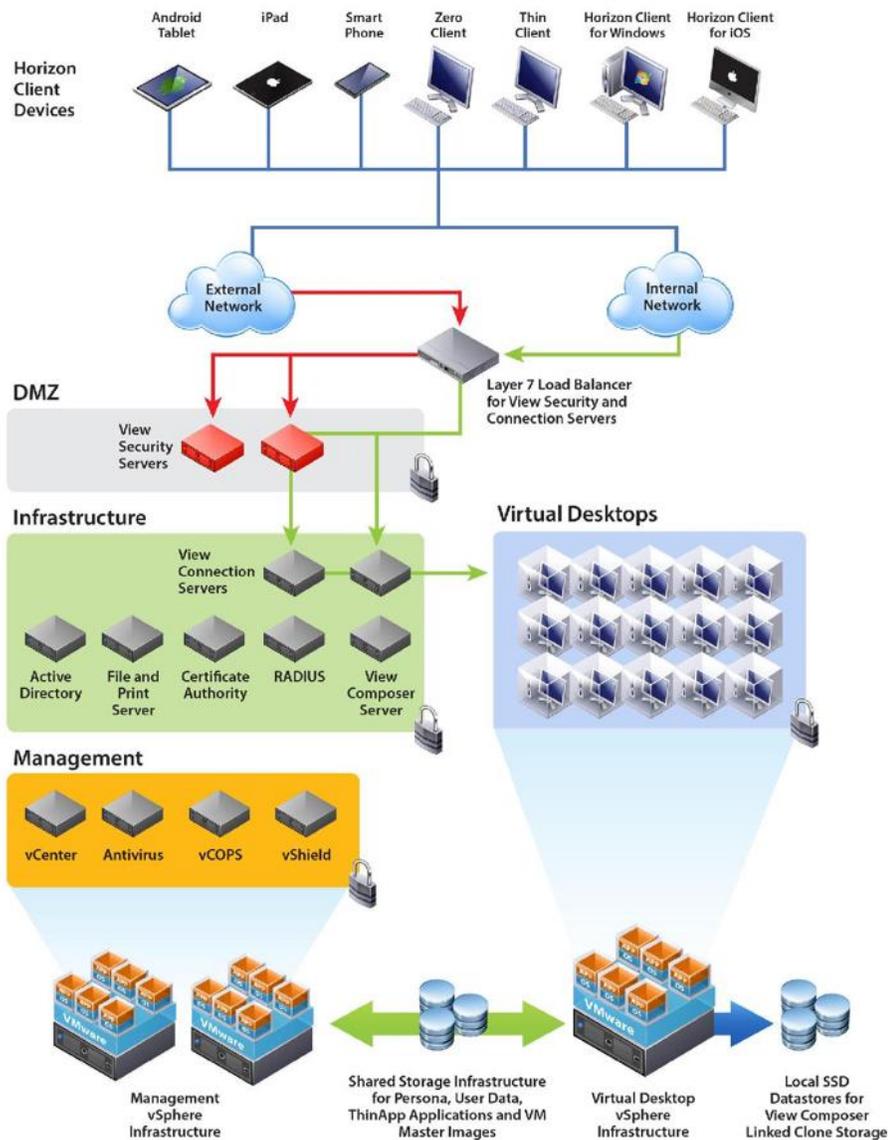
For more information about the capabilities and features of VMware Virtual SAN, see [What's New in VMware Virtual SAN](#).

VMware Horizon 6 with View

VMware Horizon with View brings the agility of cloud computing to the desktop by transforming desktops into highly available and agile services delivered from the cloud (Figure 6). View delivers virtual sessions that follow end users across devices and locations. It enables fast, secure access to corporate data across a wide range of devices, including Microsoft Windows, Mac OS, and Linux desktop computers and iOS and Android tablets.

With vCenter Server, View can be used to create desktops from virtual machines that are running on ESXi hosts, and to deploy these desktops to end users. After a desktop is created, 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 the existing Microsoft Active Directory infrastructure for user authentication and management.

Figure 6. VMware Horizon with View Components



VMware View Storage Accelerator

VMware View Storage Accelerator is 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, which greatly reduces the number of read I/O requests that are issued to the storage layer. It also addresses boot storms—when multiple virtual desktops are booted at the same time—which can cause a large number of read operations. CBRC is beneficial when administrators or users load applications or data frequently. Note that CBRC was used in all tests that were performed on the solution described here: Horizon with View running on Virtual SAN hosted on Cisco UCS.

System Configuration: Design

This section describes the configuration of the components of the reference architecture for Horizon 6 with View hosted on Cisco UCS with Virtual SAN.

As shown in Figure 7, 660 Horizon 6 with View virtual desktops running Microsoft Windows 7 were hosted on four Cisco UCS C240 M4 Rack Servers with Virtual SAN 6.0 on vSphere 6.0. With the scalability testing performed using Login VSI as a benchmarking tool, the solution exhibits linear scalability with exceptional end-user performance. According to guidance in [VMware vSphere 6.0 Configuration Maximums](#), this solution can be further scaled to 6400 virtual desktops by adding up to 64 nodes in a single Virtual SAN cluster.

Figure 7. VMware View Hosted on VMware Virtual SAN on Cisco UCS (Details)

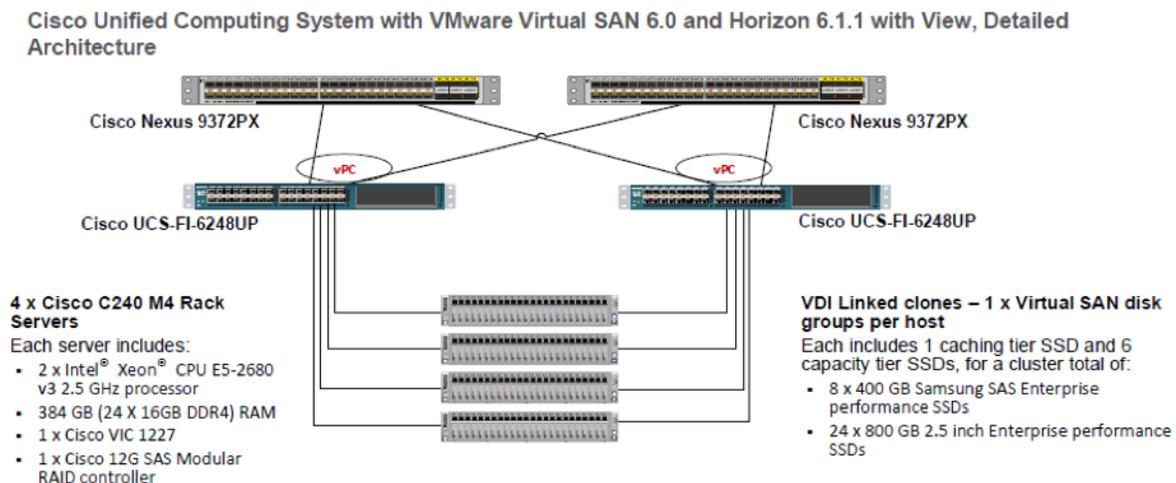


Table 2 provides guidance about Cisco UCS nodes ready for Virtual SAN. The table lists the optimal number of disks and disk groups configured while testing a linked-clone setup.

For an All-Flash Virtual SAN configuration, the caching and the capacity tiers are both made up from flash devices. Because the capacity tier is also flash memory, read cache is not used as the capacity tier can handle read operations effectively without much of a penalty. The caching tier is used for the write cache, and all virtual machine write operations use this cache device. The write cache is used to hold “hot,” or active, blocks of data, and when the blocks become “cold,” or inactive, they are moved to the capacity tier.

With an all-Flash Virtual SAN, the recommended size of the write caching tier is about 10 percent the size of the capacity tier, similar to size in a hybrid Virtual SAN configuration. The primary objective is to try to have both caching and capacity tiers wear out around the same time based on the relative endurance and capacity of the devices. If you have additional insight into your workloads, then a more customized calculation than the 10 percent rule is possible. For complete detail about design and sizing considerations, see the [VMware Virtual SAN 6.0 Design and Sizing Guide](#).

Table 2 shows a disk group configuration. Capacity can be added per drive group as one 400-GB SSD drive and six 800-GB SAS SSD drives per disk group (4.69 TB of raw capacity per disk group).

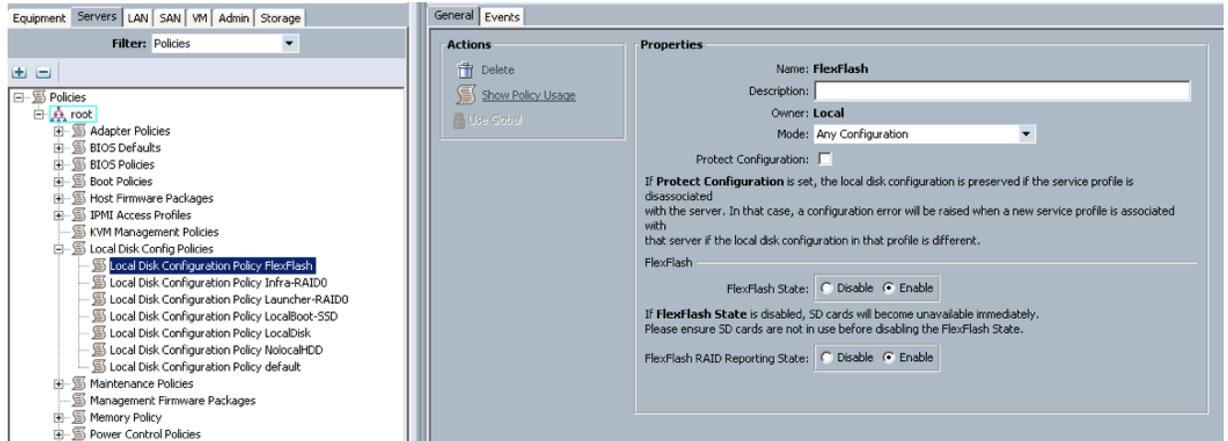
Table 2. Virtual SAN Disk Group Configuration

Type of Virtual Desktop	Number of Disk Groups per Host	SSD Drives and HDDs in Disk Groups	Used and Total Capacity
660 linked-clones (floating pool)	1	1 SSD and 6 data-tier SDDs	4.4 TB used out of 17.29 TB
660 linked-clones (dedicated pool)	1	1 SSD and 6 data-tier SSDs	5.12 TB used out of 17.29 TB

Cisco UCS Configuration

In this configuration, ESXi is booted from the on-board Cisco FlexFlash Secure Digital (SD) card. For more information, see [Cisco FlexFlash: Use and Manage Cisco Flexible Flash Internal SD Card for Cisco UCS C-Series Standalone Rack Servers](#). The Cisco FlexFlash SD card configuration is performed through a local disk policy that is applied to the service profile, as shown in the example in Figure 8.

Figure 8. Local Disk Configuration Policy: FlexFlash State Enabled



Service Profile Configuration

Table 3 summarizes the main configurable parameters of a Cisco UCS service profile.

Table 3. Service Profile Parameters

Parameter Type	Parameter	Description
Server hardware	• UUID	• Obtained from defined UUID pool
	• MAC addresses	• Obtained from defined MAC address pool
	• Worldwide port name (WWPN) • Worldwide node name (WWNN)	• Obtained from defined WWPN and WWNN pools
	• Boot policy	• Boot path and order
	• Disk policy	• RAID configuration
Fabric	• LAN	• Virtual NICs (vNICs), VLANs, and maximum transmission unit (MTU)
	• SAN	• Virtual host bus adapters (vHBAs) and Virtual SANs
	• QoS policy	• Class of service (CoS) for Ethernet uplink traffic
Operation	• Firmware policy	• Current and backup versions
	• BIOS policy	• BIOS version and settings
	• Statistics policy	• System data collection
	• Power-control policy	• Blade server power allotment

For Cisco UCS service profiles for hosts in a Virtual SAN cluster, the policy configuration shown here is recommended. This configuration does not include all Cisco UCS service profile settings. The settings shown here are configurations specific to an implementation of Cisco UCS with Virtual SAN for Horizon with View.

BIOS Policy

The BIOS policy configured for the Virtual SAN environment is aimed at achieving high performance, as shown in the example in Figure 9 and Table 4.

Figure 9. BIOS Policy Configuration for the VMware Virtual SAN Environment

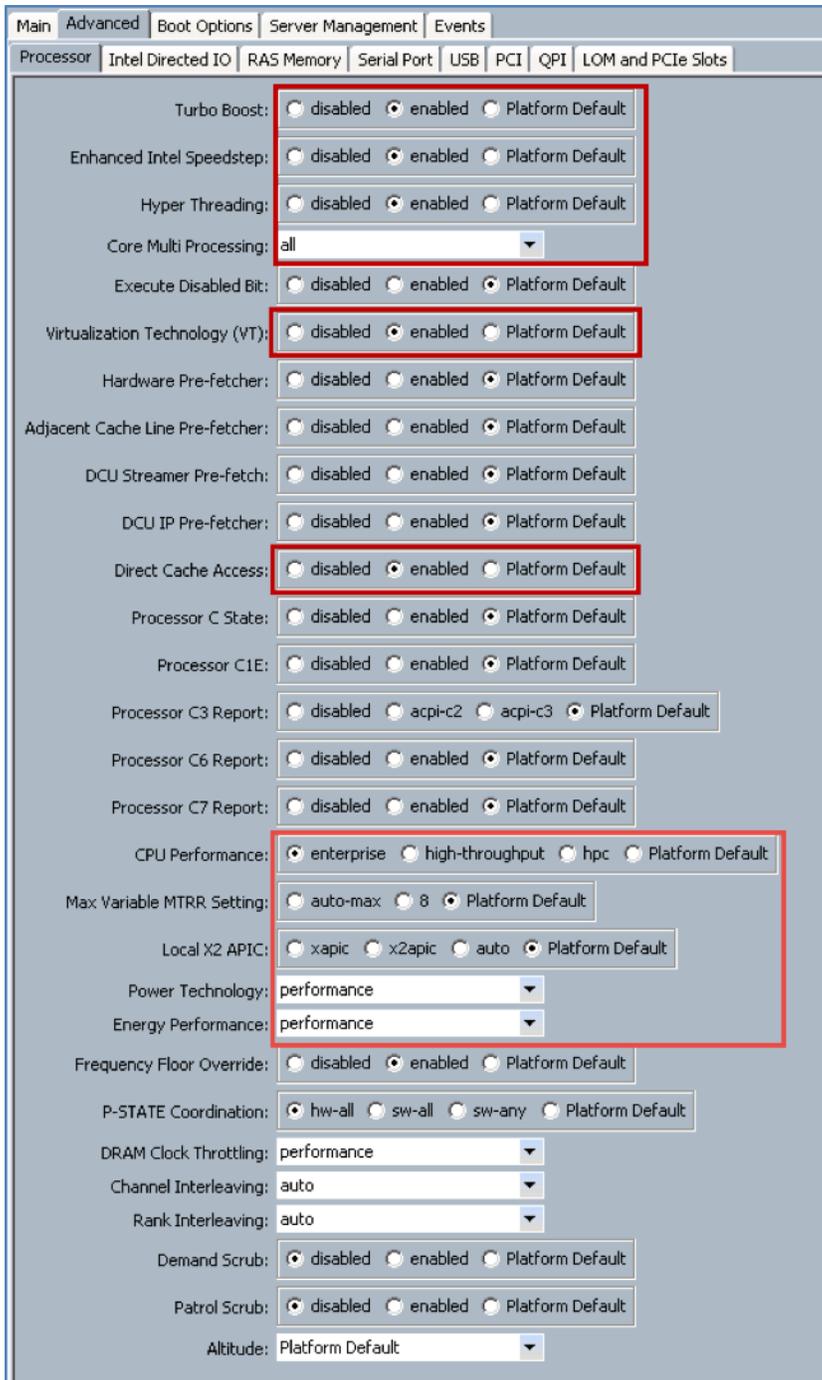


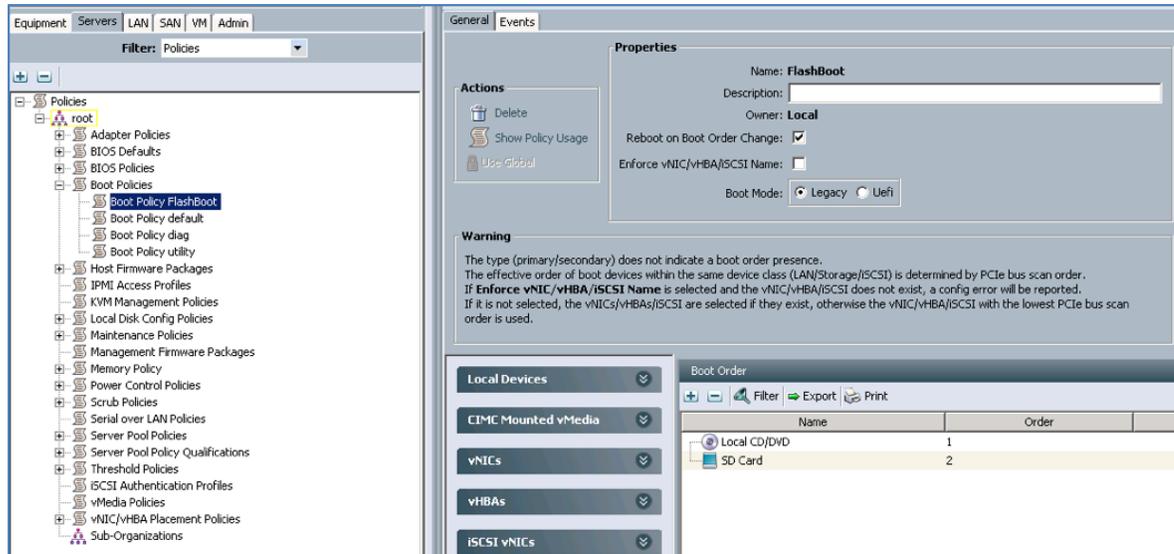
Table 4. BIOS Policy Settings for the VMware Virtual SAN Environment

Policy	Settings
Processor	<ul style="list-style-type: none"> • Turbo Boost = Enabled • Enhanced Intel Speedstep = Enabled • Hyper-Threading = Enabled • Virtualization Technology (VT) = Enabled • Direct Cache Access = Enabled • CPU Performance = Enterprise • Power Technology = Performance • Energy Performance = Enterprise
Intel Directed IO	<ul style="list-style-type: none"> • VT for Directed IO = Enabled
Memory	<ul style="list-style-type: none"> • Memory RAS Configuration = Maximum Performance • Low-Voltage DDR Mode = Performance Mode

Boot Policy

The boot policy is created with an SD card as the preferred boot option after the local CD or DVD boot option (Figure 10).

Figure 10. Boot Policy Configuration



Networking

The vSphere distributed switch is configured for all hosts in the cluster. This configuration allows virtual machines to maintain a consistent network configuration as the virtual machines migrate across multiple hosts. A separate vNIC is created for each side of the fabric to provide high availability and load balancing across both fabrics for virtual machine data, Virtual SAN, vMotion, and management traffic. These are configured as separate vNIC templates within Cisco UCS and applied as part of the service profile.

Two separate virtual switches are created. The first standard vSwitch (vSwitch0) is configured with vNICs defined for management, and another distributed vSwitch (dvSwitch) is configured with vNICs defined for virtual machine data, vMotion, and Virtual SAN traffic (Table 5).

Table 5. vNIC Template Configuration

vNIC Template Name	Fabric ID	Comments
Data-Fab-A	Fabric A	MTU = 1500; QoS Policy - Silver
Data-Fab-B	Fabric B	MTU = 1500; QoS Policy - Silver
Virtual SAN-A	Fabric A	MTU = 9216; QoS Policy - Platinum
Virtual SAN-B	Fabric B	MTU = 9216; QoS Policy - Platinum

The network control policy is set to Cisco Discovery Protocol enabled, and the dynamic vNIC connection policy is applied as connection policy, not with an adapter policy of “VMware.”

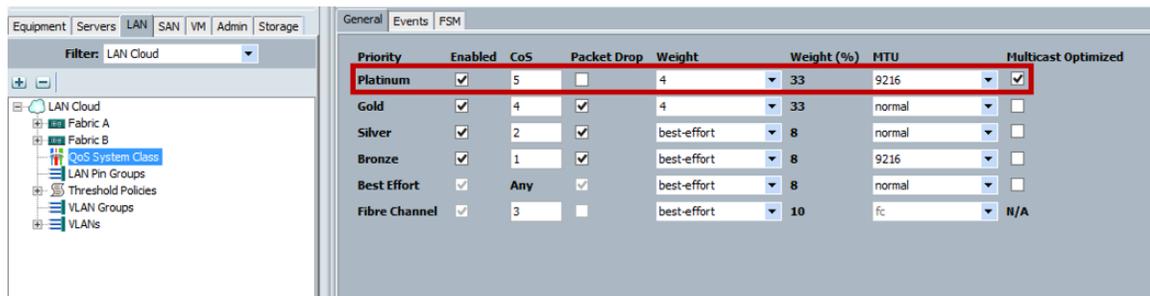
QoS Policy

The QoS policy and QoS system class mappings within Cisco UCS can be set for individual vNICs as mentioned previously. Table 6 and Figure 11 show sample vNICs with their traffic priorities. You should check the network utilization for your individual environment and set QoS policy prioritization accordingly. Set the MTU to 9216 if you require jumbo frames.

Table 6. QoS Policy Configuration

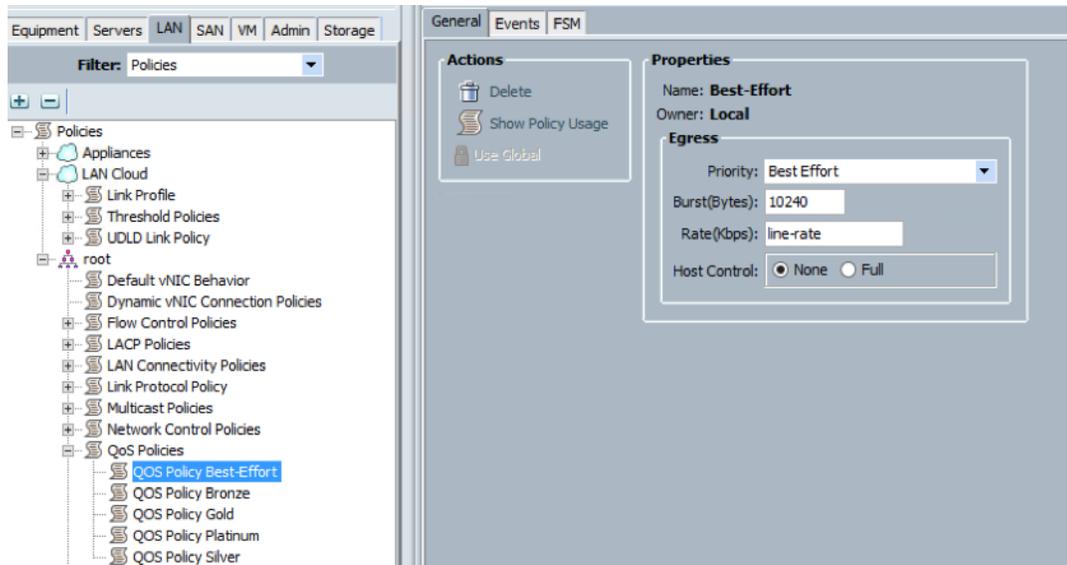
QoS Policy Name	Priority
VM Data	Gold
Virtual SAN	Platinum
vMotion	Bronze
MGMT	Silver

Figure 11. QoS System Class Configuration



For Virtual SAN traffic, this solution used a QoS policy with the priority set to Platinum and multicast optimized as highlighted in the screenshot above. (Figure 11).

Figure 12. QoS Policy



Multicast Policy

Create a multicast policy for Virtual SAN traffic as shown in Figure 13:

- IGMP Snooping State = Enabled
- IGMP Snooping Querier State = Enabled
- IGMP Snooping Querier IPv4 Address = x.x.x.x

Note: As a best practice, assign a default gateway on the same subnet for Virtual SAN traffic.

Figure 13. Multicast Policy

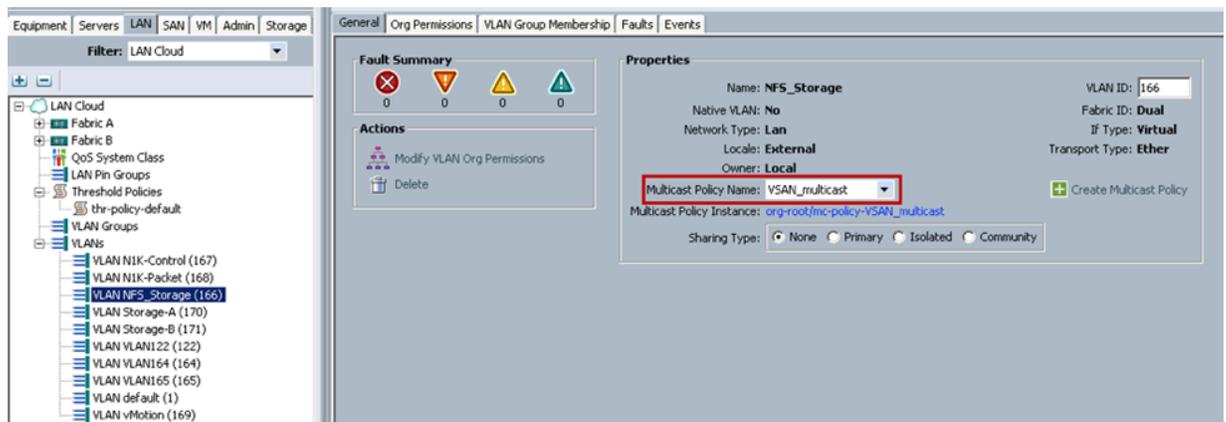


VLANs

A dedicated VLAN is recommended for the Virtual SAN VMkernel NIC, and multicast is required within the Layer 2 domain. This setting is configured as part of the VLAN as a multicast policy with snooping enabled. The following VLANs were created:

- VLAN for Virtual SAN: This is a /28 subnet with 14 IP addresses to accommodate four hosts. Attach multicast policy (Figure 14).

Figure 14. Virtual SAN VLANs



- VLAN for management components: This is a /24 subnet with 254 IP addresses to accommodate all management components plus the launcher server for Login VSI to run the test workflows.
- VLAN for virtual desktops: This is a /22 subnet with 1022 IP addresses to accommodate all 660 virtual desktops.

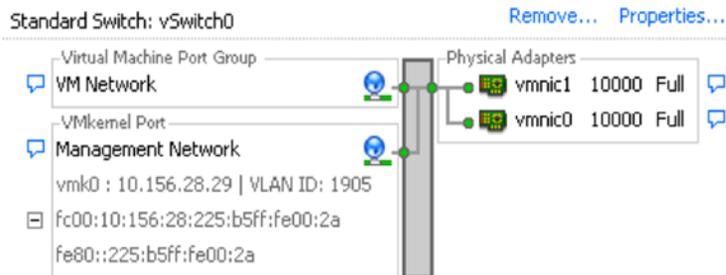
VMware ESXi Network Configuration

Properties regarding security, traffic shaping, and NIC teaming can be defined for a port group. The settings used in the port group design are shown in Figure 15.

Figure 15. Port Group Properties: Distributed vSwitch

View: vSphere Standard Switch vSphere Distributed Switch

Networking



View: vSphere Standard Switch vSphere Distributed Switch

Networking

▼ Distributed Switch: dvSwitch [Manage Virtual Adapters...](#) [Manage Physical Adapters](#)

dvSwitch ⓘ



VMware Virtual SAN Configuration

Virtual SAN is an ESXi cluster-level feature that is configured using the vSphere Web Client. The first step in enabling Virtual SAN is to select one of the two modes of disk group creation:

- **Automatic:** Enable Virtual SAN to discover all the local disks on the hosts and automatically add the disks to the Virtual SAN data store.
- **Manual:** Manually select the disks to add to the Virtual SAN shared data store.

In this setup, disk groups were created manually and the storage policies shown in Table 7 were applied, depending on whether the Virtual SAN configuration is for linked clones. These storage policies are tied to the storage requirements for each virtual machine and are used to provide different levels of availability and performance for virtual machines.

Important: Use different policies for different types of virtual machines in the same cluster to meet application requirements.

Table 7. Storage Policies for VMware View

Policy	Definition	Default (Value Applied)	Maximum
Number of disk stripes per object	Defines the number of magnetic disks across which each replica of a storage object is distributed	1	12
Flash-memory read-cache reservation	Defines the flash memory capacity reserved as the read cache for the storage object	0%	100%
Number of failures to tolerate	Defines the number of host, disk, and network failures a storage object can tolerate; for n failures tolerated, n + 1 copies of the object are created, and 2n + 1 hosts of contributing storage are required	0 (linked clones [floating pool]); 1 (full clones and linked clones [dedicated pool] replicas)	3 (in 4-host cluster)
Forced provisioning	Determines whether the object is provisioned, even when currently available resources do not meet the virtual machine storage policy requirements	Disabled	Enabled
Object-space reservation	Defines the percentage of the logical size of the storage object that needs to be reserved (thick provisioned) upon virtual machine provisioning (the remainder of the storage object is thin provisioned)	0%	100%

Note: Note: Default storage policy values are configured for linked-clones, full-clones, and replicas.

VMware Horizon 6 Configuration

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

- **View Connection Server:** A View Connection Server supports up to 2000 concurrent connections. The tests used two View Connection Servers operating in active-active mode. The two connection servers actively broker and possibly tunnel connections.
- **View block:** View provisions and manages desktops through vCenter Server. Each vCenter instance supports up to 10,000 virtual desktops. The tests used one vCenter and one Virtual SAN cluster with eight hosts. Note that the maximum number of VMware High Availability protected virtual machines allowed in a vSphere cluster is 6400 per data store.

- **View management block:** A separate vSphere cluster was used for management servers to isolate the volatile desktop workload from the static server workload. For larger deployments, separate dedicated vCenter Servers for the management block and the View block are recommended.

VMware vSphere Clusters

A 4-node Virtual SAN cluster was deployed to support 660 virtual desktops, as shown in Figure 16 and Table 8.

Figure 16. VMware View hosted on VMware Virtual SAN on Cisco UCS

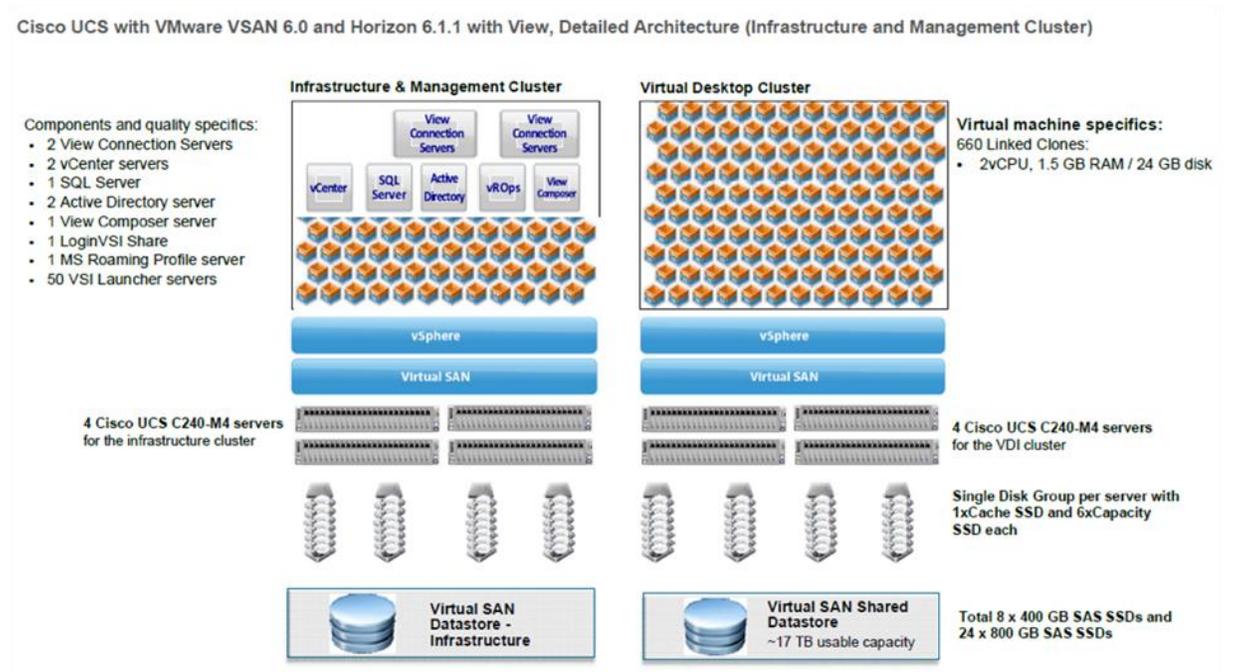


Table 8. VMware Virtual SAN Cluster Configuration

Property	Setting	Default	Revised
Cluster features	• HA	–	Enabled
	• DRS	–	Enabled
vSphere High Availability (HA)	• Host Monitoring Status	Enabled	–
	• Admission Control	Enabled	–
	• Admission Control Policy	Host failures the cluster tolerates = 1	–
	• Virtual Machine Options > Virtual Machine Restart Priority	Medium	–
	• Virtual Machine Options >Host Isolation Response	Leave powered on	–
	• VM Monitoring	Disabled	–
	• Data-Store Heartbeating	Select any, taking into account my preferences (no data store preferred)	–
vSphere DRS	• Automation Level	Fully automated (apply 1, 2, 3 priority recommendations)	–
	• DRS Groups Manager	–	–
	• Rules	–	–

Property	Setting	Default	Revised
Enhanced vMotion capability	• Virtual Machine Options	–	–
	• Power Management	Off	–
	• Host Options	Default (Disabled)	–
		Disabled	–
Swap-file location		Store in the same directory as the virtual machine	–

VMware Horizon with View Configuration

The Horizon with View installation included the following core systems.

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

Note: Note: View Security Servers were not used during this testing.

VMware View Global Policies

The View global policy settings used for all the system tests are shown in Table 9.

Table 9. VMware View Global Policies

Network Resource Pool	Host Limit (Mbps)
USB access	Allow
Multimedia redirection (MMR)	Allow
Remote mode	Allow
PC over IP (PCoIP) hardware acceleration	Allow (medium priority)

VMware View Manager Global Settings

The View Manager global policy settings used in the tests are shown in Table 10.

Table 10. VMware View Manager Global Settings

Attribute	Specification
Session timeout	600 (10 hours)
View Administrator session timeout	120 minutes
Auto update	Enabled
Display prelogin message	No
Display warning before logout	No
Reauthenticate secure tunnel connections after network interruption	No
Enable IP Security (IPsec) for security server pairing	Yes
Message security mode	Enabled
Disable single sign-on for local-mode operations	No

VMware vCenter Server Settings

View Connection Server uses vCenter Server to provision and manage View desktops. vCenter Server is configured in View Manager as shown in Table 11.

Table 11. VMware View Manager: VMware vCenter Server Configuration

Attribute	Setting	Specification
Connect using SSL	vCenter Server Settings > SSL	Yes
vCenter port	vCenter Server Settings > Port	443
View Composer port	View Composer Server Settings > Port	18443
Enable View Composer	View Composer Server Settings > Co-Installed	Yes
Advanced settings	Maximum Concurrent vCenter Provisioning Operations	40
	Maximum Concurrent Power Operations	50
	Maximum Concurrent View Composer Maintenance Operations	20
	Maximum Concurrent View Composer Provisioning Operations	12
Storage settings	Enable View Storage Accelerator	Yes
	Default Host Cache Size	2048 MB

VMware View Manager Pool Settings

The View Manager pool settings were configured as shown in Tables 12 and 13. For resiliency testing, a dedicated pool was used.

Table 12. VMware View Manager: VMware View Manager 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	/
Remote desktop power policy	Take no power action
Auto logoff time	Never
User reset allowed	False
Multisession allowed	False
Delete on logoff	Never
Display protocol	PCoIP
Allow protocol override	False
Maximum number of monitors	1
Maximum 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

Attribute	Specification
Provision all desktops upfront	Enabled

Table 13. VMware View Manager: Test Pool Configuration

Attribute	Specification
Disposable file redirection	Do not redirect
Disposable disk size	4096 MB
Select separate data stores for replica and OS	Not selected
Data stores: Storage overcommit	VSAN (Default)
Use View storage accelerator	Selected
Reclaim virtual machine disk space*	–
Disk types	OS disks
Regenerate storage accelerator after	7 days
Reclaim virtual machine disk space	–
Use Sysprep	Enabled

Note: Virtual SAN does not support the space-efficient (SE) sparse disk format.

Test Results

View running on Virtual SAN on the Cisco UCS reference architecture was tested based on real-world test scenarios, user workloads, and infrastructure system configurations.

The tests performed included the following configurations:

- Benchmarking with 660 linked-clone desktops on four Cisco UCS C240 M4 servers in an All-Flash Virtual SAN cluster
- View operations with 660 linked-clone desktops on four Cisco UCS C240 M4 servers in an All-Flash Virtual SAN cluster

These tests and the test results are summarized in the sections that follow.

Test Summary

LoginVSI is a desktop virtualization end-user experience measurement tool that uses a workload generator to automate and measure a typical office user's activity. The operations include opening a file, browsing the web, modifying files, saving, closing, printing and more. Each LoginVSI operation runs in an iterative fashion. Each iteration is a randomly sequenced workload of these applications and operations. The results of a run consist of operational latencies collected for these applications and operations for all iterations.

Each test scenario discussed here was run multiple times with a [knowledge worker workload](#) in [benchmark mode](#) to monitor the consistency and repeatability of the test results.

The tests monitored performance statistics during a virtual desktop boot storm and then waited for the systems to settle down. The tests tracked ramp-up, which is the login interval for all 500 or 1000 sessions, steady state, in which all sessions are logged on and active, and then log off.

In addition to [Login VSI Analyzer](#) results, VMware Virtual SAN Observer and ESXTOP data captured from individual test cycles were used as monitoring tools.

For more information about the applications used for these tests, see the section [For More Information](#) later in this document.

Test 1: Benchmarking with 660 Linked-Clones on Four Cisco UCS C240 M4 Servers in a VMware Virtual SAN Cluster

LoginVSI tests were run on 660 linked-clones hosted on four Cisco UCS C240 M4 servers. The tests revealed exceptional user performance, as represented by the LoginVSI Analyzer scores and latency values (Figures 17 through 23).

Test result highlights include:

- Average CPU utilization of 85 percent
- Average of up to 276 GB of RAM used out of 384 GB available
- Average 42.1-MBps network used
- Average I/O latency of 8.2 ms per host
- Average of 962 I/O operations per second (IOPS) per host

Figure 17. Login VSI Analyzer Chart for 660 Sessions with EUX Response Time

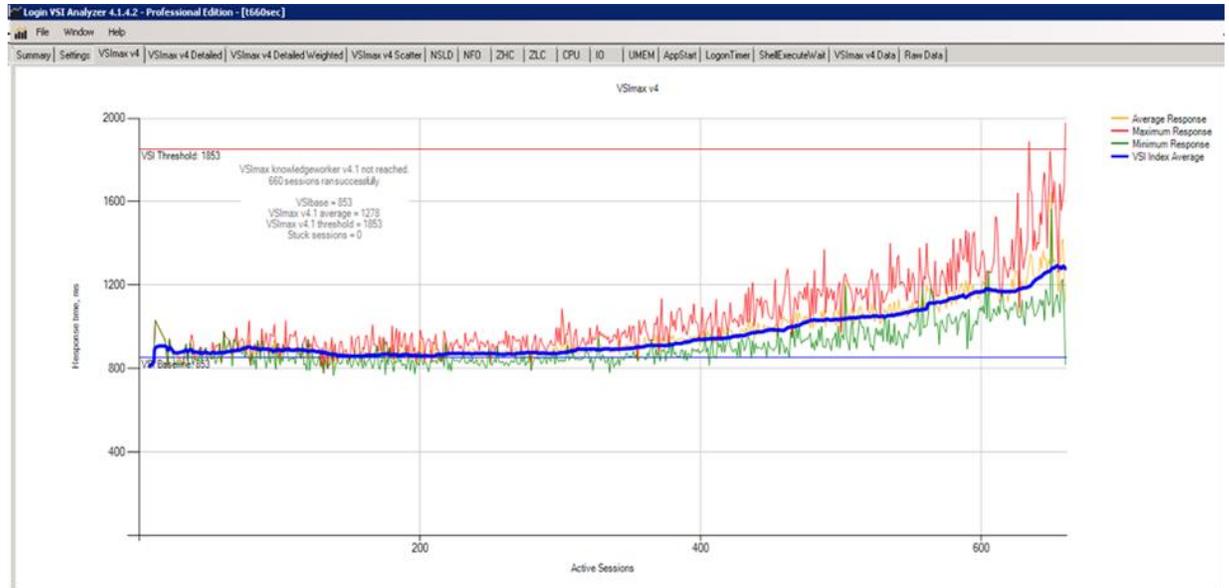


Figure 18. Host CPU Utilization from ESXTOP: 660 Linked Clones, Average CPU Utilization as a Percentage

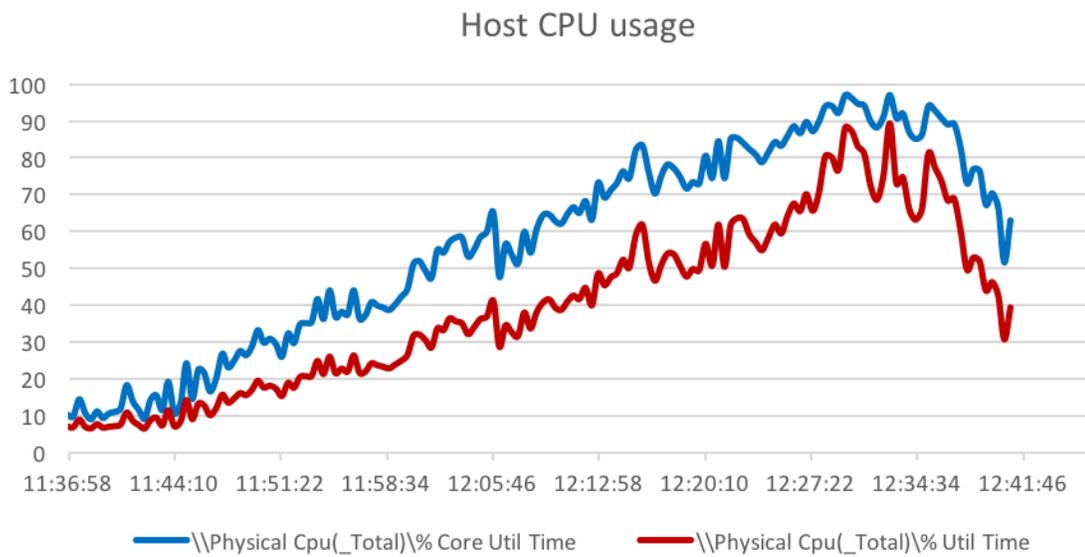


Figure 19. Host Memory Use from ESXTOP: 660 Linked Clones, Average Memory Use in GB

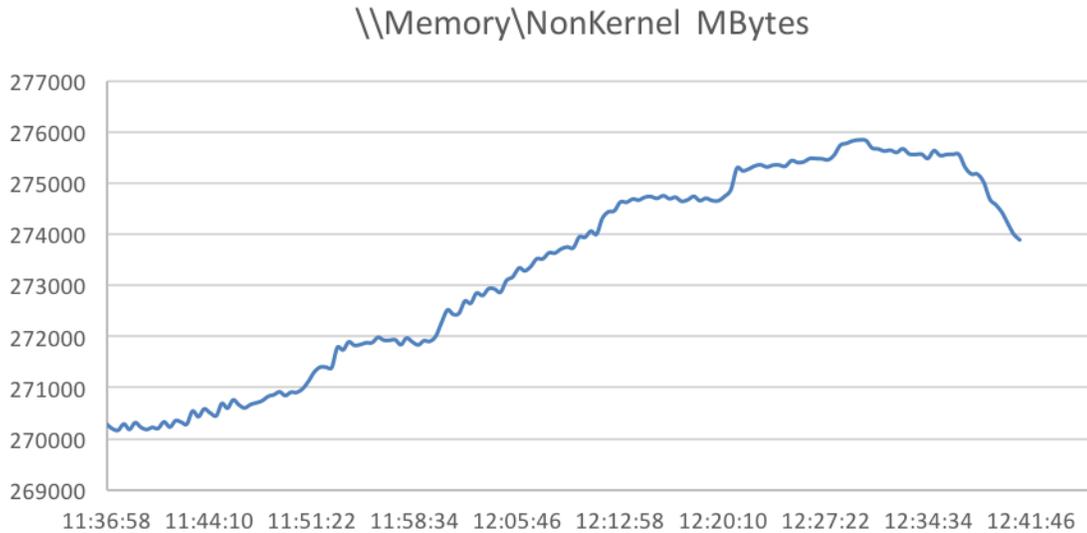


Figure 20. Host Network Utilization from ESXTOP: 660 Linked Clones, Average Network Utilization in Mbps

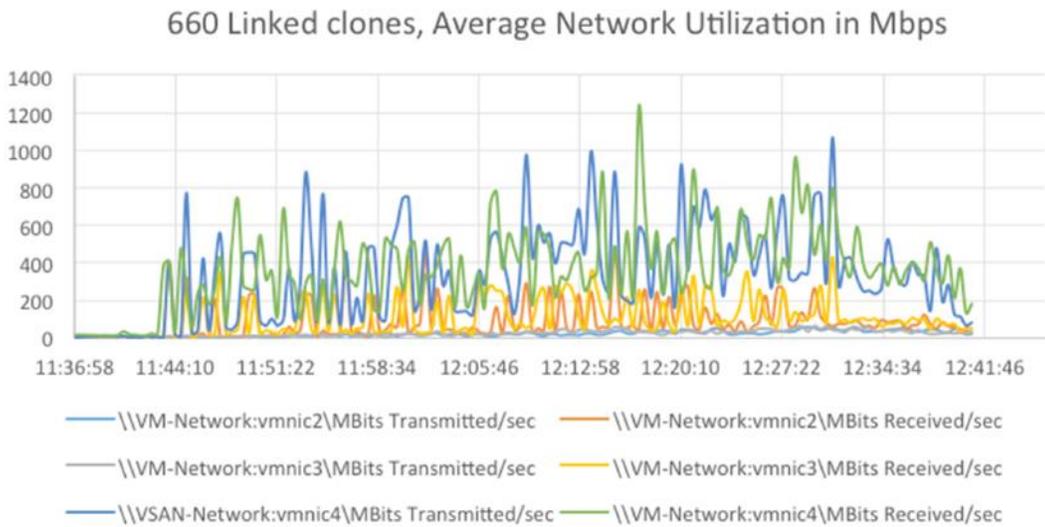


Figure 21. Host Storage Latency from ESXTOP: 660 Linked Clones, Average Read-Write Latency in Milliseconds

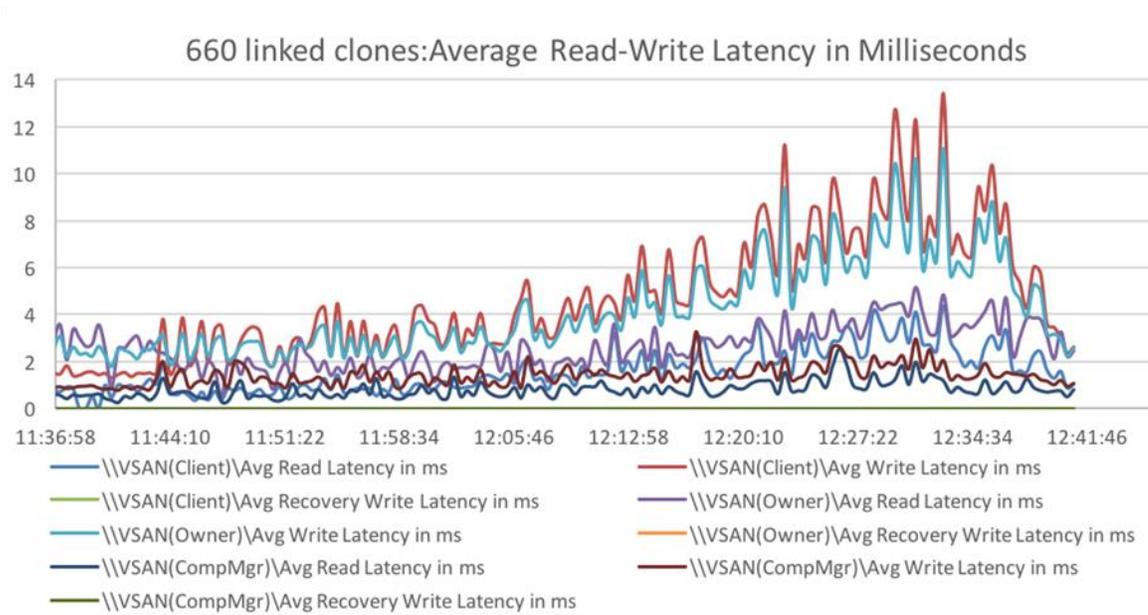


Figure 22. Host Storage Read-Write Rate from ESXTOP: 660 Linked Clones, Average Read-Write Rate in MBps

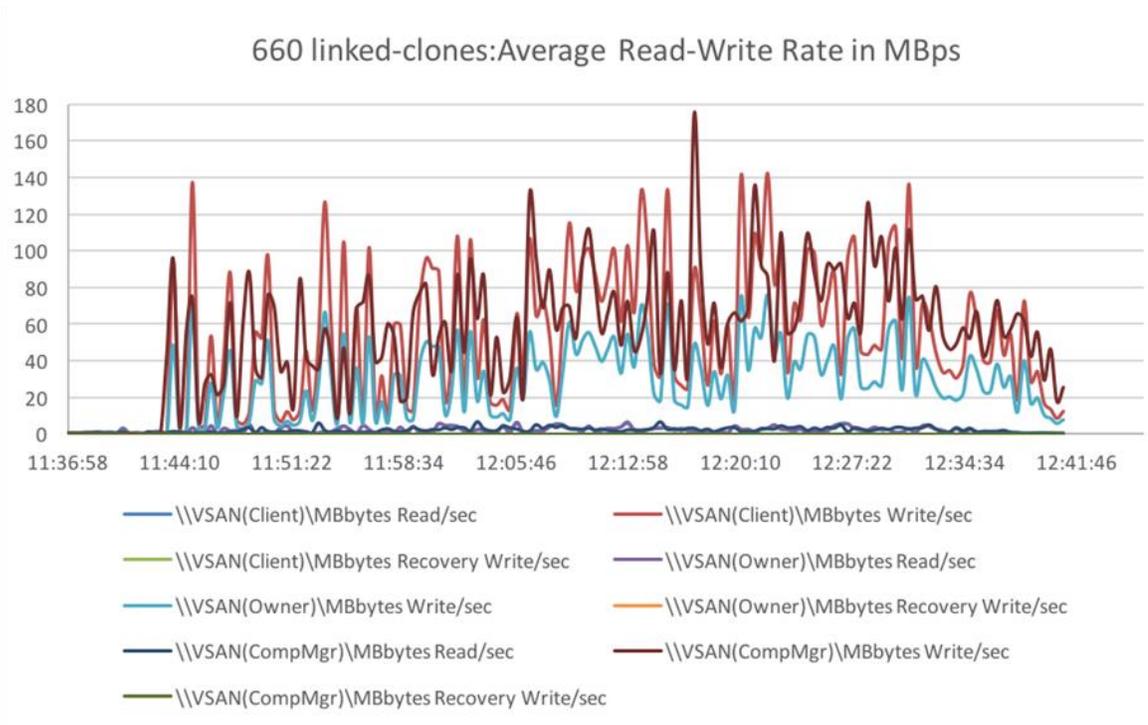


Figure 23. Host IOPS and Latency Graph from Virtual SAN Observer: 660 Linked Clones



Disk latency values shown in Figure 23 are obtained from Virtual SAN Observer. The average read and write latency is 9.2 ms for the host shown here, and on average 8.6 ms across all hosts. These values, which are below the target threshold of 20 ms, correlate with the low application response times measured by Login VSI, and the overall results of a better end user experience. In these tests, the average IOPS value generated per host was 962.

Test 2: VMware View Operations with 660 Linked Clones on Four Cisco UCS C240 M4 Servers in an All-Flash VMware Virtual SAN Cluster

In addition to LoginVSI benchmarking tests, VMware View operations tests were conducted to measure the effect of these administrative tasks on the environment, as shown in Table 14.

The configuration consisted of a 4-node Virtual SAN cluster with one disk group (with one caching SSD and six data SSD drives) with 660 linked-clone desktops.

Table 14. VMware View on Cisco UCS C240 M4: Operations Test Results for 660 Linked Clones

Operation Test Results	Provisioning	Refresh	Recompose	Power-on	Deletion
Execution Time	69 min	51 min	101 min	6 min	33 min

Times for these VMware View operations are measured through log entries found at C:\Program Data\VMware\VDM\logs\log-YEAR—MONTH—DAY for the vCenter Server.

Figures 24 through 28 show the CPU utilization during these operations.

Figure 24. Provisioning Operation Graph

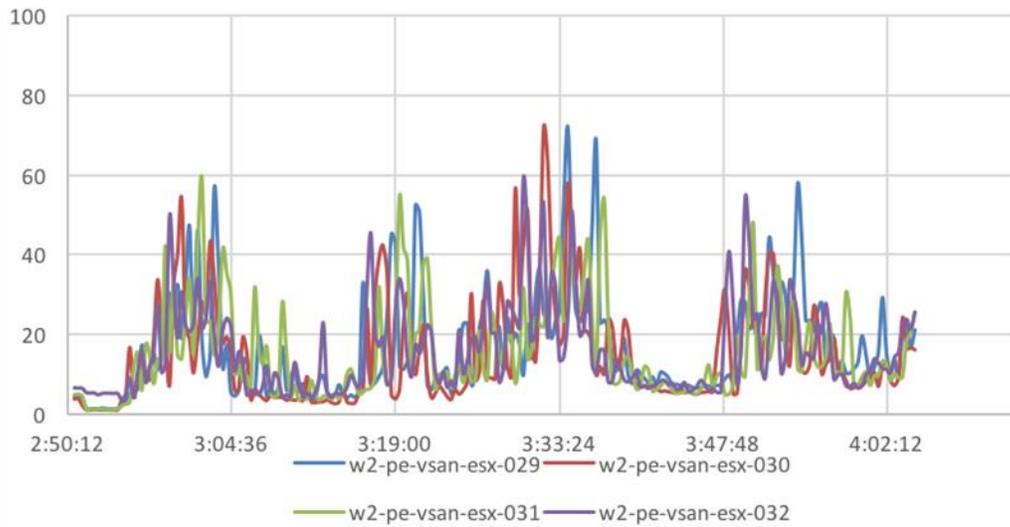


Figure 25. Desktop Refresh Operation Graph

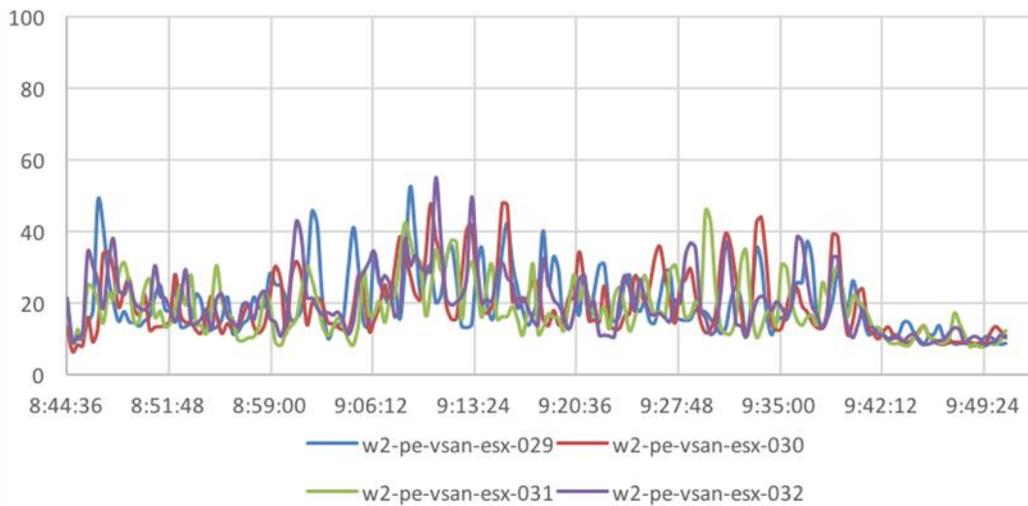


Figure 26. Recomposition Operation Graph

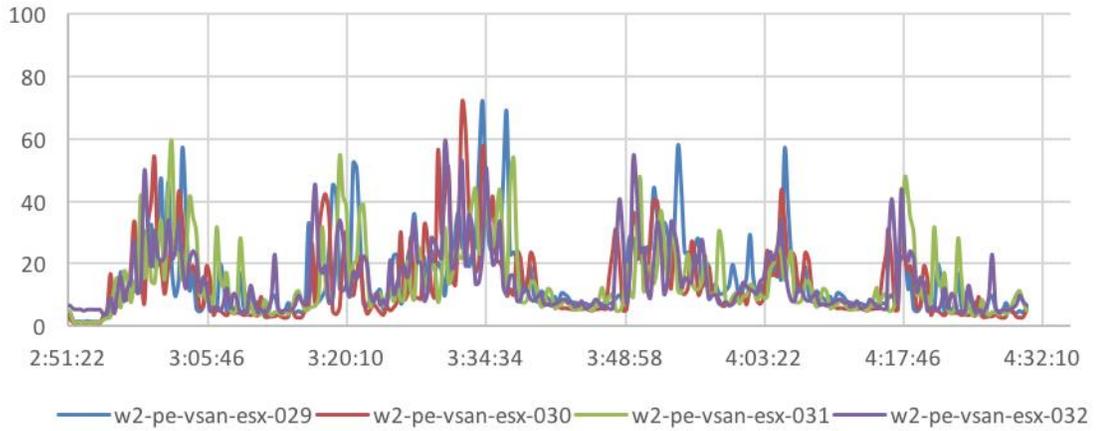


Figure 27. Power-on Operation Graph

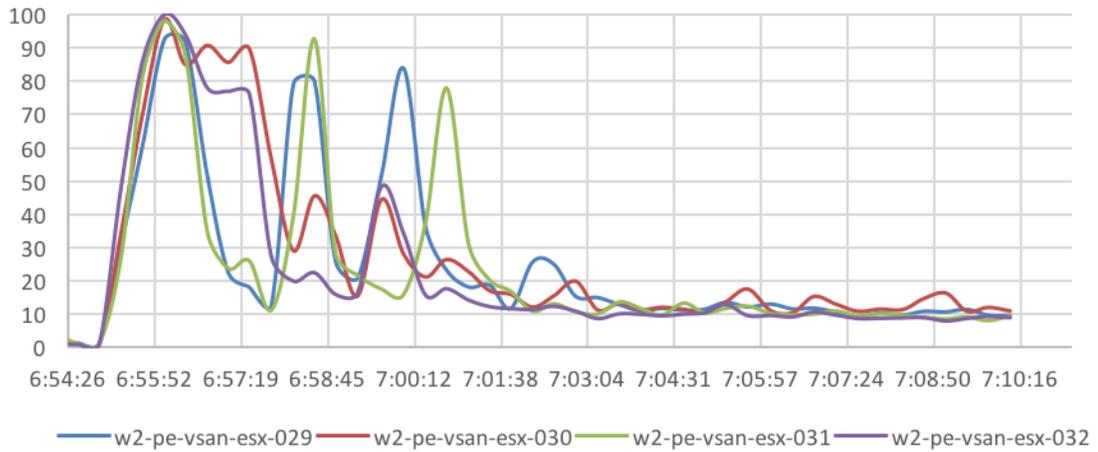
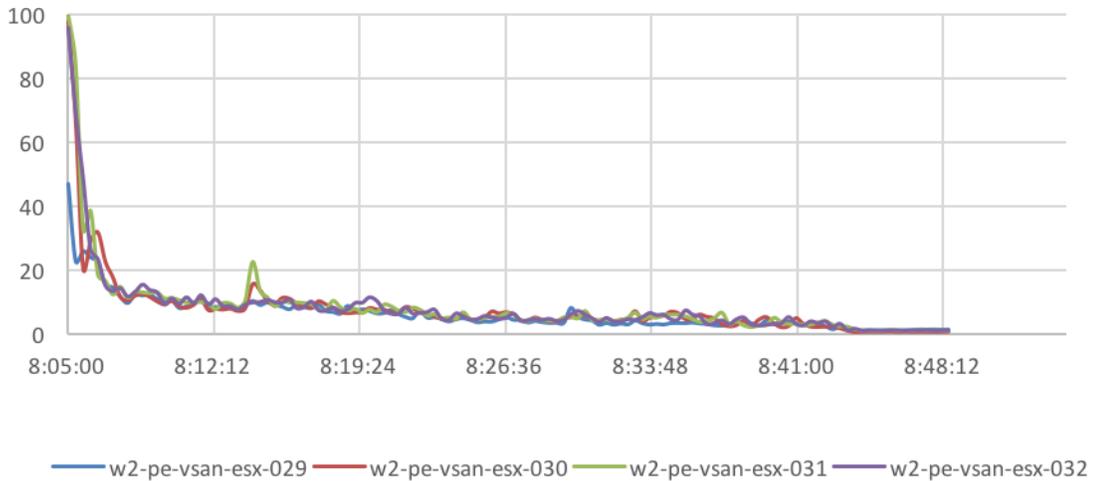


Figure 28. Deletion Operation Graph



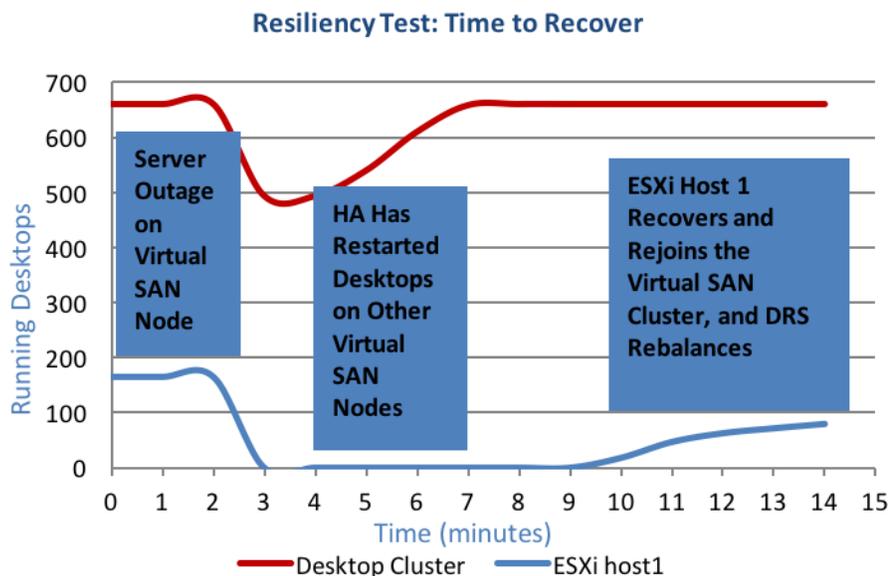
VMware Virtual SAN Failure Simulation

In some cases during ongoing operations in a Virtual SAN environment, failure of either an individual disk or a host may affect virtual machine availability, depending on the storage policies applied. The tests described in this section simulate these failure scenarios to provide an understanding of how Virtual SAN maintains storage data that is highly available under different conditions.

Resiliency Test: One-Node Failure

A single Virtual SAN node hardware failure was simulated for a cluster with four hosts and 660 running virtual desktops, all under simulated workload. Figure 29 shows the results.

Figure 29. Resiliency Testing: Recovery Time for a Single-Node Failure



An ESXi host with 166 running virtual desktops was reset, and all the virtual machines on that node became unavailable. vSphere High Availability restarted the virtual desktops on the other Virtual SAN cluster nodes, and all were ready for user login within four minutes of the simulated failure.

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

VMware Virtual SAN Observer

Virtual SAN Observer is designed to capture performance statistics for a Virtual SAN cluster and provide access through a web browser for live measurements. It also can generate a performance bundle over a specified duration. Virtual SAN Observer is part of Ruby vSphere Console (RVC), which is a Linux console user interface for ESXi and vCenter. RVC is installed on vCenter and is required for running Virtual SAN Observer commands.

Following best practices, an out-of-band vCenter appliance is used in this reference architecture to run Virtual SAN Observer commands. This setup helps ensure that the production vCenter instance is not affected by the performance measurements.

The Virtual SAN Observer commands that were used for this solution are shown in Table 15.

Table 15. Virtual SAN Observer Commands

Virtual SAN Observer Command	Description
<code>vsan.resync_dashboard 10.0.115.72.54 -r 0</code>	Observe data migration while placing hosts in full-migration maintenance mode.
<code>vsan.disk_object_info</code>	Verify disk object information.
<code>vsan.vm_object_info</code>	Verify virtual machine object information.
<code>vsan.disks_info hosts/10.0.115.72.54</code>	Obtain a list of disks on a specific host.
<code>vsan.obj_status_report</code>	Obtain health information about Virtual SAN objects. This command is helpful in identifying orphaned objects.
<code>vsan.reapply_vsan_vmknfc_config</code>	Re-enable Virtual SAN on VMkernel ports while performing network configuration-related troubleshooting.
<code>vsan.observer {cluster name} -r -o -g /tmp -i 30 -m 1</code>	Enable and capture performance statistics used for benchmark testing. For more information, see Enabling or Capturing Performance Statistics Using VMware Virtual SAN Observer .

For a more comprehensive list of Virtual SAN Observer commands, see [VMware Virtual SAN Quick Monitoring and Troubleshooting Reference Guide](#).

System Sizing

The reference architecture uses the sizing specifications described in this section.

Virtual Machine Test Image Build

The virtual machine image was used to provision desktop sessions in the View environment for linked clones (Table 16). The image conformed to testing tool standards and was optimized in accordance with the [VMware View Optimization Guide for Microsoft Windows 7 and 8](#).

The VMware OS Optimization Tool was used to make the changes.

Table 16. Virtual Machine Test Image Build

Attribute	Linked Clones
Desktop operating system	Microsoft Windows 7 Enterprise SP1 (32-bit)
Hardware	VMware Virtual Hardware Version 11
vCPUs	2
Memory	1536 MB
Memory reserved	1536 MB
Video RAM	35 MB
3D graphics	Off
NICs	1
Virtual network adapter 1	VMXNet3 adapter
Virtual SCSI controller 0	Paravirtual
Virtual disk: VMDK 1	24 GB
Virtual disk: VMDK 2 (nonpersistent disk)	1 GB
Virtual floppy drive 1	Removed
Virtual CD/DVD drive 1	–
Applications	<ul style="list-style-type: none">• LoginVSI Version 4.1.4.2<ul style="list-style-type: none">◦ Adobe Acrobat 11◦ Adobe Flash Player 16◦ Doro PDF 1.82◦ FreeMind• Microsoft Internet Explorer 9• Microsoft Office 2010
VMware tools	9.10.0-2476743
VMware View Agent	6.1.1-2769635

Management Blocks

Table 17 summarizes the sizing for the management blocks.

Table 17. Management Block Sizing

Server Role	vCPU	RAM (GB)	Storage (GB)	Operating System	Software Version
Domain controller	2	6	40	Windows Server 2012 R2 (64-bit)	
Microsoft SQL Server	2	8	140	Windows Server 2012 R2 (64-bit)	SQL Server 2012 (64-bit)
VMware vCenter Appliance for Virtual SAN Observer (out of band)	8	24	300	SUSE Linux Enterprise Server (SLES) 11 (64-bit)	vCenter 6.0
VMware View Connection Server	4	10	60	Windows Server 2012 R2 (64-bit)	View Connection Server 6.1.1-2769403
VMware View Composer Server	4	10	60	Windows Server 2012R2 64-bit)	View Composer 6.1.1-2768165
Login VSI server	4	8	60	Windows Server 2012 R2 (64-bit)	Version 4.1.4.2

Host Configuration

Table 18 summarizes the host configuration.

Table 18. Host Configuration

Component	Value
CPU	<ul style="list-style-type: none"> 2-core 2.30-GHz Intel Xeon processor E5-2698 v3 CPU, 135W, 16 cores, 40-MB cache, DDR4, 2133-MHz, with Intel Hyper-Threading enabled
RAM	<ul style="list-style-type: none"> 384 GB (24 x 16 GB)
NICs	<ul style="list-style-type: none"> Cisco UCS VIC 1227 mLOM network adapter (2 x 10-Gbps ports) Firmware version: 2.2(5a) Driver version: enic 2.1.2.71
BIOS	<ul style="list-style-type: none"> C240M4.2.0.6a.0.051220151501
Disks	<ul style="list-style-type: none"> 6 x 800-GB 2.5-inch enterprise performance SAS SSD drive 1 x 400-GB 2.5-inch enterprise performance SAS SSD drive
VMware ESXi version	<ul style="list-style-type: none"> VMware-ESXi-6.0.0-2809209-Custom-Cisco-6.0.0.1
Storage adapter	<ul style="list-style-type: none"> Firmware package version: 24.7.0-0047

Bill of Materials

Table 19 lists the bill of materials for the solution described in this document.

Table 19. Bill of Materials

Area	Component	Quantity
Host hardware	Cisco UCS C240 M4 Rack Server	4
	Intel Xeon processor E5-2680 v3 CPU at 2.50 GHz	8
	16-GB DDR4 2133-MHz RDIMM, PC4-17000, dual rank, 1.2V	96
	Cisco 12-Gbps SAS modular RAID controller with 1-GB cache	4
	Cisco UCS VIC 1227 dual-port 10-Gbps mLOM	4
	32-GB SD card	8
	800-GB 2.5-inch enterprise performance SAS SSD drive	24 (for linked-clones capacity)
	400-GB 2.5-inch enterprise performance SAS SSD drive*	4 (for caching)
Network switch	Cisco UCS 6248UP 48-Port Fabric Interconnect	2
	Cisco Nexus® 9372PX Switch	2
Software	VMware ESXi 6.0	4
	VMware vCenter Server 6.0	1
	VMware Horizon 6.1.1	1
	Microsoft Windows 2012 R2	4
	Microsoft SQL Server 2012 R2	1

***Note:** Enterprise performance SSDs were used for the caching tier due to their availability in our labs; for better price effectiveness, use enterprise value drives as shown on the Virtual SAN Hardware Compatibility List (HCL) for the caching tier.

Table 20 lists available solution bundles for desktop virtualization.

Table 20. Available Desktop Virtualization Bundles

Cisco UCS Virtual SAN–Ready Nodes: Solution Bundle Part Numbers
<ul style="list-style-type: none"> • UCS-VSAN-C240M4-LC-SK1 • UCS-VSAN-C240M4-FC-SK1 • UCS-VSAN-C220M4-LC-SK1 • UCS-VSAN-C240M4-LC-EXP1 • UCS-VSAN-C240M4-FC-EXP1 • UCS-VSAN-C220M4-LC-EXP1 • UCS-VSAN-C240M4-LC-GPUK1-EXP1 • UCS-VSAN-C240M4-LC-GPUK2-EXP1

Conclusion

A solution implementing a VMware Virtual SAN 6.0 All-Flash configuration and VMware Horizon 6 on Cisco UCS C240 M4 Rack Servers provides linear scalability with exceptional end-user performance. A simpler management experience is achieved with Cisco UCS Manager centrally managing the infrastructure. Virtual SAN configuration, management, and monitoring is integrated into VMware vSphere. This solution provides cost-effective scaling for virtual desktop deployments of 3 to 64 nodes.

The reference architecture demonstrates the following main capacities:

- Linear scalability is achieved with Virtual SAN as the storage solution on Cisco UCS for hosting VMware View-based virtual desktops. The reference architecture hosted 660 desktops on four Cisco UCS C240 M4 nodes keeping all aspects of end-user performance consistently acceptable, with disk latency of less than 15 ms and application response times of less than 3 ms as measured by Login VSI.
- Optimal performance is achieved while performing all virtual desktop operations such as refresh, recompose, deploy, power-on, and power-off operations. Times measured for these operations fall within industry-measured benchmarks and testify to the joint solution's scalability.
- To achieve better performance and future expandability and reduce the failure domain in a Virtual SAN environment, choose your disk-group design with care. The use of multiple smaller disk groups reduces the failure domain compared to a single-disk-group design, although this approach adds cost and reduces the number of disk slots available.
- Virtual SAN provides highly available and resilient storage for hosting View virtual desktops. Multiple maintenance and failure scenarios tested demonstrate the resiliency of the joint solution.

For More Information

- [Virtual SAN–Ready Nodes for Cisco UCS](#)
- [Cisco UCS C240 M4 High-Density Rack Server \(SFF Disk-Drive Model\) data sheet](#)
- [Cisco UCS C240 M3 Rack Server with VMware Virtual SAN 5.5 and Horizon 6 with View](#)
- [What's New in VMware Virtual SAN](#)
- [Cisco FlexFlash: Use and Manage Cisco Flexible Flash Internal SD Card for Cisco UCS C-Series Standalone Rack Servers](#)
- [VMware Virtual SAN Compatibility Guide](#)
- [Changing the Default Repair-Delay Time for a Host Failure in VMware Virtual SAN](#)
- [I/O Analyzer](#)
- [Ruby vSphere Console](#)
- [Enabling or Capturing Performance Statistics Using VMware Virtual SAN Observer](#)
- [VMware View Optimization Guide for Microsoft Windows 7 and 8](#)
- [VMware Virtual SAN Quick Monitoring and Troubleshooting Reference Guide](#)
- [Working with VMware Virtual SAN](#)
- [VMware Virtual SAN–Ready System Recommended Configurations](#)
- [Enabling or Capturing Statistics Using VMware Virtual SAN Observer for VMware Virtual SAN Resources](#)



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)