High Performance VDI using SanDisk® SSDs, The VMware Horizon View™ and Virtual SAN
A Deployment and Technical Considerations Guide

May 2014

SanDisk®
a Western Digital brand

vmware®

Table of Contents

1. Executive Summary

This document will provide the reader with a high-level overview of VMware®'s Virtual SAN (Storage Area Network) technology, which is being used to support desktop workloads in a virtualized environment (VDI). It will describe an architecture supporting Virtual Desktop Infrastructure (VDI) that has been demonstrated as an effective alternative to traditional SAN storage for this type of application. The primary focus of this paper will be on meeting the typical challenges associated with the storage subsystem of a virtualized infrastructure, with respect to performance capacity, and operational considerations.

Up to now, the cost of flash storage has kept it primarily relegated for use in improving the performance of the highest priority workloads. SanDisk is changing the economics of this, and is opening up the use of solid-state drives (SSDs) to a broader set of deployments — such as on-line transaction processing (OLTP) systems, VDI deployments, and Analytics (e.g. Big Data) — at more accessible price points.

Desktop virtualization may be implemented in a variety of delivery models, each of which has value for certain roles and responsibilities. This Reference Architecture is targeted at the mainstream knowledge worker who requires access to multiple applications and multimedia. Other architectures may be preferable for mobile devices (e.g. smartphones) and/or remote users having specific needs (e.g. Bring Your Own Device (BYOD)).

VMware Virtual SAN enables server drives in a cluster to be used as a shared storage pool among the vSphere nodes in that cluster.

Solid-state drives (SSDs) have two purposes in Virtual SAN implementations: to provide a read cache and a write buffer. This approach dramatically improves the performance of virtual machines. In some respects, a Virtual SAN can be compared to a number of "hybrid" storage solutions in the market, which also use a combination of SSD and Hard Disk Drive (HDD) storage to boost the performance of the I/O and which have the ability to scale out based on low-cost HDD storage.

VMware Horizon View provides users with remote access to secure virtual desktops that are stored in the data center. This provides high availability for desktop services, as well as ensures greater business agility through flexible provisioning of those services, compared with those offered by traditional PCs, while greatly reducing the desktop total cost of ownership (TCO).

2. Introduction

2.1 Audience

This reference architecture is part of the SanDisk's Enterprise Storage Solutions library of reference documents and is intended for use by individuals who are responsible for architecting, designing, managing, and/or supporting virtualized infrastructures. Consumers of this document should be familiar with concepts pertaining to VMware vSphere®, VMware Horizon™ View, and basic server and storage technology.

2.2 Purpose

This document will cover the following subject areas:

- Solution Overview – VMware and SanDisk

    - VMware® Virtual SAN™

    - VMware Horizon View

    - SanDisk SSDs

- Solution Architecture and Deployment

    - Test Results & Environment

    - Configuration

- Deployment Best Practices

    - VMware Horizon View

    - VMware vSphere 5.5

    - SanDisk SSDs

- Technical Support and additional resources

2.3 Customer Types

Customers that are considering implementing a Virtual Desktop Infrastructure (VDI) in order to reduce and/or contain their IT support costs are the focus of this Reference Architecture. In addition to virtualizing the desktop, this approach uses VMware's Virtual SAN – a software-defined-storage architecture – to further optimize usage of the server and storage infrastructure.

2.4 Workloads Targeted

This Reference Architecture was designed and tested against the virtual desktop of a knowledge worker using applications such as Microsoft® Office (Word, Excel, PowerPoint); web browsing (Internet Explorer, Firefox); email (Outlook), Adobe for document formatting; and multimedia applications (streaming audio/video). Details regarding the workload specifics can be found in the Test Environment section and Appendix of this document.

3. VMware and SanDisk – Solution Overview

Virtual SAN is a new storage solution from VMware that is fully integrated with vSphere. It automatically aggregates server drives in a cluster to create shared storage that can be rapidly provisioned from VMware vCenter™ during VM creation. It is a platform for VM Storage Policies designed to simplify virtual machine (VM) storage placement decisions for vSphere administrators. It is integrated with core vSphere features such as vSphere HA, vSphere DRS and vMotion. It can also be used to address quality of service (QoS) requirements by creating VM Storage Policies that define the level of performance and availability required on a per–virtual machine basis
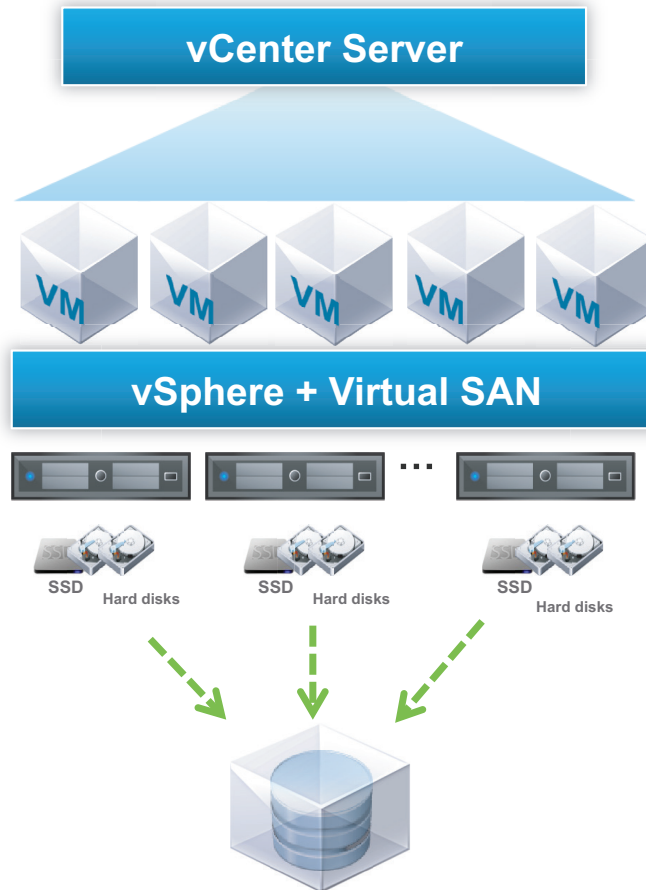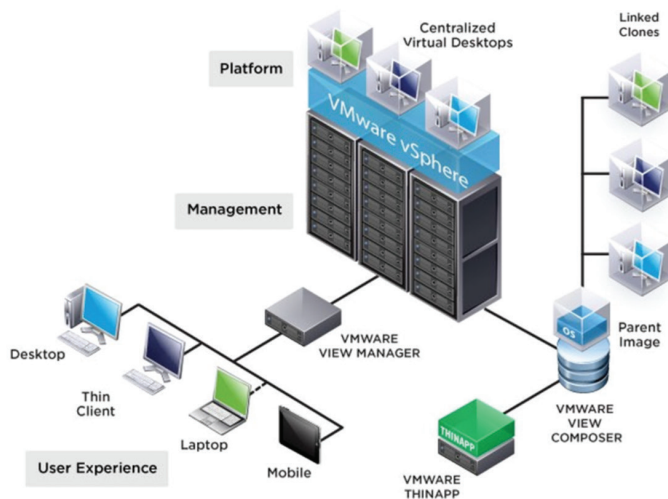

VMware Virtual SAN High Level Diagram

SanDisk SSDs are a valuable enhancement to Virtual SAN environments. They allow for the acceleration of I/O among the VMware ESX servers – using SSDs for high performance read/write caching and HDDs for cost-effective data persistence. Each vSphere host must have at least one SSD when participating in the Virtual SAN cluster. The more SSD capacity the host has, the greater the performance boost, because more I/O can be cached to the SSD.

VMware's Horizon View is used for VDI, which connects multiple types of end user devices to their own personal virtual desktop running in the vSphere infrastructure in the data center.

Horizon View Desktop Agent sends screen updates to the end user's device, while the end user sends back keyboard and mouse movements to update the applications being hosted at the central site. The end user's Horizon View desktop consists of that person's operating system, applications and data – all stored securely in the company's private data center. Horizon View offers linked VM clones to save disk space when creating multiple VMs based on the same VM golden image.

VMware Horizon View High Level Diagram

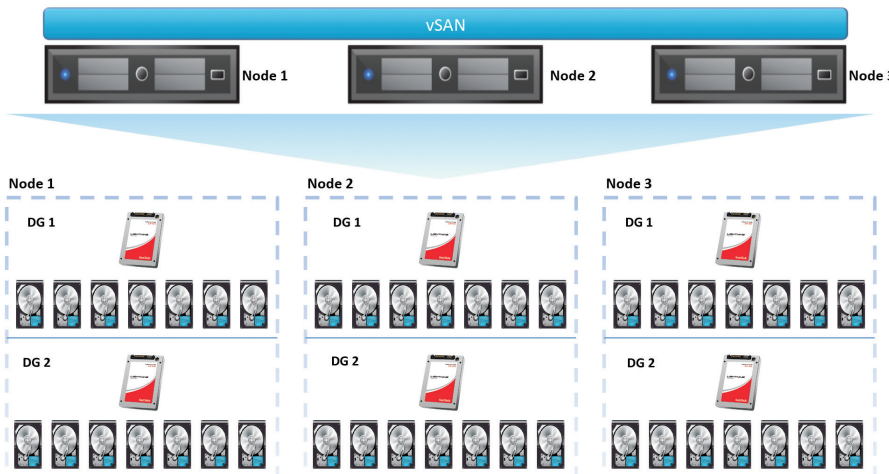**VMware Horizon View High Level Diagram**

4. Solution Architecture and Deployment

The following section will discuss in detail about the testing process, results and configuration used to validate the solution.

3 Node Virtual SAN Diagram

- Disk Groups (DG) are created in Pass-Through mode in each node.
- Each node (ESXi host) configured with 2 DG. Total 6 DGs are presented to VSAN datastore.

4.1 Testing Process and Workload Used:

For the validation testing in the Virtual SAN environment, we used the View Planner standard benchmark workload and used the Group A (CPU sensitive) and Group B (CPU and IO sensitive) score to determine the Virtual SAN environment capability for hosting VDI.

The view desktop is created using Windows® 7, 64-bit operating system standard image. Necessary configuration changes are done according to the View Planner installation and configuration user guide. All the applications, part of the View Planner pre-selected workload requirements, are installed inside this image.

View client desktop image is created with Windows® 7, 32-bit operating system and configured, according to the View Planner installation and configuration user guide.

A three-node Virtual SAN cluster is created, and floating pool view desktops are provisioned on top of it. View clients' desktops and other Horizon View infrastructure VMs, such as View Planner Appliance, vCenter, AD-DNS, DHCP, VMware Horizon View, are provisioned in separate data store and computing nodes so that View desktop performance can be measured on top of Virtual SAN data store and computing nodes.

4.2 Results Observed:

The following workloads were executed to get the results

- VMware View Planner Standard Benchmark – This is a View Planner that provided the standard benchmark run. Our test results met the 95th percentile response time QoS (quality of service) threshold (CPU sensitive operation < 1 sec., and Disk and CPU sensitive operation < 6 sec.)

- No Network Operation Workload – In this workload, we excluded Internet Explorer and Firefox-related operations, and executed the run. In this case as well, test results met the 95th percentile response time QoS threshold (CPU sensitive operation < 1 sec., and Disk and CPU sensitive operation < 6 sec.)

4.3 Test Result 1:

We have run the View Planner "Standard Benchmark_5i" workload for 276 VM in a three-node Virtual SAN cluster. The 95th percentile response time for Group A, and Group B met the QoS (quality of service) threshold (<1 sec. for Group A and <6 sec. for Group B). Although the three-node cluster provisioned was having an adequate amount of disk capacity available to accommodate more desktops, but once the CPU usage level was within 95-100% range, we didn't scale any further with the "Standard Benchmark_5i" workload.

In this test; View Clients, View Desktops, and View Planner harness were configured to run in the local area network (LAN) environment and all the virtual machines (VMs) were configured to run from same virtual local area network (VLAN) switch. Following are the application response time for Group A and Group B workload.
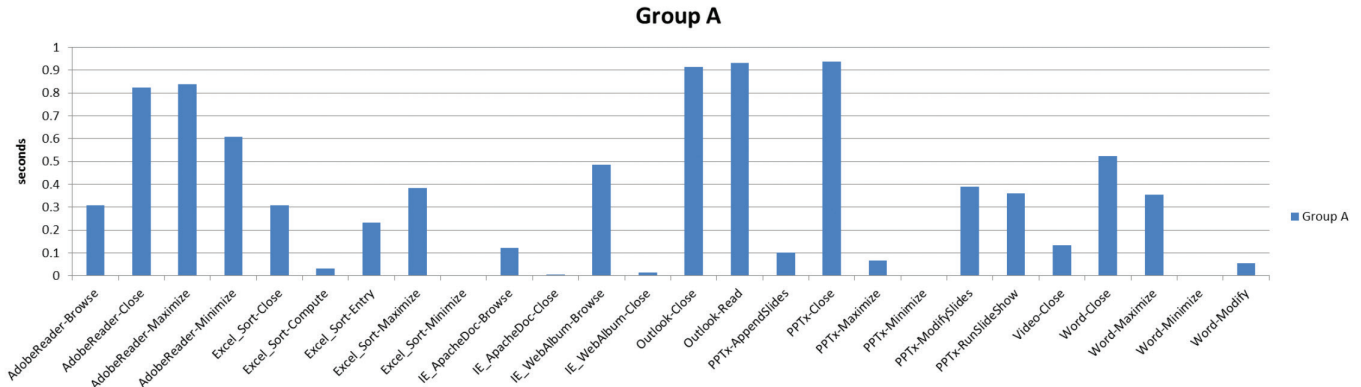

Fig 1: Application Response Time of "Group A" Operations for 276 VM
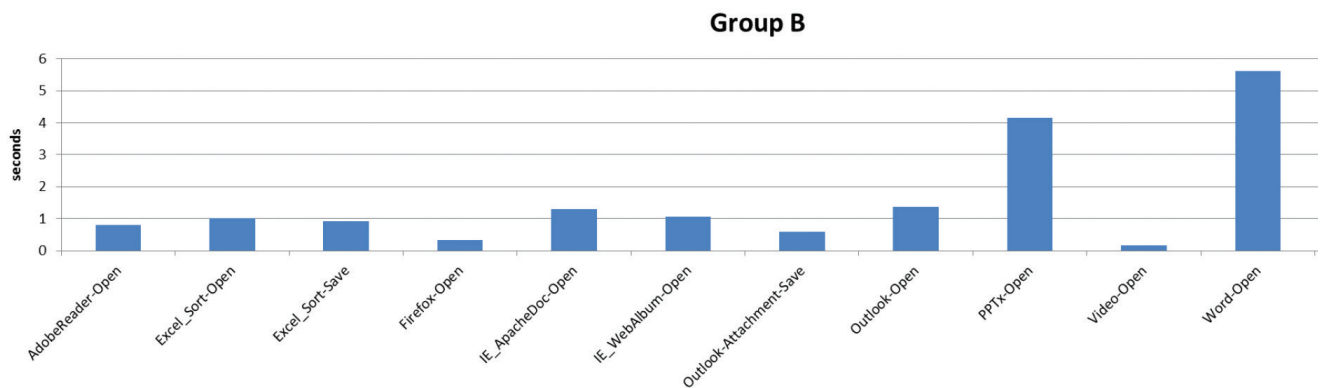

Fig 2: Application Response Time of "Group B" Operations for 276 VM

Further to application matrices, the below graph shows the utilization of CPU, drive and latency at drive level captured using View Planner report.
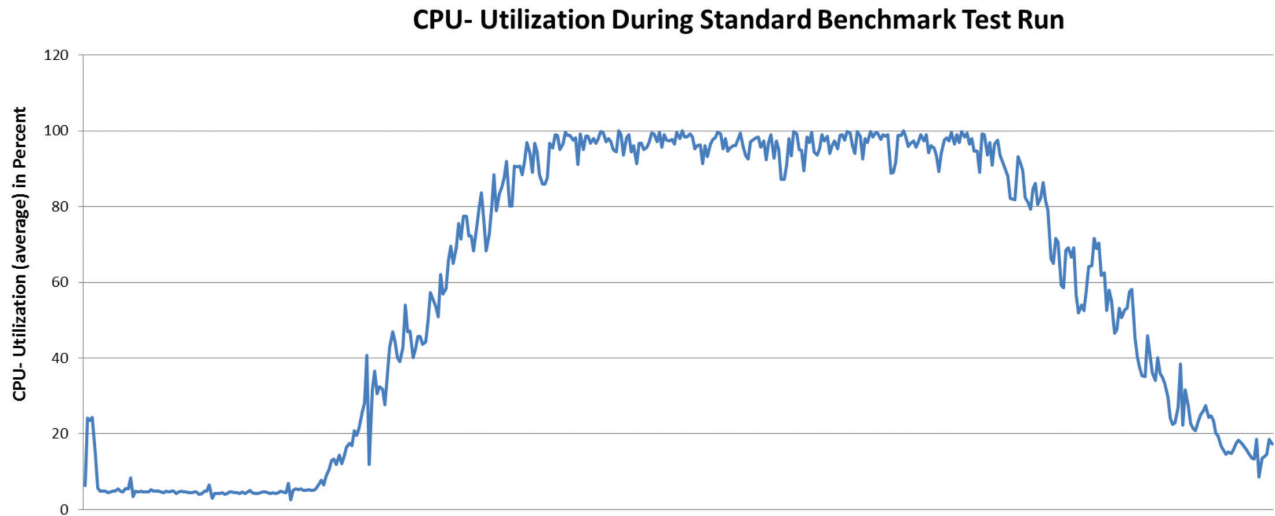

Figure 3: CPU Utilization on one host during Standard Benchmark workload run
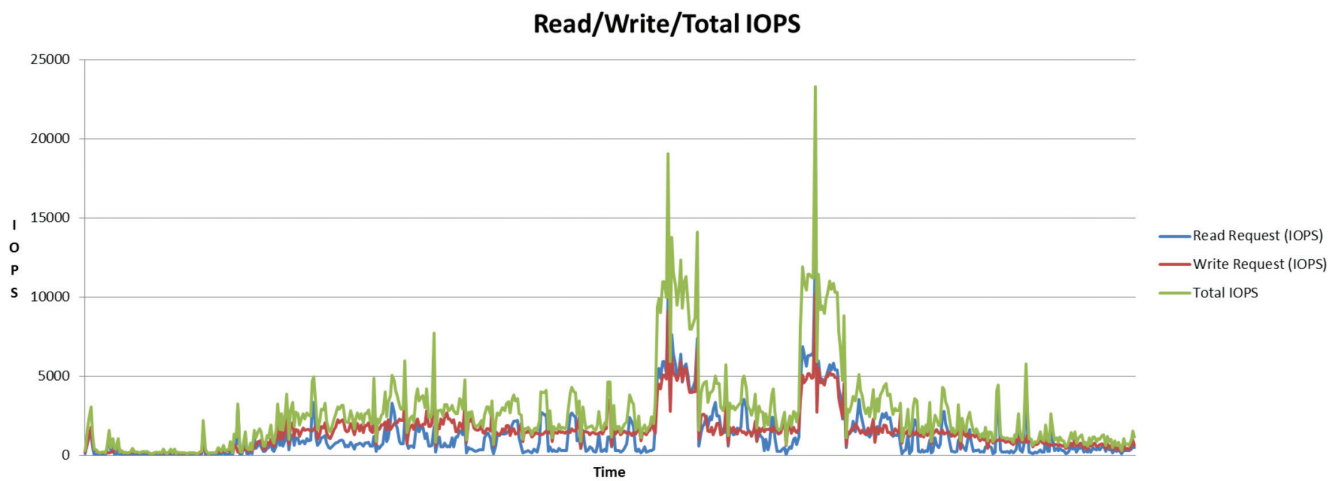
Figure 4: Drive Utilization on one host during Standard Benchmark workload run
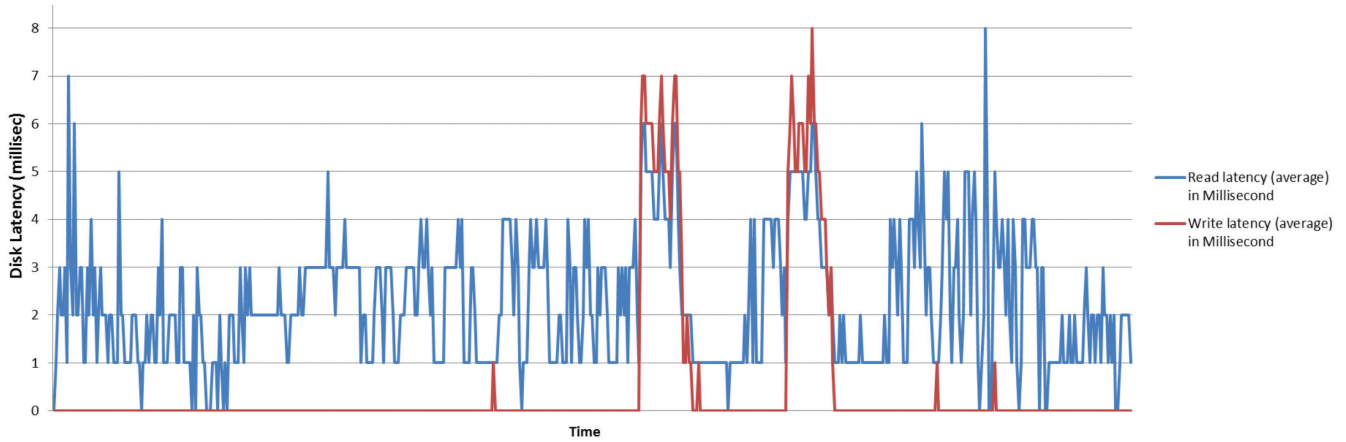
Figure 5: Drive Latency on one host during Standard Benchmark workload run

4.4 Test Result 2:

In this test we excluded the Internet Explorer (IE) and Firefox-based browser workload. The goal of this test is to run the workloads that are executed within the desktops and it does not include any network-related activities. We created a profile with all the preselected application workload (excluding all the workload containing either IE or Firefox operation) and set five iterations with 5-second think time.

Figure 6: Application Response Time of "Group A" Operations for 300 VM

Figure 7: Application Response Time of "Group B" Operations for 300 VM

4.5 Test Bed / Configuration

3 Node Virtual SAN Diagram

- Disk Groups (DG) are created in Pass-Through mode in each node.
- Each node (ESXi host) configured with 2 DG. Total 6 DGs are presented to VSAN datastore.

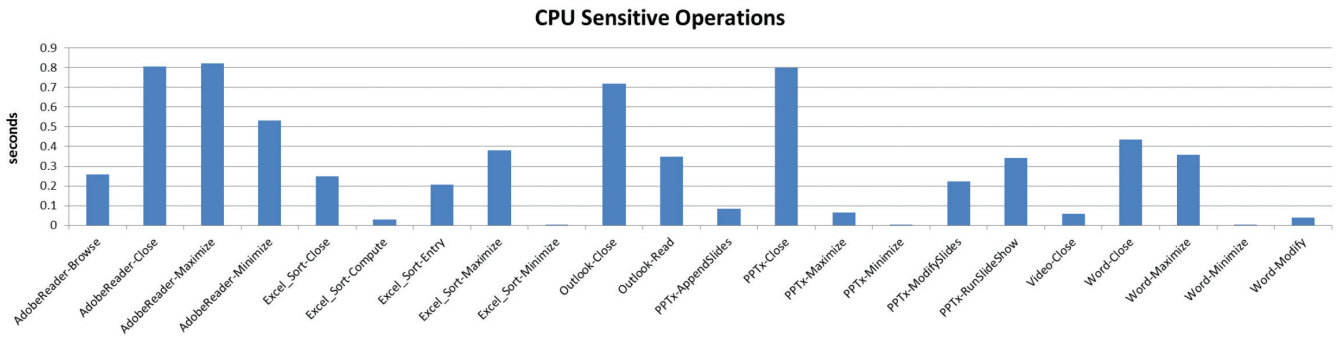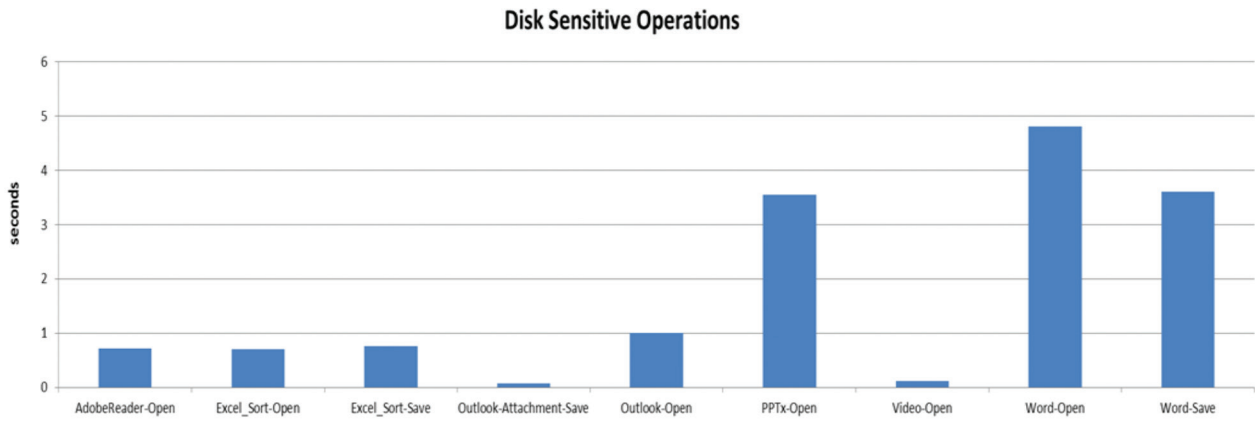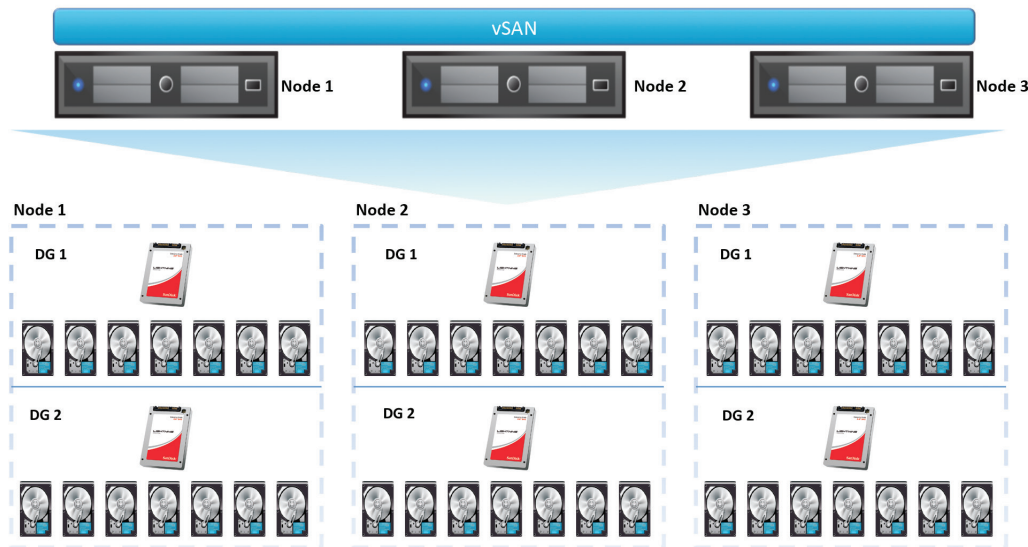A three-node Virtual SAN cluster is built using Virtual SAN beta version (build number 1439689). The Virtual SAN Datastore is presented to deploy the Horizon View desktops. The remaining VMs such as VMware Horizon View Manager, vCenter Server, and Composer are deployed on different clusters. The View Planner appliance and View Clients drive the workload and are configured to run from separate clusters. This is to ensure that Virtual SAN Datastore only carries the workload that is related to VMware Horizon View desktops.

Two disk groups are created in each node with one SanDisk SSD and seven HDDs. No storage policies were applied (default unchanged) for VDI desktops to achieve redundancy or tolerance.

The view client VMs, and Horizon View Desktop VMs are configured in a distributed switch and run from the same VLAN. The infrastructure components are configured in different VLANs, but were part of same LAN infrastructure. The View desktops connections are made using PCoIP protocol.

Using VMware Horizon View administrator, the default desktop pool is created and the workload is driven from the View Planner harness for different combination, as mentioned in the "Test Results" section above.
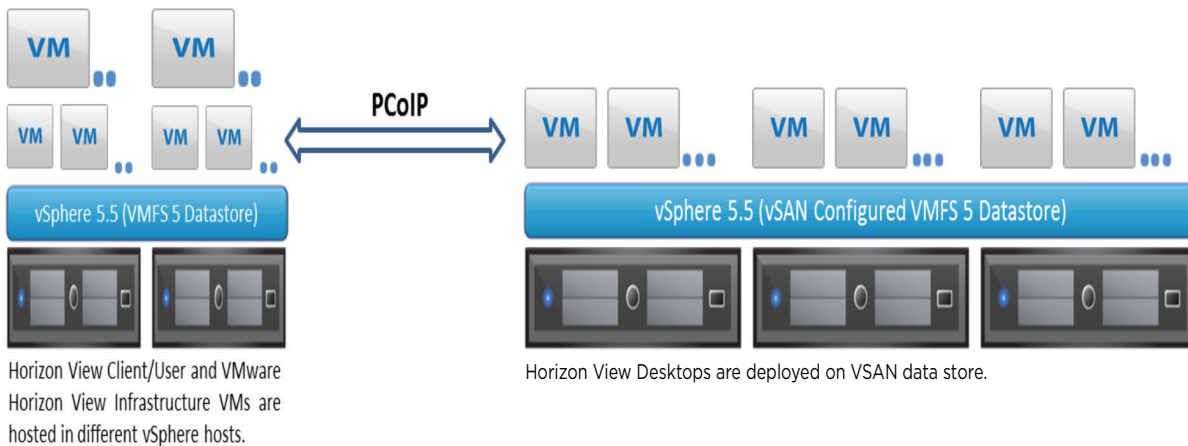
View Planner Test Execution Setup

Following table describes the VM configuration used in the testing:

VM Name	OS	Role(s)	Memory (GB)	No. of vCPU
AD-DNS-DHCP	Windows 2008 - 64 bit	DHCP server, Active Directory and Domain	4	1
Horizon View Manager	Windows 2012 - 64 bit	VMware horizon view connection manager	8	2
View Composer	Windows 2012 - 64 bit	View Composer, SQL Server 2008	4	2
vCenter	Windows 2008 R2 - 64 bit	vCenter for Horizontal View	32	4
View Desktop	Win 7 - 64 bit	Horizon View End User Desktops	2	1
View Client	Win 7 - 32 bit	Horizon View End User Clients	2	2
VMware View Planner	Cent OS - 32 bit	VDI Workload Generator	3	1

5. Deployment Best Practices

5.1 Pass-through vs RAID

Virtual SAN uses no hardware or software RAID. Instead, RAID type functionality is performed at the VM level. The preferred drive-set up for Virtual SAN is to use JBOD disks, with a pass-through capable controller. Pass-through is the optimal configuration for Virtual SAN; however, Virtual SAN drives can also be configured as RAID0 drives. Usage of RAID in a Virtual SAN configuration may allow you to leverage existing infrastructure that may not have a pass-through capable controller, although RAID configurations will increase set-up complexity.

5.2 VMware ESXi™ on SD card to make more room for storage

This approach is used in order to save the server-side slot by putting the ESXi binaries on a SanDisk SD card. With the SD card, all the slots available in the server can be used to configure the disk groups. It allows the server to support more disks for greater capacity.

5.3 Number and Depth of Drive Group

Although there is no standard rule of thumb, using more drive groups will provide more total IOPS in the configuration. The SSD in a drive group is meant for "read cache" and "write buffer." The ratio of SSDs to HDDs can be altered directly within a drive group, based on capacities of each, and also by the number of drive groups that are created. Determining the optimal configuration will depend on application behavior. This choice of configurations is very relative, and requires careful consideration of the workload that is supported.

5.4 Fault Tolerance of View Desktops

Virtual SAN provides storage policies for setting the redundancy and tolerance levels of VMs. In general, VDI environments require much less fault tolerance than general-purpose server-side applications – where the failure of the application will affect many end users. The level of fault tolerance needed for your VDI environment is dependent upon how critical these desktops are to your organization. It should be noted that the lower the need for fault tolerance (i.e. redundant VMs), the higher the density of VMs that any given physical server can support. As a best practice, VDI redundancy can be kept at a minimum, and some may choose to deploy no redundant VMs for VDI at all.

5.5 Virtual SAN "Tips":

Virtual SAN 101	

• Minimum of 3 compute nodes – max of 32
• Up to 5 drive groups per compute node
• Each drive group has:
  - 1 x SSD
  - Up to 7 HDD
	• Drives do not use RAID
	• Level of redundancy is a VM level storage policy
	• SSD is a READ cache and WRITE buffer
	• # & depth of drive groups defines performance

6. Conclusions

VMware's Virtual SAN technology has been demonstrated to be an extremely effective alternative to traditional SAN storage for a VDI environment. Using VMware's View Planner Standard Benchmark, our test results met the requirements for VMware's VDI performance measures – achieving 95th percentile response time for QoS (CPU-sensitive operation < 1 sec., and Drive and CPU-sensitive operation < 6 sec.).

Similarly, in running VMware's No Network Operation Workload (excluding Internet Explorer and Firefox-related operations), our test results met the 95th percentile response time QoS (CPU-sensitive operation < 1 sec., and Drive and CPU-sensitive operation < 6 sec.).

The three-node Virtual SAN test environment scaled to 300 virtual desktops – at which point, CPU usage was within 95-100% that prevented additional VDI support. If additional CPU resources were added, the storage configuration used in this test could successfully achieve - (only 100 desktops per host is currently supported on VSAN) more than 300 desktops.

The use of SD cards as the ESXi boot device had the result of increasing the number of storage slots available on each host, potentially eliminating the need for external storage.

Our testing was run on pre-release Virtual SAN software. At time of this publication, Virtual SAN had achieved General Availability and the range of the hardware compatibility list (HCL) has been greatly increased and additional capabilities have been added.

– All of which has the result of simplifying the deployment of Virtual SAN into existing IT infrastructures.

7. Acknowledgements

Martyn Wiltshire and Felix Najer of SanDisk IT were early beta-site users of Virtual SAN and were instrumental in the setup, testing and benchmarking of the cluster. Biswapati Bhattacharjee of SanDisk Technical Marketing provided invaluable assistance in the tuning and virtualization areas. Dave Montgomery of SanDisk Product Management contributed to the content and review of this paper, as well as providing overall guidance. We would also like to thank VMware's Performance Benchmarking group and Wade Holmes, VMware Technical Marketing, for their thorough review of our test methodologies and results.

Appendix A – Resources:

1)   SanDisk Enterprise Storage Solutions
       http://www.sandisk.com/Enterprise

2)   VMware Virtual SAN Software
       http://www.vmware.com/products/virtual-san/

3)   VDI Solution – VMware Horizon View
       http://www.vmware.com/products/horizon-view/

4)   VDI benchmark - VMware View Planner
       http://www.vmware.com/products/view-planner

5)   Using SSDs in specific industry solutions
       http://www.sandisk.com/enterprise/industry-solutions/

6)   General Information about SSDs in Enterprise IT – SanDisk IT Blog
       http://itblog.sandisk.com/

Appendix B - VMware View Planner 3.0.1 Overview:

VMware View Planner is a VDI workload generator that automates and measures a typical office user's desktop activity. The automated applications are Microsoft® Office, PDF browse, watching a video - and the operations on these applications are opening a file, browsing the web, modifying files, saving, closing, and such.

The central piece in the View Planner architecture is the Harness, the automated controller. It controls everything, from the management of participating desktop and client virtual machines, to starting the workload and collecting results, to providing a monitoring interface through a web user interface.

The harness is a CentOS Linux-based appliance virtual machine (VM) that interacts with Horizon View Manager components. It also runs a web server to present a user-friendly web interface.

As shown in the diagram, the appliance interacts with a VMware vCenter and Horizon View Manager, and with Active Directory/DNS to control desktop virtual machines and clients/users. It also communicates with client virtual machines to initiate remote protocol connections such as PCoIP, RDP, etc.

The appliance is responsible for starting the workload simulation in desktop virtual machines.

Upon completion, results are uploaded and stored in a database inside the appliance. Results can be viewed using the web interface or extracted from the database at any time.
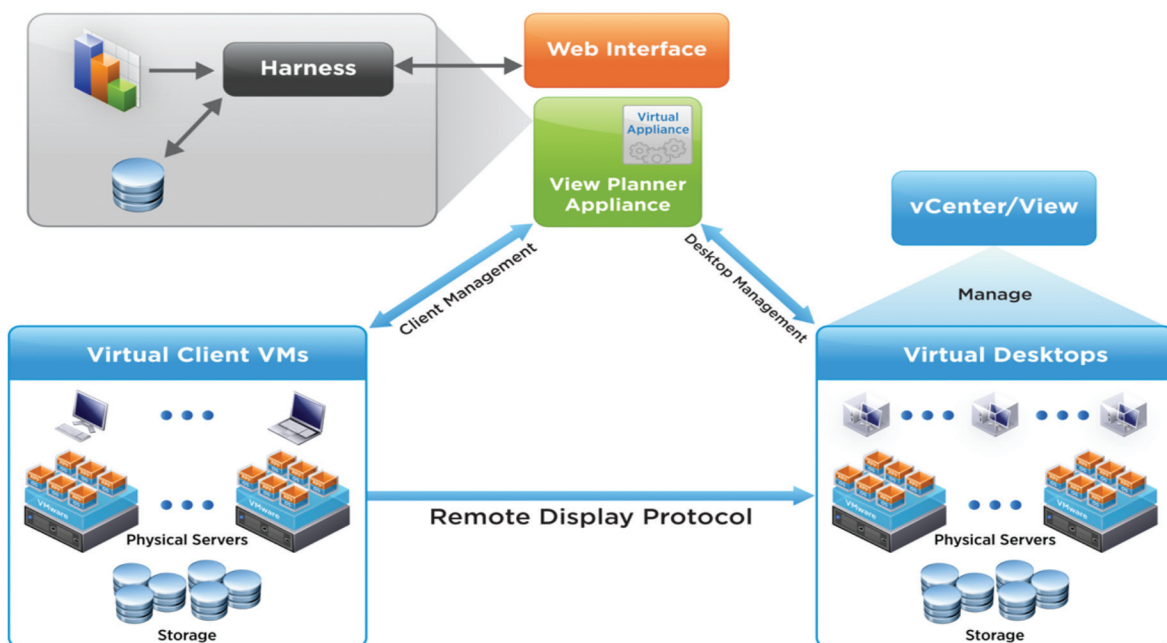

Fig 8: VMware View Planner Overview

Appendix C - Test Setup Configuration Details:

		MILEXLID01	MILEXLID02	MILEXLID03	
		DELL R720	DELL R720	DELL R720	
		Ivy Bridge	Sandy Bridge	Sandy Bridge	
Cores		20@2.8 GHz	16@2.7 GHz	16@2.7 GHz	
RAM		260 GB	260 GB	260 GB	
OS (on 2 x SD)		2 x 16 GB - RAID1*	2 x 16 GB - RAID1*	2 x 16 GB - RAID1*	
Drive Group 1	**# of SSD**	1	1	1	
	SSD Model **SSD Size (GB)**	Lightning® Mixed-Use (800GB)	Lightning® Mixed-Use (800GB)	Lightning® Mixed-Use (800GB)	
	# of Spindles **Spindle Spec** **Size of Spindles**	7 10K RPM SAS Dell 600	7 10K RPM SAS Dell 600	7 10K RPM SAS Dell 600	
Drive Group 2	**# of SSD**	1	1	1	
	SSD Model **SSD Size (GB)**	Lightning® Mixed-Use (800GB)	Lightning® Mixed-Use (800GB)	Lightning® Mixed-Use (800GB)	
	# of Spindles **Spindle Spec** **Size of Spindles**	7 10K RPM SAS Dell 600	7 10K RPM SAS Dell 600	7 10K RPM SAS Dell 600	
Total Usable (GB)		8400	8400	8400	25200
Total Usable (TB)		8.4	8.4	8.4	25.2

*Controller used was DELL H310 in pass through mode.

Appendix D - SSD versus HDD – What's the difference?

Enterprise storage needs are rapidly changing, causing the technology infrastructure required to meet these needs to evolve continuously. Hard Disk Drives (HDDs) are cheap, plentiful, have lots of capacity and are usually purchased as a bulk commodity. They will remain a key element of enterprise IT infrastructure for the foreseeable future. HDD technology offers a range of price/performance options (and choices between 15K, 10K, 7.2K or 5.4K RPM drives), but the technology for the drives remains fundamentally the same - rotating disks and a magnetic head on a moving arm to read and write data.

HDD technology is well-suited for many storage scenarios, but it has some intrinsic limitations. Data is stored on spinning platters with "read/write" arms that have to physically move across specific sectors of the hard drive. As a result, HDDs are inherently good at reading data that is stored sequentially within and across sectors, because the movement of the "read/write" arm is minimized.

But the nature of today's computing makes it increasingly likely that data will be stored wherever there is space on the drive. The increasingly random nature of data storage is being driven by virtualization—and it can cause significant performance problems for traditional hard drives.

"Virtual machines" (VMs) create, store and access data across a pool of virtualized storage. That data is stored wherever the capacity exists, in a random fashion. As the number of VMs trying to access data simultaneously increases, it creates a problem known as the "I/O blender" effect – resulting in data being stored in an increasingly random fashion, forcing the read/write arm of a traditional HDD to move constantly. And each one of those moves adds time delays to the system. As the number

of VMs increases and as the amount of data being searched grows, the delays quickly add up. This creates a bottleneck in storage – and reduces the responsiveness of data center applications.

The good news for data center managers and for storage experts is that SSDs present an attractive alternative to HDDs. Because they have no moving parts, SSDs remove one of the major drawbacks of HDDs. SSDs also eliminate sources of latency, as well as heat. The physics of how SSDs access data means that they're better suited for today's random-access data needs.

- SSDs are fast, much faster than HDDs for most read and write operations—as much as four to five times faster (4-5 X faster) on reads/writes per second. As compared to 7200 RPM SATA 2.5" hard drive. Based on published specifications and internal benchmarking tests.

- SSDs excel at random reads—so they aren't as susceptible to latency caused by the demands of virtualization. The result is more predictable performance as an infrastructure scales without overprovisioning.

- SSDs are ideal for weathering events like a "boot storm," a phenomenon that occurs when many users are trying to log in simultaneously. HDDs suffer from boot-up latency that makes them especially susceptible during "boot storm" events. However, SSDs' speed and capacity means that boot storms are less likely to occur than with HDD deployments.

SSDs are more reliable than HDDs. Unlike HDDs, SSDs have no moving parts. That makes them more resistant to shock, vibration and temperature variations. As a result, they have a higher mean time between failure (MTBF) rating than HDDs.

Types of SSDs – Read-Intensive, Mixed-Use, and Write-Intensive

The core component of an SSD is NAND flash memory. The two basic types of NAND flash are SLC (single-layer cells) and MLC (multi- layer cells). SLC NAND stores one bit of data per cell with faster write times and longer endurance. MLC NAND stores multiple bits of data per cell with slower write times and less endurance. However, SanDisk's proprietary controller and advanced parallel processing architecture optimize the way that the bits are written across the MLC NAND, reducing wear and lengthening product life – improving overall durability.

Unlike magnetic media on HDD storage, data stored on flash needs to be erased before new data can be written or "programmed" — this is known as the Program-Erase Cycle (PEC). The maximum number of PECs a drive can sustain before its performance and capacity degrades is dependent on two items – 1) the choice of SLC or MLC NAND and, 2) the amount of overprovisioning – providing extra NAND capacity to automatically replace worn out cells. Cell endurance is defined as the number of times the media can be rewritten (programmed and erased), and it is typically on the order of several thousand PECs per NAND cell.

Since program-erase cycles are difficult to translate into data storage terminology, SSD endurance is specified by vendors as the number of Drive Writes per Day (DWPD). One DWPD is the equivalent of "filling up" the drive to maximum capacity, then erasing the entire drive – this process requires a PEC for every NAND cell on the SSD.

The three classifications for enterprise-grade SSDs that SanDisk produces are Read-Intensive (RI), Mixed-Use (MU) and Write-Intensive (WI). The primary distinctions among these drive types are their endurance specifications, SAS or SATA interface, and cost. In general, as seen in this short table:

Read-Intensive   =   Less than 3 DWPD
Mixed-Use   =   Between 3 and 10 DWPD
Write-Intensive   =   Greater than 10DWPD

When to use RI, MU or WI SSDs

Every application workload in your data center has different I/O characteristics, because different applications have different read/write ratios and different I/O patterns – sequential or random. Understanding the I/O characteristics of your application workloads is critical to selecting an SSD that best meets your application requirements.

SSD storage is more cost-effective than HDD storage for read-intensive applications (RI) with random I/O patterns, whereas Write-Intensive (WI) applications with predominantly sequential I/O patterns tend to be better suited to HDD storage. Furthermore, understanding the application I/O characteristics will also help you in choosing the appropriate SSD storage (RI vs. MU vs. WI).

Here are some real-world examples to guide your choice of SSDs:

Read-intensive (RI) scenarios—SSD technology is ideal for online content providers, especially video-on-demand (VOD) and image-retrieval applications. With thousands of users accessing high-definition (HD) content simultaneously, the customer experience is completely dependent on how quickly data can be accessed and delivered. RI SSDs provide high capacity, typically with lower performance and a lower cost per GB of capacity. They are often used in applications for which a large amount of data needs to be stored and accessed (read) frequently.

Mixed-Use (MU) scenarios—Email has become the default means of communication for most businesses and organizations today. Ensuring that data gets in and out of the system in a timely manner is critically important to keeping business running. Email is more frequently read (around 70%) than written (30%), and it is an ideal workload for MU SSDs.

Write-intensive (WI) scenarios—For On-line Transaction Processing (OLTP) or analyzing search logs from on-line applications, the speed with which data can be written to a storage environment becomes critical. Delays can result in transactions that "time out" or that simply deliver a poor experience to customers or employees. WI SSDs provide higher performance and endurance, but they typically carry a higher cost per GB of capacity. They are frequently used in environments in which only a few drives are needed for high I/O performance requirements.