FlashStack Converged Infrastructure Solution

For VMware Horizon View 6.0

With Pure Storage FlashArray 400, VMware vSphere 5.5, Cisco UCS B-Series blade servers and Cisco Nexus 5500-series Switches

Ravi Venkat, Solutions Architect, Pure Storage

Version: 1.0

Table of Contents

Introduction

Pure Storage has introduced a converged infrastructure platform named FlashStack which is built upon trusted hardware from Cisco and Pure Storage. FlashStack leverages Cisco's extremely flexible and expandable Unified Computing System to provide the compute horsepower, Cisco Nexus Switches for networking and Pure Storage's own FlashArray 400 Series as the storage foundation. Enhanced by a single point of contact for support, FlashStack will provide a modular and powerful infrastructure that is simple to use, deploy and maintain while providing an architecture that suits a wide variety of application requirements.

This document describes a reference architecture for deploying VMware vSphere environments on Pure Storage FlashArray (FA-400 series) using VMware vSphere 5.5U1, VMware View 6.0, VMware View Planner 3.0.1 (a performance testing tool for Virtualized desktop simulation), and Microsoft Windows 7 as the guest operating system. Pure Storage has validated the reference architecture within its lab – this document presents the hardware and software configuration, the workload configuration and the results and further offers implementation, sizing guidance.

Goals and Objectives

The goal of this document is to showcase the ease of deploying a large number of virtual desktops on the Pure Storage FlashArray. We will demonstrate the scalability of VMware View based virtual desktop building blocks with VMware View Planner as a performance benchmark tool running Standard benchmark profile(heavy user workload). VMware View based Windows 7 desktops will be deployed incrementally in number to find the high-end scalability characterization of the FlashStack for a single server, and later an entire chassis with particular focus and emphasis on storage. In addition, we highlight the benefits of the Pure Storage FlashArray including data reduction and low latency that directly impacts the user experience and provides customers a real alternative for their large-scale VMware deployment projects.

Audience

The target audience for this document includes storage and virtualization administrators, consulting data center architects, field engineers, and desktop specialists who want to implement VMware-based virtual desktops on the Pure Storage FlashArray. A working knowledge of VMware vSphere, VMware View, VMware View Planner, server, storage, networks and data center design is assumed but is not a prerequisite to read this document.

Reference Architecture Design Principles

The guiding principles for implementing this reference architecture are:

- **Repeatable**: Create a scalable building block that can be easily replicated at any customer site. Publish the version of various firmware under test and weed out any issues in the lab before customers can deploy this solution.
- **Virtualized**: Implement every infrastructure component in a VM.
- **Available**: Create a design that is resilient and not prone to failure of a single component. For example, we include best practices to enforce multiple paths to storage, multiple NICs for connectivity, and high availability (HA) clustering including dynamic resource scheduling (DRS) on vSphere.
- **Efficient**: Take advantage of inline data reduction and low latency of the Pure Storage FlashArray by pushing the envelope on VMs per server density.
- **Simple**: Avoid tweaks to make the results look better than a normal out-of-box environment.

Infrastructure Components of the Reference Architecture

The IT industry has been transformed over the past decade over the concept of virtualizing and centralizing servers to enable IT to deliver a more secure, manageable, less costly, and ultimately more mature compute model. While the industry as a whole is not 100% virtualized, the availability of commodity scalable server architectures with increasingly large amounts of CPU power and centralized memory have made the promise of virtualizing those remaining power, business critical workloads much closer to reality. The same principles are being applied to Virtual desktops as more and more customers are hosting their end user desktops in the data center to get a secure, reliable, and efficient ways of managing the end user needs.

The Pure Storage FlashArray

Pure Storage is the leading all-flash enterprise array vendor, committed to enabling companies of all sizes to transform their businesses with flash.

Built on 100% consumer-grade MLC flash, Pure Storage FlashArray delivers all-flash enterprise storage that is 10X faster, more space and power efficient, more reliable, and infinitely simpler, and yet typically costs less than traditional performance disk arrays.

	FA-405	FA-420	FA-450
FRONT VIEW (dual controllers)			
REAR VIEW (dual controllers)			
CAPACITY	• Up to 40+ TBs effective capacity • 2.75-11 TBs raw capacity	• Up to 125+ TBs effective capacity • 11-35 TBs raw capacity	• Up to 250+ TBs effective capacity • 34-70 TBs raw capacity
	Effective capacity assumes HA, RAID, and metadata overhead, GB-to-GiB conversion, and includes benefit of data reduction with always-on inline deduplication, compression & pattern removal. Average data reduction is calculated at 6-to-1. Some customers see data reduction in excess of 20-to-1. Effective capacity has no upper limit and will vary depending on workload.		
PERFORMANCE	• Up to 100,000 **32K** IOPS @ <1ms average latency • Up to 3 GB/s bandwidth	• Up to 150,000 **32K** IOPS @ <1ms average latency • Up to 5 GB/s bandwidth	• Up to 200,000 **32K** IOPS @ <1ms average latency • Up to 7 GB/s bandwidth
	Why does Pure Storage quote 32K, not 4K IOPS? The industry commonly markets 4K IOPS benchmark to make numbers look high, but real-world environments are dominated by IO sizes of 32K or larger. Pure Storage has optimized the FlashArray for the real-world. FlashArray adapts automatically to 512B-32KB IO for superior performance, scalability, and data reduction.		
HOST CONNECTIVITY	• 8 Gb/s Fibre Channel • 10 Gb/s Ethernet iSCSI • Replication ports	• 8 Gb/s Fibre Channel • 10 Gb/s Ethernet iSCSI • Expansion slot (FC or iSCSI) • Replication ports	• 16 Gb/s Fibre Channel • 10 Gb/s Ethernet iSCSI • Expansion slot (FC or iSCSI) • Replication ports

Figure 1. Pure Storage FlashArray 400 Series Specifications

Pure Storage FlashArray sets the benchmark for all-flash enterprise storage arrays. It delivers:

Consistent Performance FlashArray delivers consistent <1ms average latency. Performance is optimized for the real-world applications workloads that are dominated by I/O sizes of 32K or larger vs. 4K/8K hero performance benchmarks. Full performance is maintained even under failures/updates.

Less Cost than Disk Inline de-duplication and compression deliver 5 – 10x space savings across a broad set of I/O workloads including Databases, Virtual Machines and Virtual Desktop Infrastructure.

Mission-Critical Resiliency FlashArray delivers >99.999% proven availability, as measured across the Pure Storage installed base and does so with non-disruptive everything without performance impact.

Disaster Recovery Built-In FlashArray offers native, fully-integrated, data reduction-optimized backup and disaster recovery at no additional cost. Setup disaster recovery with policy-based automation within minutes. And, recover instantly from local, space-efficient snapshots or remote replicas.

Simplicity Built-In FlashArray offers game-changing management simplicity that makes storage installation, configuration, provisioning and migration a snap. No more managing performance, RAID, tiers or caching. Achieve optimal application performance without any tuning at any layer. Manage the FlashArray the way you like it: Web-based GUI, CLI, VMware vCenter, Rest API, or OpenStack.

Pure Storage FlashArray FA-400 Series includes FA-405, FA-420, and FA-450. A FlashArray is available for any application, and any budget.

Purity Operating Environment 4.0

Purity is a storage operating environment built from the ground-up for flash. Purity is provided at no additional cost with every FlashArray, and runs consistently across the entire FlashArray hardware family.

The heart of Purity is the Purity Core, a foundation layer which virtualizes the entire array into a large pool of flash. It manages, protects, and extends the life of the flash via FlashCare technology, ensures consistent performance, and provides advanced data services via the Adaptive Metadata Fabric.

On top of the core run three key Purity Services: FlashReduce, FlashProtect and FlashRecover, ensuring array-wide efficiency and data protection.

Finally, Purity is topped-off with technologies to deliver management simplicity and enterprise supportability.

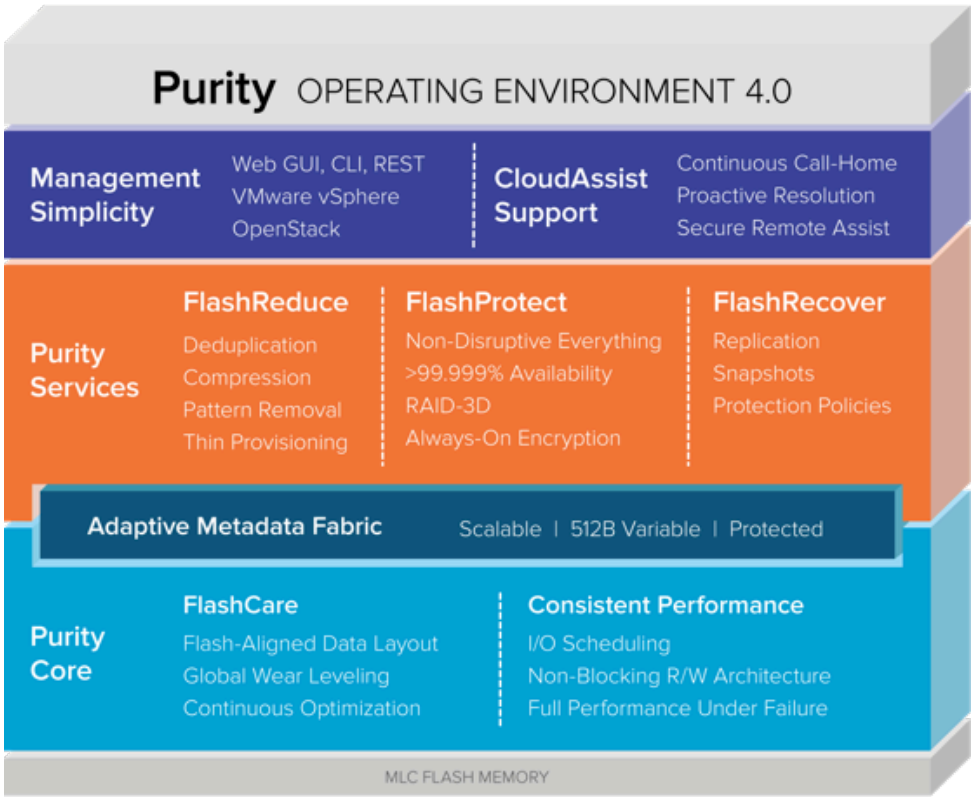

Figure 2. Building block of the Purity Operating Environment 4.0

Cisco Unified Computing System

The Cisco Unified Computing System™ (Cisco UCS™) is a next-generation data center platform that unites compute, network, storage access, and virtualization into an organized structure aimed to reduce total cost of ownership and introduce vastly improved infrastructure deployment mechanisms at scale. UCS incorporates a unified network fabric with scalable, modular and powerful x86-architecture servers. With an innovative and proven design, Cisco UCS delivers an architecture that increases cost efficiency, agility, and flexibility beyond what traditional blade and rack-mount servers provide. Cisco makes organizations more effective by addressing the real problems that IT managers and executives face and solves them on a systemic level.

Figure 3. Cisco Unified Computing System

Greater Time-on-Task Efficiency

Automated configuration can change an IT organization's approach from reactive to proactive. The result is more time for innovation, less time spent on maintenance, and faster response times. These efficiencies allow IT staff more time to address strategic business initiatives. They also enable better quality of life for IT staff, which means higher morale and better staff retention—both critical elements for long-term efficiency.

Cisco UCS Manager is an embedded, model-based management system that allows IT administrators to set a vast range of server configuration policies, from firmware and BIOS settings to network and storage connectivity. Individual servers can be deployed in less time and with fewer steps than in traditional environments. Automation frees staff from tedious, repetitive, time-consuming chores that are often the source of errors that cause downtime, making the entire data center more cost effective.

Easier Scaling

Automation means rapid deployment, reduced opportunity cost, and better capital resource utilization. With Cisco UCS, rack-mount and blade servers can move from the loading dock and into production in a "plug-and-play" operation. Automatically configure blade servers using predefined policies simply by inserting the devices into an open blade chassis slot. Integrate rack-mount servers by connecting them to top-of-rack Cisco Nexus® fabric extenders. Since policies make configuration automated and repeatable, configuring 100 new servers is as straightforward as configuring one server, delivering agile, cost-effective scaling.

Virtual Blade Chassis

With a separate network and separate management for each chassis, traditional blade systems are functionally an accidental architecture based on an approach that compresses all the components of a rack into each and every chassis. Such traditional blade systems are managed with multiple management tools that are combined to give the illusion of convergence for what is ultimately a more labor-intensive, error-prone and costly delivery methodology. Rack-mount servers are not integrated and must be managed separately or through additional tool sets, adding complexity, overhead, and the burden of more time.

Architecturally, Cisco UCS blade and rack-mount servers are joined into a single virtual blade chassis that is centrally managed yet physically distributed across multiple blade chassis, rack-mount servers, and even racks and rows. This capability is delivered through

Cisco® fabric interconnects that provide redundant connectivity, a common management and networking interface, and enhanced flexibility. This larger virtual chassis, with a single redundant point of management, results in lower infrastructure cost per server, with fewer management touch points, and lower administration, capital, and operational costs.

Cisco Nexus 5500UP Switches

Cisco Nexus 5500UP Switches, using cut-through architecture, supports line-rate 10 Gigabit Ethernet on all ports while maintaining consistently low latency independent of packet size and services enabled. It supports a set of network technologies known collectively as Data Center Bridging (DCB) that increases the reliability, efficiency, and scalability of Ethernet networks. These features allow the switches to support multiple traffic classes over a lossless Ethernet fabric, thus enabling consolidation of LAN, SAN, and cluster environments. Its ability to connect Fibre Channel over Ethernet (FCoE) to native Fibre Channel protects existing storage system investments while dramatically simplifying in-rack cabling. The Nexus 5500UP series switch we used was a Nexus 5548UP (see Figure 4 below), which is a 1RU 10 Gigabit Ethernet, Fibre Channel, and FCoE switch offering up to 960 Gbps of throughput and up to 48 ports. The switch has 32 unified ports and one expansion slot.

Figure 4. Cisco Nexus 5548UP Switch

VMware vSphere 5.5

VMware vSphere is the industry-leading virtualization platform for building cloud infrastructures. It enables IT to meet SLAs (service-level agreements) for the most demanding business critical applications, at the lowest TCO (total cost of ownership). vSphere accelerates the shift to cloud computing for existing data centers and also underpins compatible public cloud offerings, forming the foundation for the industry's only hybrid cloud model. With the support of more than 3,000 applications from more than 2,000 ISV partners, vSphere is the trusted platform for any application.

- VMware vSphere Hypervisor Architecture provides a robust, production-proven, high-performance virtualization layer. It enables multiple virtual machines to share hardware resources with performance that can match (and in some cases exceed) native throughput.
- VMware vSphere Virtual Symmetric Multiprocessing enables the use of ultra-powerful virtual machines that possess up to 64 virtual CPUs.
- VMware vSphere Virtual Machine File System (VMFS) allows virtual machines to access shared storage devices (Fibre Channel, iSCSI, etc.) and is a key enabling technology for other vSphere components such as VMware vSphere Storage vMotion®.
- VMware vSphere Storage APIs provide integration with supported third-party data protection, multipathing and storage array solutions.

VMware Horizon View 6.0

Horizon is a family of desktop and application virtualization solutions, which provide a streamlined approach to deliver, protect, and manage Windows desktops and applications to the end user so they can work anytime, anywhere, on any device.

Key Features

Horizon 6 leverages desktop virtualization with View and builds on these capabilities, allowing IT to deliver virtualized and remoted desktop and applications through a single platform and supports users with access to all their Windows and online resources through one unified workspace.

Horizon 6 supports the following key functionalities:

• Desktops and Applications Delivered through a Single Platform – Deliver virtual or remote desktops and applications through a single platform to streamline management and easily entitle end users.

• Unified Workspace – Securely delivers desktops, applications, and online services to end users through a unified workspace, providing a consistent user experience across devices, locations, media, and connections.

• Closed Loop Management and Automation – Consolidated control, delivery and protection of user compute resources with cloud analytics and automation, cloud orchestration and self-service features.

• Optimization with the Software-Defined Data Center – Allocates resources dynamically with virtual storage, compute, and networking to manage and deliver desktop services on demand.

• Central Image Management – Central image management for physical, virtual, and BYO devices.

• Hybrid-cloud flexibility – Provides an architecture built for onsite and cloud-based deployment.

VMware View Architecture and Components

This section describes the components and VMware products that interact with View.

View includes seven main components:

- View Connection Server
- View Security Server
- View Composer Server
- View Agent

- Horizon Clients
- View Persona Management
- ThinApp

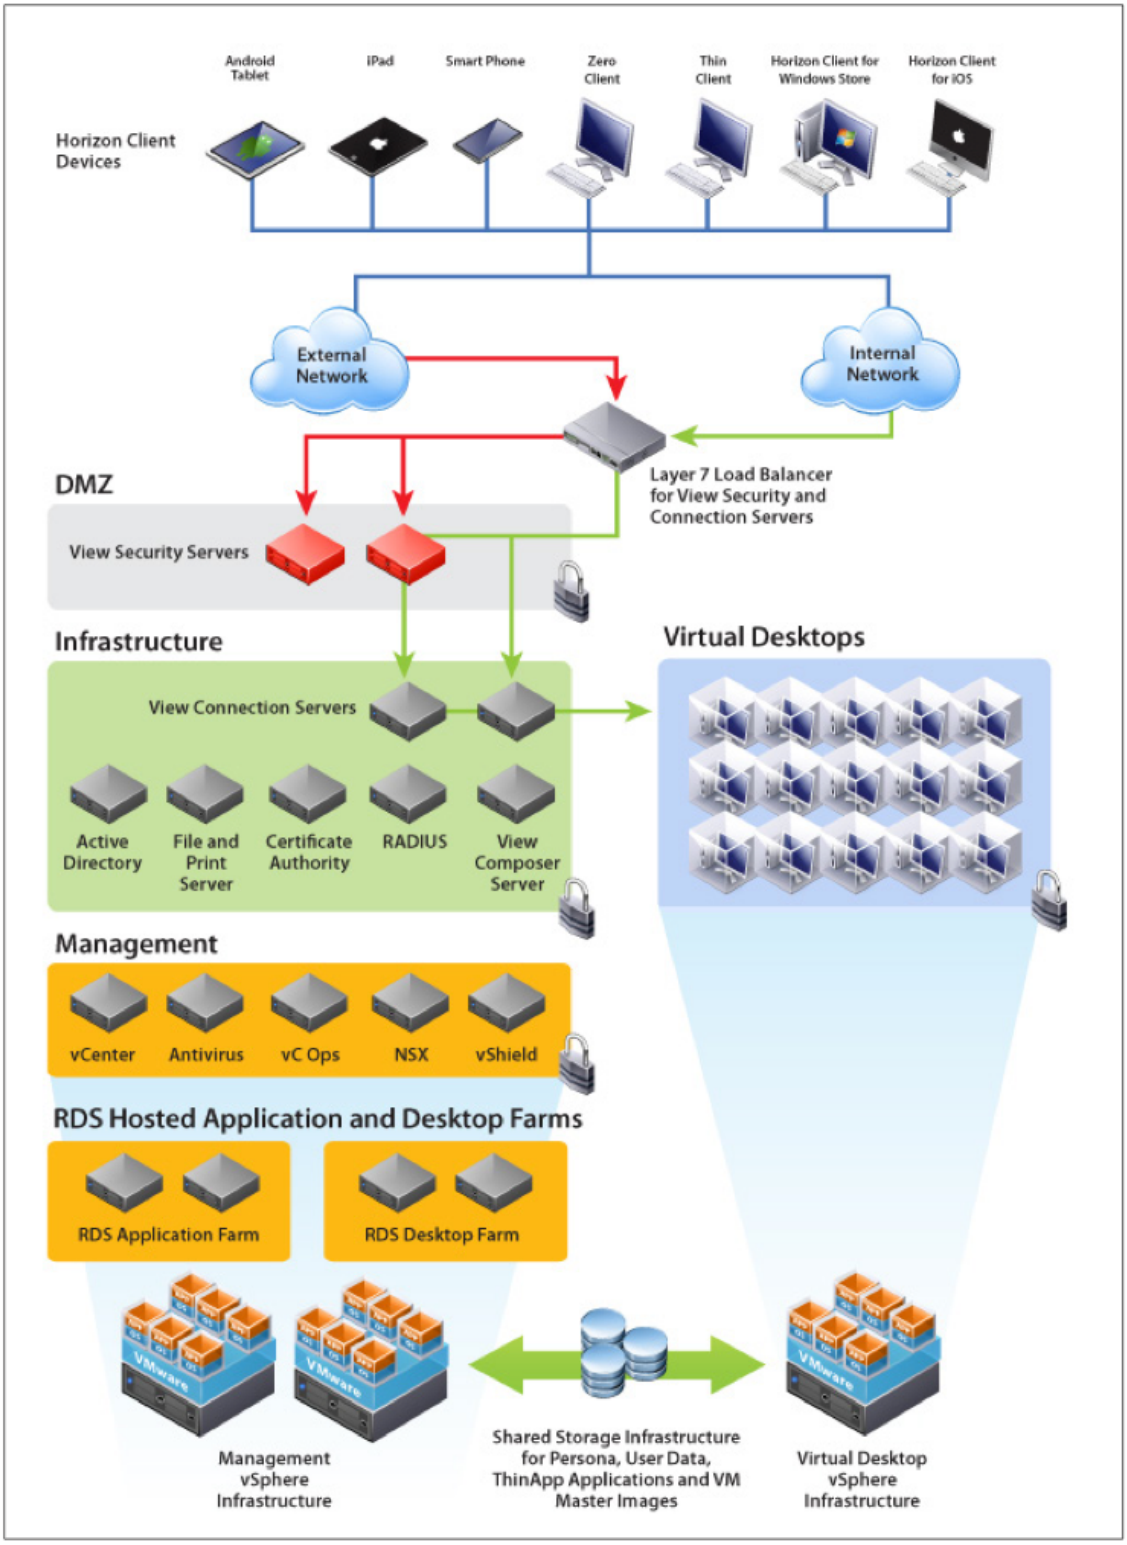

Figure 5. VMware Horizon with View deployment and components

View Connection Server

View Connection Server streamlines the management, provisioning, and deployment of virtual desktops. As an administrator, you can centrally manage thousands of virtual desktops from a single console. End users connect through View Connection Server to securely and easily access their personalized virtual desktops. View Connection Server acts as a broker for client connections by authenticating and directing incoming user desktop requests.

View Security Server

A View security server is an instance of View Connection Server that adds an additional layer of security between the Internet and your internal network. Outside the corporate firewall, in the DMZ, you can install and configure View Connection Server as a View security server. Security servers in the DMZ communicate with View Connection Servers inside the corporate firewall. Security servers ensure that the only remote desktop traffic that can enter the corporate data center is traffic on behalf of a strongly authenticated user. Users can only access the desktop resources for which they are authorized.

View Composer Server

View Composer Server is an optional service that enables you to manage pools of “like” desktops, called linked- clone desktops, by creating master images that share a common virtual disk. Linked-clone desktop images are one or more copies of a parent virtual machine that share the virtual disks of the parent, but which operate as individual virtual machines. Linked-clone desktop images can optimize your use of storage space and facilitate updates. You can make changes to a single master image through the vSphere Client. These changes trigger View Composer Server to apply the updates to all cloned user desktops that are linked to that master image, without affecting users’ settings or persona data.

View Agent

The View Agent service communicates between virtual machines and Horizon Client. You must install the View Agent service on all virtual machines managed by vCenter Server so that View Connection Server can communicate with them. View Agent also provides features such as connection monitoring, virtual printing, persona management, and access to locally connected USB devices. View Agent is installed in the guest operating system.

Horizon Clients

Horizon Clients are available for Windows, Mac, Ubuntu Linux, iOS, and Android to provide the connection to remote desktops from your device of choice.

By installing Horizon Client on each endpoint device, your end users can access their virtual desktops from devices such as smartphones, zero clients, thin clients, Windows PCs, Macs, and iOS- and Android-based mobile devices. Unity Touch for Horizon Clients

makes it easier to run Windows apps on iPhone, iPad, and Android devices. Horizon Clients enable users to:

• Connect to View Connection Server or a View security server

• Log in to their remote desktops in the data center

• Edit the list of servers they connect to

View Persona Management

View Persona Management is an optional component included with Horizon with View that provides persistent, dynamic user profiles across user sessions on different desktops. You can deploy pools of stateless, floating desktops and enable users to maintain their designated settings between sessions. User profile data is downloaded as needed to speed up login and logout time. New user settings are automatically sent to the user profile repository during desktop use.

ThinApp

ThinApp is an optional software component included with Horizon that creates virtualized applications. In a Horizon implementation, these virtual packages reside on a ThinApp repository in a network share. As an administrator, you can copy a full ThinApp package from the repository to the virtual desktop. You can also place a shortcut on the virtual desktop that points to the ThinApp package on the repository. Applications on remote desktops can be natively installed applications, ThinApp virtual applications, or shortcuts to ThinApp virtual applications. You can permanently copy a ThinApp virtual application to a remote desktop or add a shortcut that points to the virtual application on the ThinApp repository. As part of Horizon, ThinApp simplifies repetitive administrative tasks and reduces storage needs for virtual desktops by maintaining applications independent of the underlying OS.

VMware View Planner 3.0.1

VMware View Planner is a tool designed to simulate a large-scale deployment of virtualized desktop systems and study its effects on an entire virtualized infrastructure. The tool is scalable from a few virtual machines running on one VMware ESX host up to hundreds of virtual machines distributed across a cluster of ESX hosts.

View Planner assists in the setup and configuration of the testing infrastructure, runs a set of application operations selected to be representative of real-world user applications, and reports data on the latencies of those operations.

A selection of applications to run in the workload:

- Microsoft Word
- Microsoft Excel
- Microsoft PowerPoint

- Microsoft Outlook
- Microsoft Internet Explorer
    - Document browse
    - Picture album browse
- Mozilla Firefox
- Adobe Reader
- Archiving software
- Video playback software
- Custom applications

VMware View Planner consists of the following components:

- Number of desktop virtual machines running on one or more ESX hosts.
- Number of client virtual machines running on one or more ESX hosts (only used in the case of remote-mode and passive-mode runs; not used for local-mode runs).
- Single controller appliance running on an ESX host. Figure 6 below shows a conceptual overview of a typical View Planner layout.

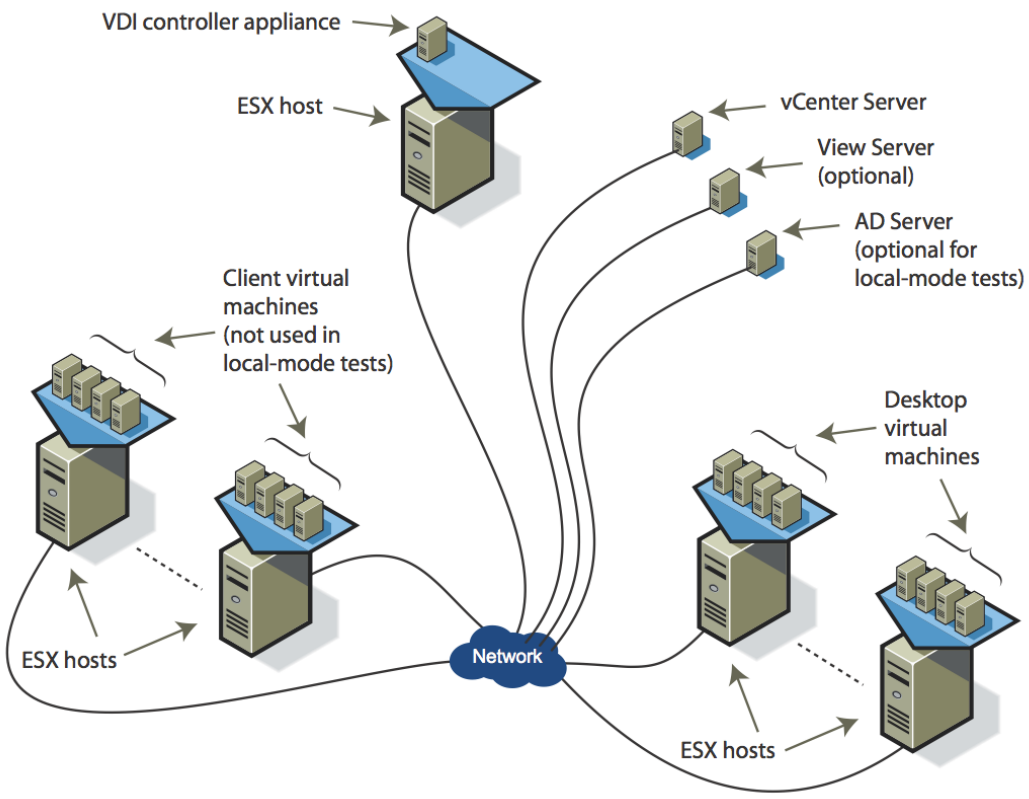

Figure 6. Conceptual diagram of VMware View planner layout

View Planner Workload Applications and User Operations

The standardized View Planner workload mix consists of nine applications running in the desktop virtual machines and performing a combined total of 44 user operations. These user operations are separated into the three groups shown in Table 1: interactive operations (Group A), I/O operations (Group B), and background load operations (Group C). The operations in Groups A and B are used to determine Quality of Service, while the operations in Group C are used to generate additional load.

When View Planner is run as a benchmark, the standardized workload mix must be used.

When View Planner is run in flexible mode (that is, not as a benchmark), user operations can be removed from the standardized mix and custom user operations and even custom applications can be added.

Table 1. View Planner 3.0.1 User Operations

Group A (Interactive Operations)	Group B (I/O Operations)	Group C (Background Load)
AdobeReader: Browse	AdobeReader: Open	7-Zip: Compress
AdobeReader: Close	Excel_Sort: Open	Outlook: Restore
AdobeReader: Maximize	Excel_Sort: Save	PowerPoint: SaveAs
AdobeReader: Minimize	Firefox: Open	Video: Play
Excel_Sort: Close	IE_ApacheDoc: Open	
Excel_Sort: Compute	IE_WebAlbum: Open	
Excel_Sort: Entry	Outlook: Attachment-Save	
Excel_Sort: Maximize	Outlook: Open	
Excel_Sort: Minimize	PowerPoint: Open	
Firefox: Close	Video: Open	
IE_ApacheDoc: Browse	Word: Open	
IE_ApacheDoc: Close	Word: Save	
IE_WebAlbum: Browse		
IE_WebAlbum: Close		
Outlook: Close		
Outlook: Read		
PowerPoint: AppendSlides		
PowerPoint: Close		
PowerPoint: Maximize		
PowerPoint: Minimize		
PowerPoint: ModifySlides		
PowerPoint: RunSlideShow		
Video: Close		
Word: Close		
Word: Maximize		
Word: Minimize		
Word: Modify		

Whether used as a benchmark or in flexible mode, an iteration is the completion of all the user operations in the currently selected set.

Reference Architecture Solutions Overview

FlashStack consists of a combined stack of hardware (storage, network and compute) and software (Cisco UCS Manager, VMware vCenter/ESXi and Pure Storage GUI.

- **Network**: Cisco Nexus 5500UP and Cisco UCS Fabric Interconnect 6248UP for external and internal connectivity of IP and FC network.
- **Storage**: Pure Storage FlashArray FA-405 with Fibre Channel connectivity
- **Compute**: Cisco UCS B200 M3 Blade Server

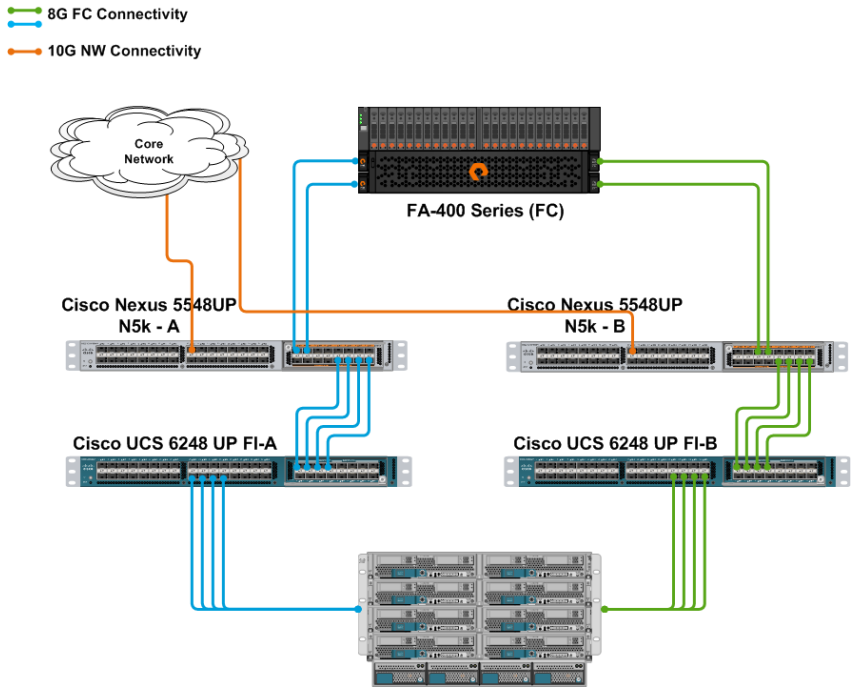

Figure 7. FlashStack Connectivity Diagram

Figure 7 shows a detailed topology of the reference architecture configuration. A major goal of the architecture is to build out a highly redundant and resilient infrastructure. Thus, we used powerful servers with dual Fibre Channel ports connected redundantly to two SAN switches that were connected to redundant FC target ports on the FlashArray. The servers were hosted in a vSphere HA cluster and had redundant network connectivity.

Cisco UCS Server Configuration

A pair of Cisco UCS Fabric Interconnects 6248UP, and one chassis with eight identical Intel CPU-based Cisco UCS B-series B200-M3 blade servers were deployed for hosting the virtual desktops. The UCS manager, UCS Fabric Interconnects and the components in the chassis were upgraded to 2.2.2c firmware level.

The server had the same Cisco VIC 1240 cards and they were connected four ports from each Cisco Fabric extender of the Cisco UCS chassis to the Cisco Fabric Interconnect, they were in turn connected to Cisco Nexus 5548UP Switch for upstream connectivity to access the Pure Storage FlashArray LUNs. The server configuration is described in the Table 2 below.

Component	Description
Processor	2 X Intel Xeon E5-2680 @ 2.7GHz (16 Cores total, 32 Logical CPUs total)
Memory	256 GB @ 1600 MHz (16 X 16GB), regular voltage
HBA	2 X 8G ports on Cisco UCS VIC 1240 (UCSB-MLOM-40G-01)
NIC	4 X 10G ports on Cisco UCS VIC 1240 (UCSB-MLOM-40G-01)
BIOS	Intel Virtualization Tech, Intel ATS Support, Intel VT-D features were enabled, Turbo Boost
vSphere	ESXi 5.5U1, Build 1746018 (Cisco Customized image)

Table 2: Desktop host server configuration

Cisco UCS Service Profile configuration

In order to facilitate rapid deployment of UCS servers, a service profile template was created with the following characteristics [more details can be found in the deployment guide].

1. We configured boot from SAN policy so that the server booted from a Pure Storage boot LUN (see Figure 8 below)

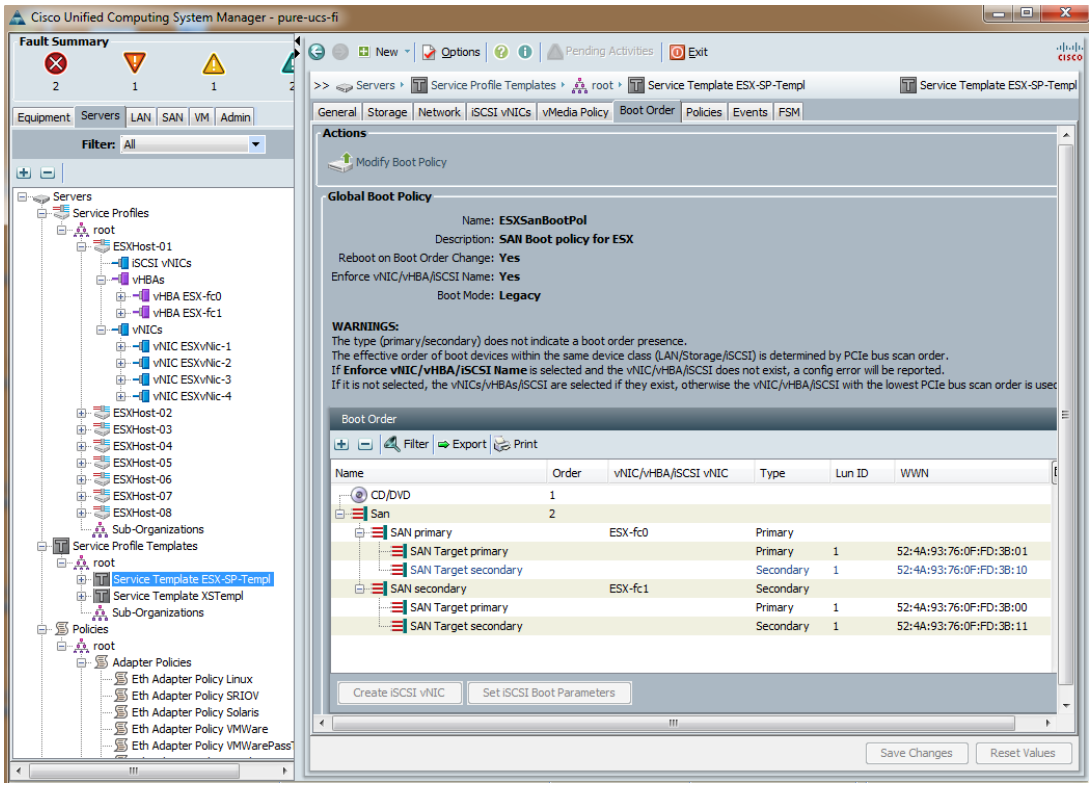

Figure 8. Cisco UCS service profile template with Boot from SAN policy configuration

2. We kept every other setting to the default, we didn't tweak any parameters
3. The Ethernet and FC adapter policy was set to VMware policy
4. The BIOS defaults were used for the B200-M3 blade servers
5. We configured two vHBA FC adapter and four vNIC Eth adapters on the Cisco VIC cards
6. We deployed eight service profiles from the template and associated it with the blade servers in the chassis. Figure 9 below shows the Cisco UCS manager snapshot of service profile setup for the tests.

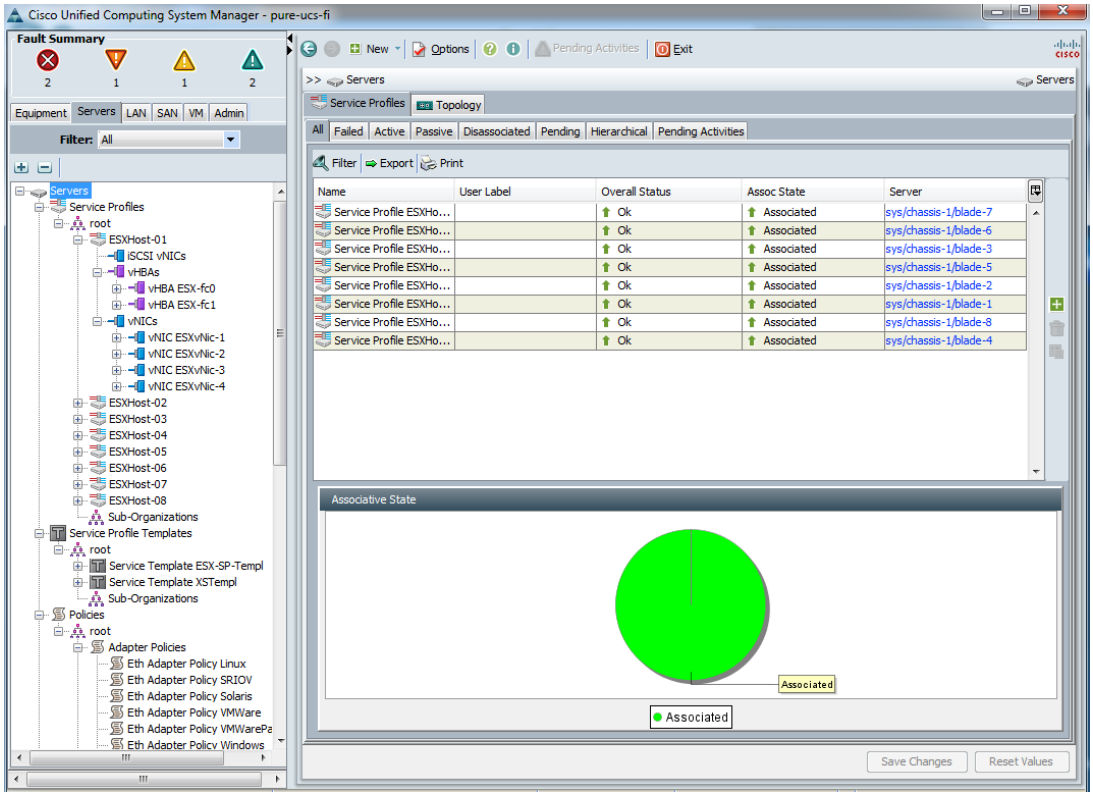

Figure 9. Cisco UCS service profile configuration

VMware vSphere Configuration and Tuning

In this section, we discuss the ESXi 5.5U1 cluster configuration, network configuration and ESXi tuning for the system configuration.

ESXi Cluster and storage configuration

A datacenter and a cluster with eight hosts were configured with VMware High Availability (HA) clustering and Distributed Resource Scheduling (DRS) features. DRS was set to partially automated mode with power management turned off. The host EVC policy was set to Intel Westmere. The default BIOS for B200-M3 was chosen for all the service profiles. We had to create two datastore for the ESXi cluster for making the HA cluster datastore heartbeat to work correctly.

Due to the simplicity of both the Pure Storage FlashArray and the Cisco UCS, configuration of VMware ESXi best practice configuration is accordingly simple. ESXi uses its Native Multipathing Plugin architecture to manage I/O multipathing to underlying SAN storage volumes. Pure Storage FlashArray volumes (while not actually an ALUA array—it indeed is active/active) volumes are claimed by default by the Storage Array Type Plugin (SATP) for ALUA devices. Therefore all devices (by default) would inherit the Most Recently Used (MRU) Path Selection Policy (PSP). This would limit I/O to a single path and would be a colossal detriment to performance, as only leveraging a single path/port to the array would remove the active/active nature and performance advantage of the FlashArray.

All the ESXi servers were configured to change the default PSP for Pure devices from MRU to Round Robin (with advanced configuration to alternate paths after every I/O). The following command was run on each ESXi server prior to the presentation of FlashArray devices:

```
esxcli storage nmp satp rule add -s "VMW_SATP_ALUA" -V "PURE" -M "FlashArray" -
P "VMW_PSP_RR" -O "iops=1"
```

Figure 10 shows a properly configured Pure Storage LUN.

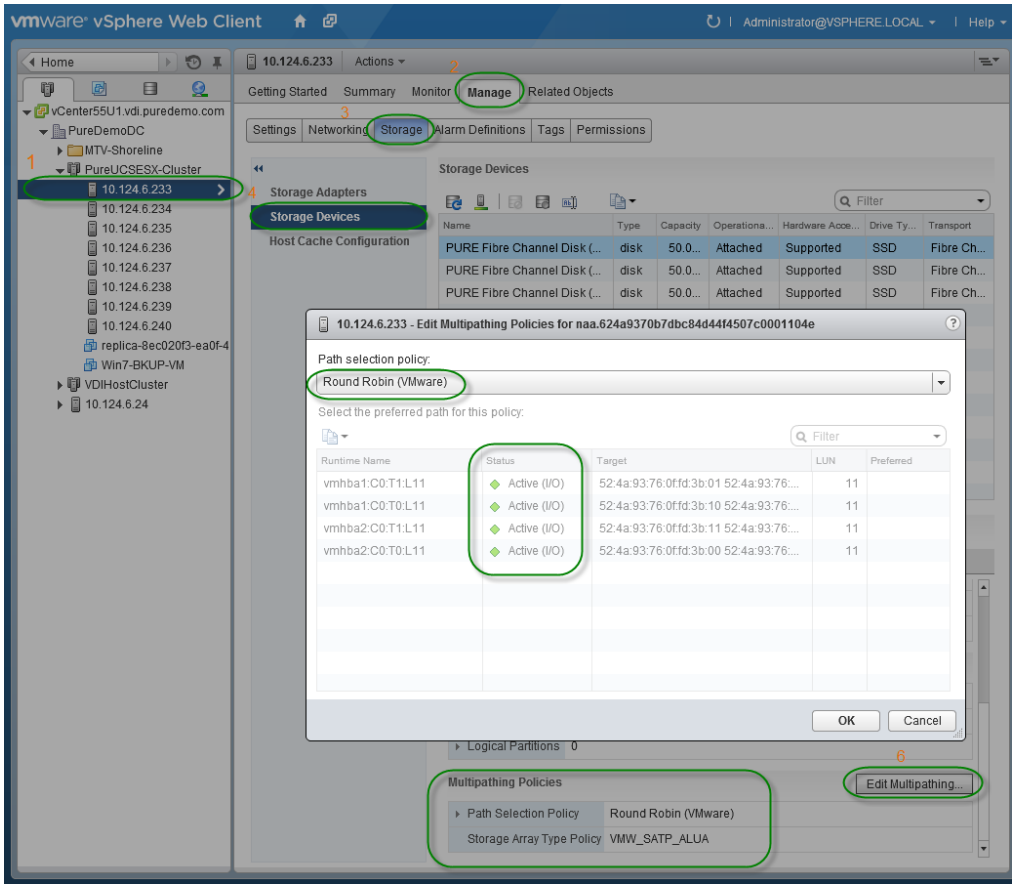

Figure 10. VMware vSphere web client showing Pure Storage LUN multipathing configuration

ESXi Network configuration

Two virtual switches each containing two vmnics were used for each host. Although this design could have taken advantage of distributed vSwitch (DVS), we went with standard vSwitch. The redundant NICs were teamed in active/standby mode and VLAN configurations were done on the upstream Cisco Nexus 5500 UP switches. The Nexus switch provided an internal private network and had a DNS helper which redirected to the infrastructure DNS and DHCP. The virtual switch configuration and properties are shown in figure 11 and figure 12.

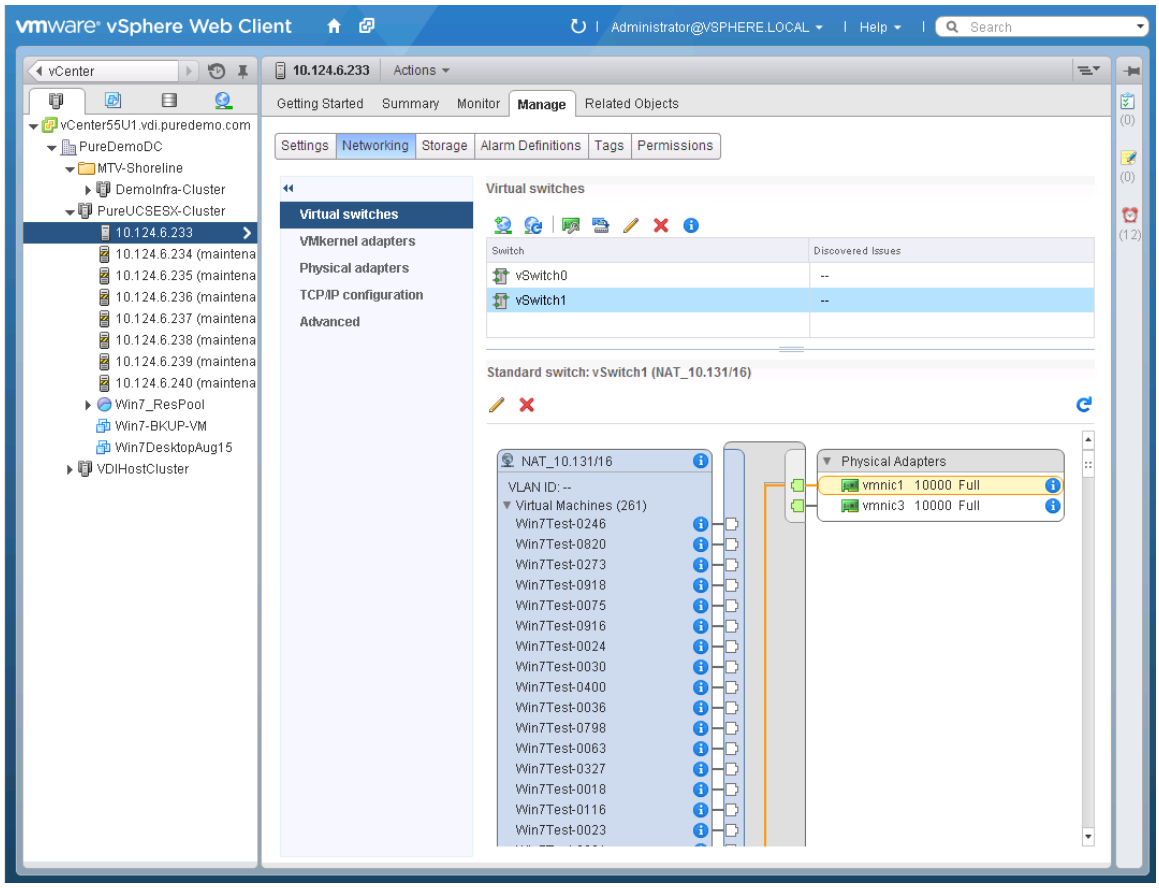

Figure 11. ESXi server network configuration on all servers (vSwitch1 for Desktops)

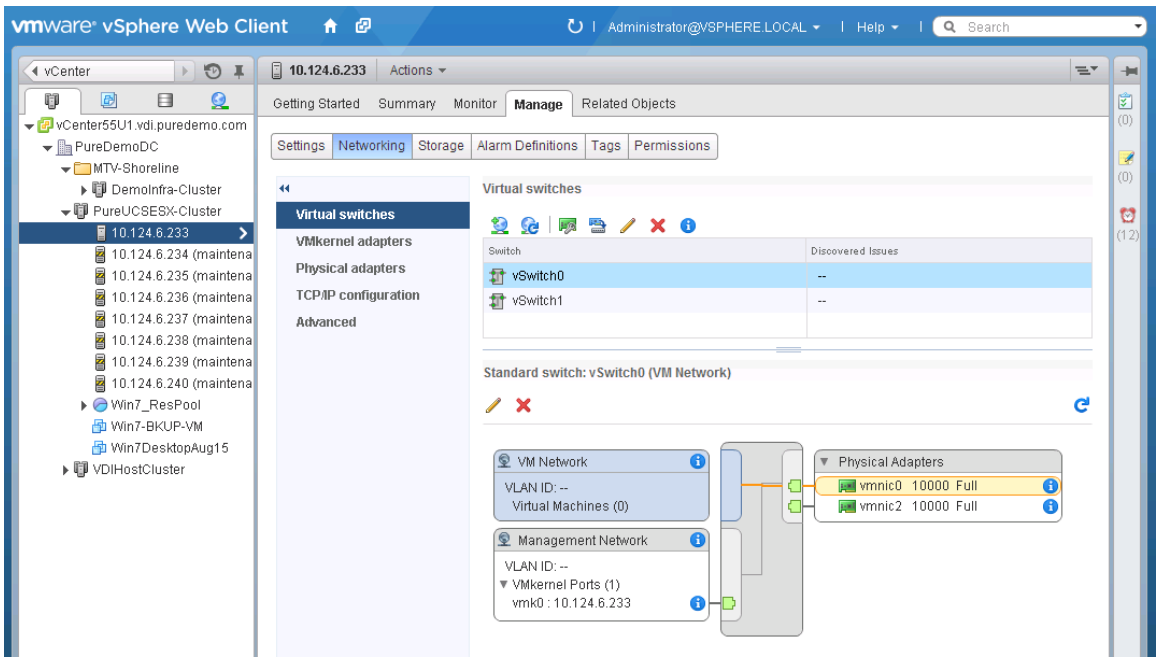

Figure 12. ESXi server network configuration on all servers (vSwitch0 for management)

VMware Horizon View 6.0 Configuration

VMware Horizon View 6.0 configurations were quite minimal; some of the tuning is highlighted in the section.

1. Use SE sparse Virtual disks format

   VMware Horizon View 5.2 and above supports a new vmdk disk format called Space Efficient (SE) sparse virtual disks which was introduced in vSphere 5.1. The advantages of SE sparse virtual disks can be summarized as follows:

   - Benefits of growing and shrinking dynamically, this prevent vmdk bloat as desktops rewrite data and delete data.
   - Available for View Composer based linked clone desktops (Not for persistent desktops)
   - VM hardware version 9 or later
   - No need to do refresh/recompose to reclaim space (Eliminated)
   - No need to set black out periods, as we handle UNMAPs efficiently

We recommend using this new disk format for deploying desktops on Pure Storage due to the space efficiencies and preventing vmdk bloats. Appendix A has screen shots of how to configure VM disk space reclaim in VMware Horizon View 6.0.

2. Disable View Storage Accelerator

   The View storage accelerator, VSA, is a feature in VMware View 5.1 onwards based on VMware vSphere content based read caching (CBRC). There are several advantages of enabling VSA including containing boot storms by utilizing the host side caching of commonly used blocks. It even helps in steady state performance of desktops that use same applications. As Pure Storage FlashArray gives you lots of IOPS at very low latency, we don't need the extra layer of caching at the host. The biggest disadvantage is the time it takes to recompose and refresh desktops, as every time you change the image file it has to rebuild the disk digest file. Also it consumes host side memory for caching and consume host CPU for building digest files. For shorter desktop recompose times, we recommend turning off VSA. Appendix A has screen shot which has VSA disabled in the vCenter Server storage settings.

3. Tune maximum concurrent vCenter operations

   The default concurrent vCenter operations on the vCenter servers are defined in the View configuration's advanced vCenter settings. These values are quite conservative and can be increased to higher values. Pure Storage FlashArray can withstand more operations including –

   Max Concurrent vCenter provisioning operation (recommended value >= 50)

Max Concurrent Power operations (recommended value >= 50)

Max concurrent View composer operations (recommended value >= 50)

The higher values will drastically cut down the amount of time spent for doing these operations.

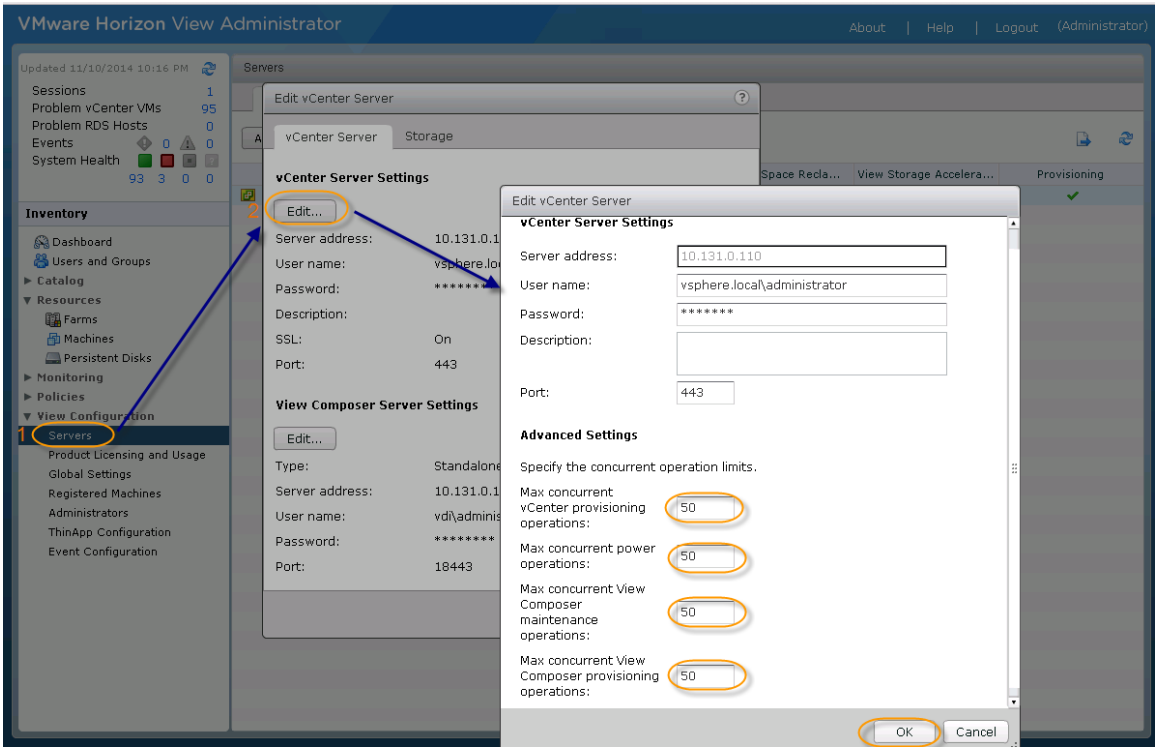

Figure 13. Advanced vCenter settings in Horizon View

Some caveats include –

1. These settings are global settings and will affect all the pools. Pools with other disk arrays will suffer if you set these values higher, so enabling these will have adverse effects.

2. vCenter configuration especially number of vCPUs, amount of memory, and the backing storage has implications from these settings. In order to attain the performance levels we have described in this white paper, it is important to note the vCenter configurations.

VMware Horizon View Linked Clone based Windows 7 desktop pool

VMware Horizon View 6.0 (v6.0.1, build 2088845) was installed on a Windows 2008 R2 VM with 4 vCPU/16GB of memory. The VMware Horizon View Composer 6.0.1 (build 2078421) was installed on a separate Windows 2008 R2 VM with 2 vCPU/4 GB of

memory for linked clone deployment. We used the View connection server to deploy all the Windows 7 desktops for View Planner testing. The automated floating desktop pool settings (for View composer based linked clones) to deploy 1,000 linked clone Windows 7 desktops are shown below.

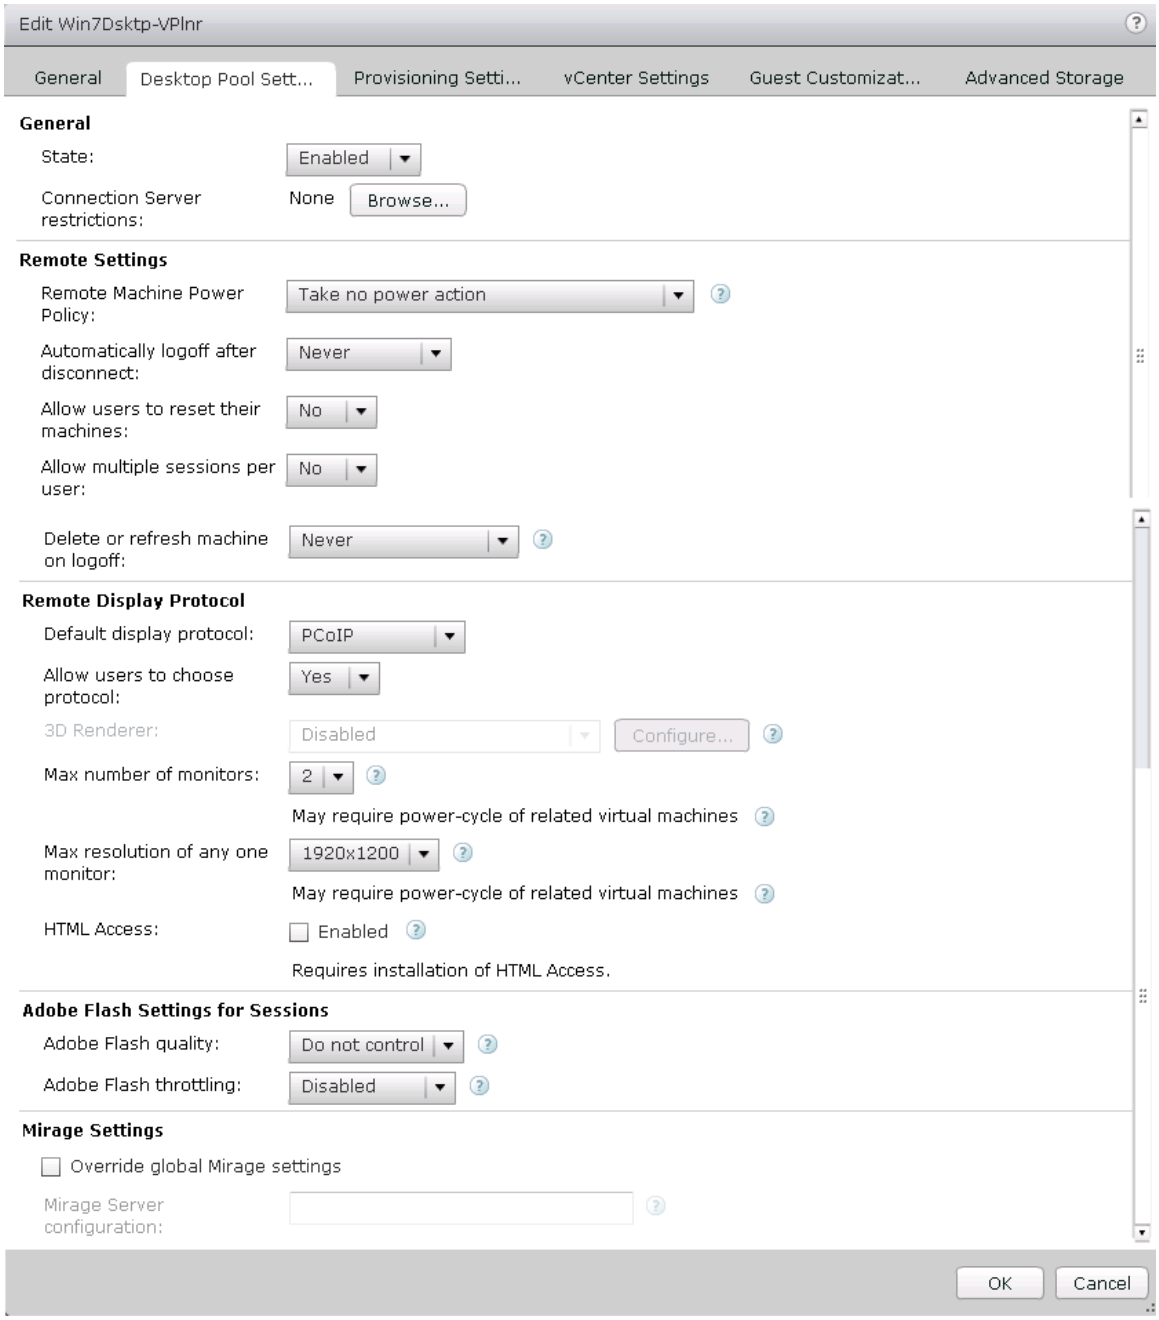

Figure 14. VMware Horizon View Windows 7 desktop pool setting

The entire pool creation of 1000 linked clone Windows 7 desktop took less than 40 minutes. The Pure Storage FlashArray sustained over 40,000 IOPS and the latency was less than 0.5msec with over 1.5GB/sec bandwidth. Figure 15 below shows the pool creation from Pure Storage dashboard perspective.

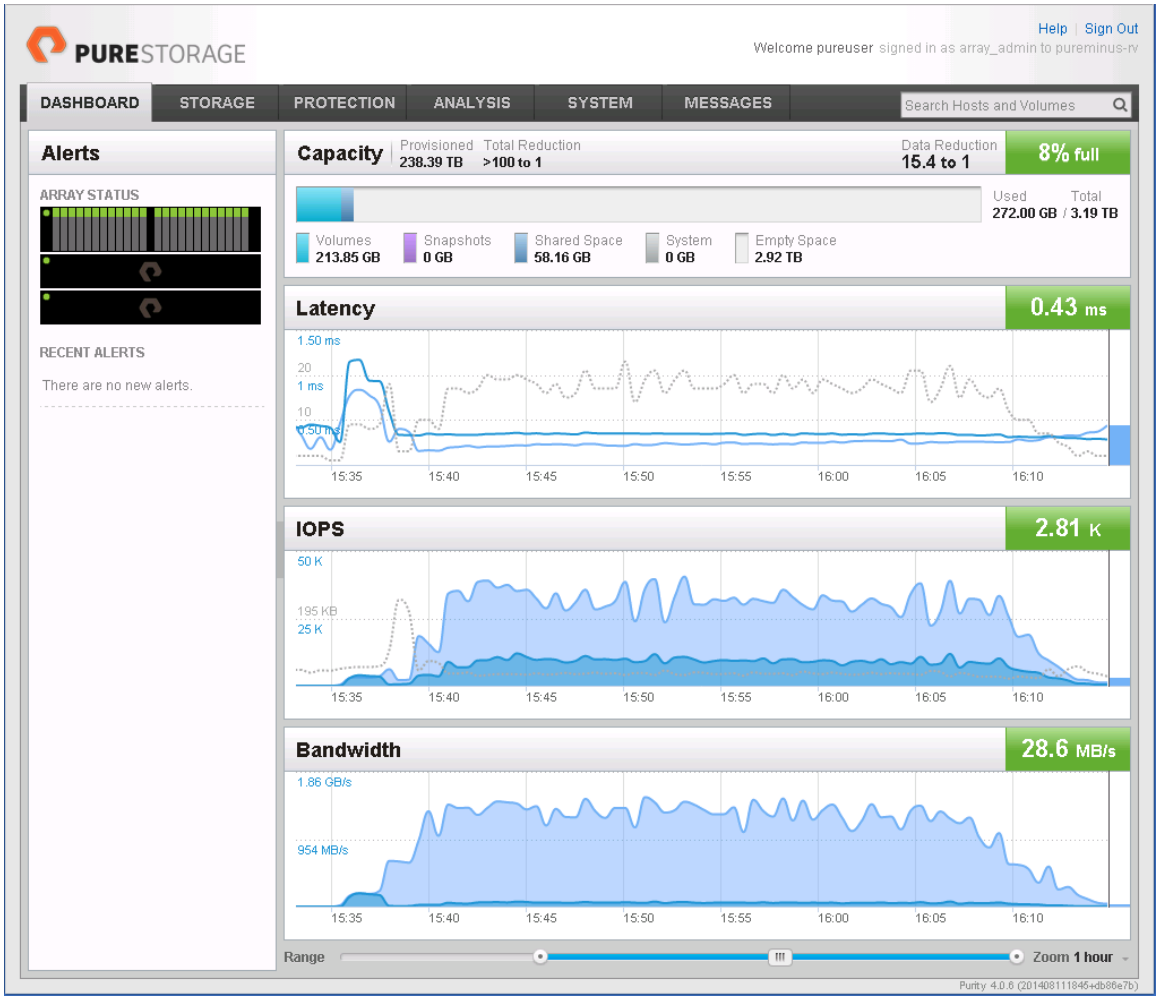

Figure 15. VMware Horizon View Windows 7 desktop pool creation storage utilization

Microsoft Windows 7 Desktop Configuration

The View Planner document provided guidelines for configuring the base Windows 7 image. The Windows 7 Enterprise (64-bit) desktop had the following configuration –

Component	Description
Desktop	Windows 7 Enterprise (64-bit)
Hardware Version	9
vCPU	1
Memory	2 GB
vNIC Adapter	1 X vmxnet3 adapter
vStorage Adapter	1 X ParaVirtual storage adapter
Virtual Disks	30 GB base image + 1 GB Data disk
VMware Horizon View Agent	Viewagent-x86_64-6.0.1-2089044
Installed Applications	MS Office 2007 including Outlook, Adobe Reader 9,Firefox, Internet Explorer, Archive-7Zip, Windows Media player

Table 3: Windows 7 Desktop configuration

Desktop Testing Tool – VMware View Planner 3.0.1

VMware View Planner is a tool designed by VMware to simulate a large-scale deployment of virtualized desktop systems and study its effects on an entire virtualized infrastructure. The tool is scalable from a few virtual machines running on one VMware vSphere host up to thousands of virtual machines distributed across a cluster of vSphere hosts.

View Planner runs a set of application operations selected to be representative of real-

world user applications, and reports data on the latencies of those operations. In our tests, we used this tool to simulate a real world scenario, and then accepted the resultant application latency as a metric to measure end user experience.

We used the StandardBenchmarkProfile mode with a think time of 2 seconds and the default five iterations for our entire test runs. The ramp-up time used was minimum (2*vm+60, 600) seconds.

View Planner has three run modes based on what is getting tested, including passive mode, remote mode and local mode. We did local mode testing with VMware Horizon View based Linked Clone desktops with the following settings:

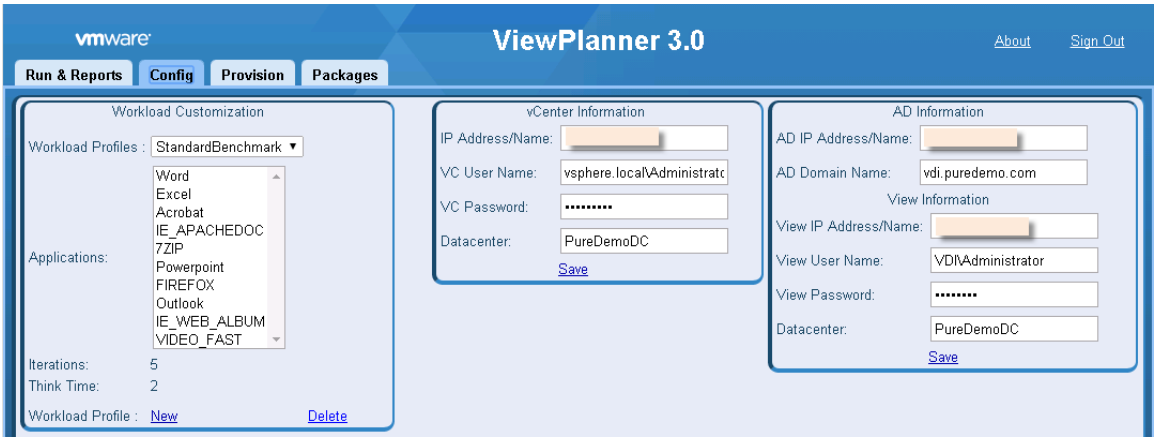

Figure 16. VMware View Planner 3.0.1 configuration details

Pure Storage FlashArray Configuration

The FlashArray FA-405 contains no special configurations or value changes from any normal configuration. The FlashArray contains two drive bays fully populated with 256 GB SSDs with two NVRAM devices per shelf.

The hosts are redundantly connected to the controllers with two FC connections to each controller from two HBAs on each host over the Fibre Channel protocol for a total of eight logical paths.

A cluster group was configured with all the ESXi hosts and a private volume was created for boot from SAN for each host. A single 50 TB LUN was shared across the entire host group for hosting the desktops. A screen shot of the Pure Storage Storage tab is shown in figure 17.

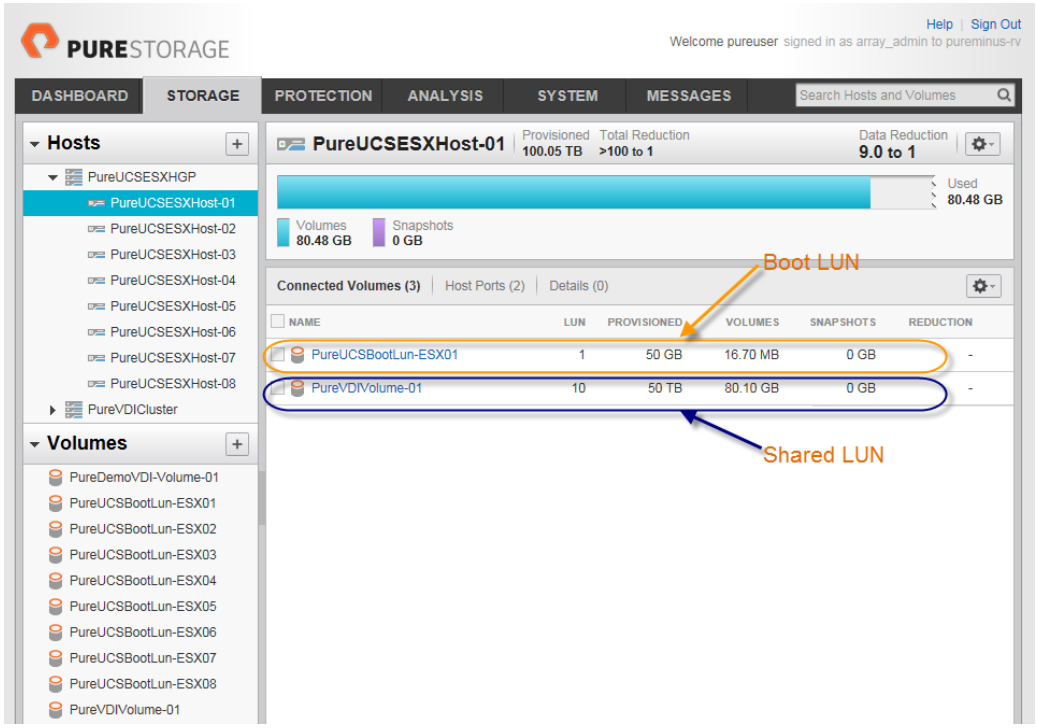

Figure 17. Pure Storage FlashArray Storage configuration

Solution Validation

In order to deploy 100s of desktops a proper hardware and software configuration, a good test plan and success criteria has to be in place. This section talks about the test infrastructure, hardware configuration and infrastructure VM setup we had in place for this reference architecture.

Test Setup

Figure 18 shows a topological view of the test environment for our reference architecture. The VMware Horizon View infrastructure components were placed on a dedicated infrastructure host. The boot LUNs for each ESXi hosts and all the desktops were hosted on a single 5.5 TB (raw) FlashArray FA-405.

The tested configuration included:

- One 5.5 TB Pure Storage FlashArray FA-405 in HA configuration, including two controllers and one disk shelf:
  - 1 x 50 TB volume was provisioned on Pure Storage FlashArray
  - 8 X 50 GB Boot volume for the eight Cisco UCS Blade servers
- Eight Cisco UCS B-series blade server based on dual socket Intel E5-2680 (2.7 GHz) processor with 256 GB of memory running ESXi 5.5U1 (Cisco UCS customized image) were used to host the desktops
- One dedicated infrastructure server was used to host the all of the infrastructure virtual machines:
  - Active directory VM, DNS VM, and DHCP VM
  - VMware Horizon View Connection server
  - VMware Horizon View Composer server
  - VMware vSphere Virtual Center server
  - Microsoft SQL server for both virtual center and View event database
  - VMware View Planner 3.0.1 appliance

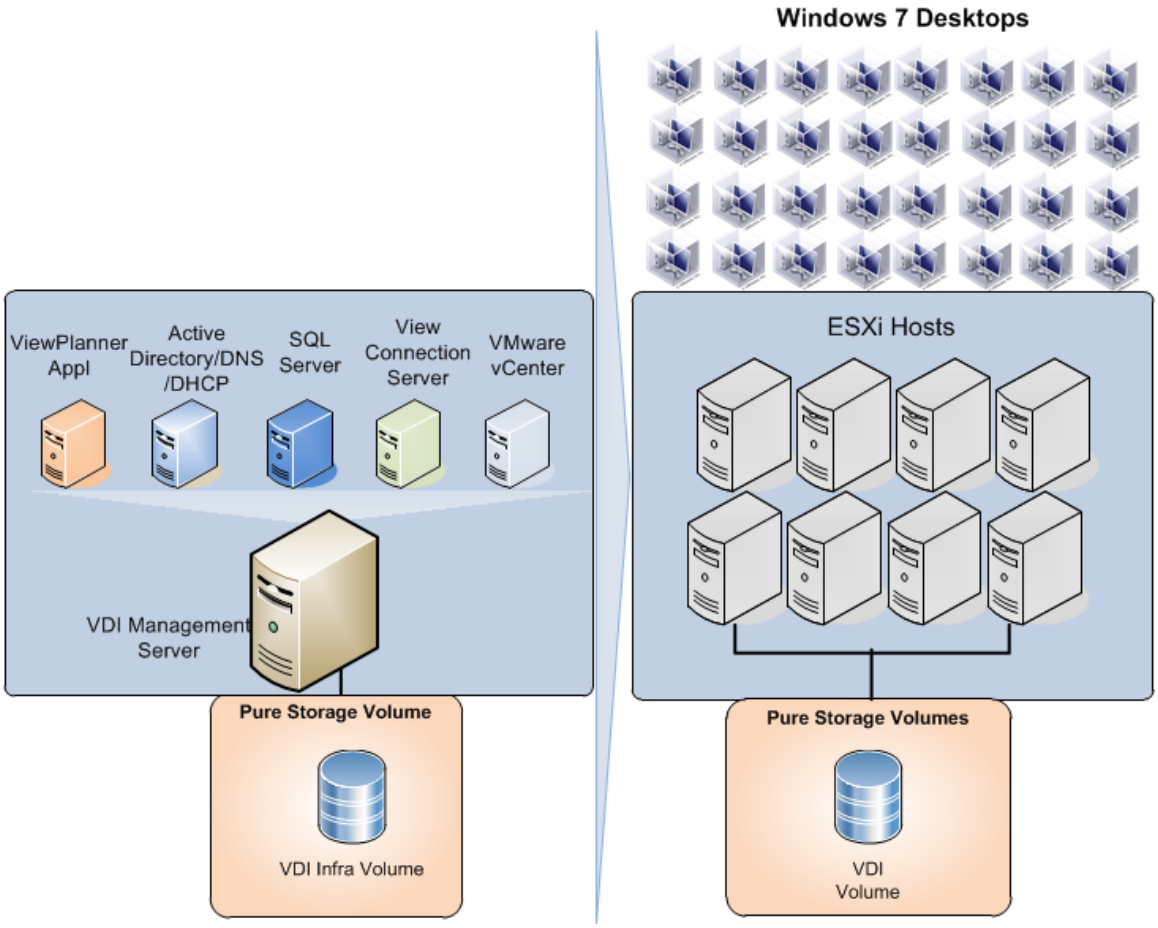

Figure 18. Pure Storage FlashArray Storage configuration

Test Plan and Success Criteria

The test plan was to first determine the number of desktops that can be deployed on a single server and use that as a building block for scaling and deploying desktops on the entire blade server chassis. To achieve that a single server scale testing was done to find out the breaking point of the server starting with 50 desktops on one server and scaling it in increments of 10. We used VMware Horizon View 6.0 to deploy a desktop pool with 1000 desktops from a template image. Then we moved a certain number of desktops over to individual servers and ran the tests.

The success criteria is as follows:

1. The server utilization has to be less than 100% CPU and 100% Memory consumption with no memory ballooning or swapping

2. The VMware View planner score should be consistent and must have passing score (response time from each desktops should be less than 1 second and 6

seconds for Group-A and Group-B score respectively)

3. The backend Pure Storage FlashArray is keeping up with the I/O needs and are getting sub-millisecond latency

Once we determined the single server scalability numbers we then scaled the number of servers from 1 to 2 to 4 and finally 8 servers or a complete chassis. At each point we captured the ESXi server CPU/Memory/Disk utilization along with the Pure Storage FlashArray disk utilization data. The VMware View planner provided a detailed report on each of the test run which would determine the test pass or failure.

Scalability Results

This section highlights the test results for the scalability testing starting with single server scale testing which would be used in the subsequent two, four, eight servers testing.

Single Server testing

Figure 19 below describes the test setup for a single server testing. A single Cisco UCS B200-M3 blade server is deployed and the number of desktops is varied starting from 50 users. Using VMware View planner as the workload generator the test data is verified to see if it meets the success criteria. The number of desktops at the break-even point is noted and it is used as the building block for subsequent desktop deployment.

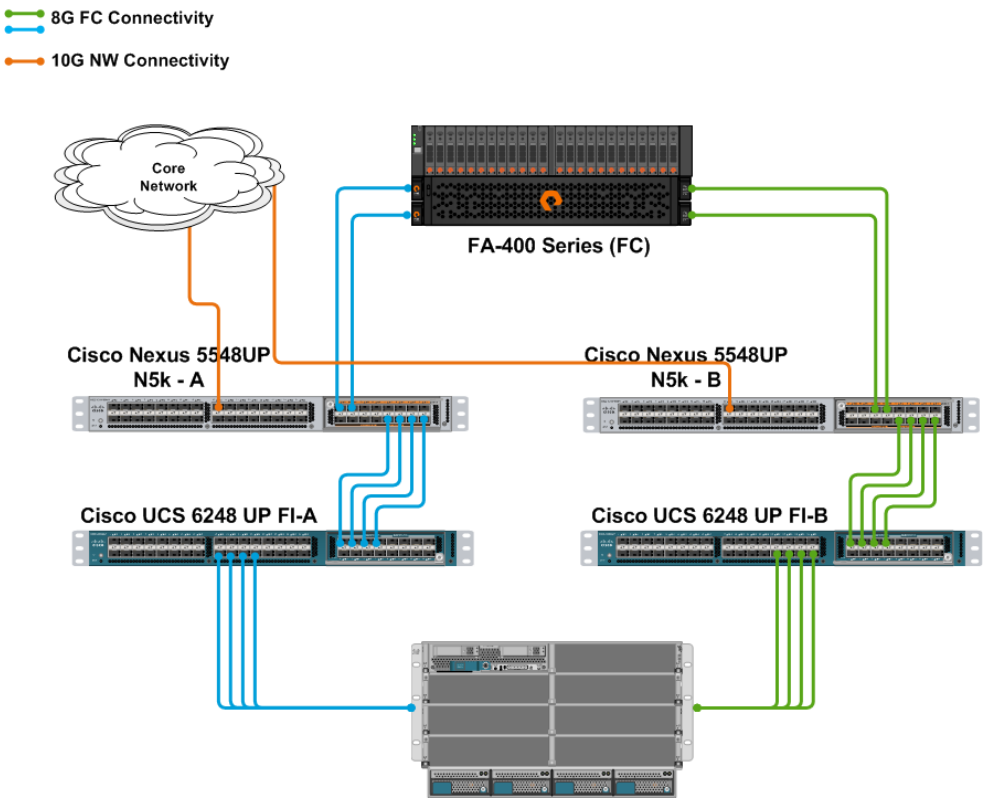

Figure 19. Single server scale test configuration

Test Results

Running the VMware View Planner Benchmark with medium workload on a single server (with all other servers in the Cluster in maintenance mode) we were able to run 100 desktops sessions and the View planner Group-A 95% response time was 0.52 second and Group-B 95% response time was 2.896 seconds (Lower is better). The resource utilization of the ESXi server and storage utilization is shown below.

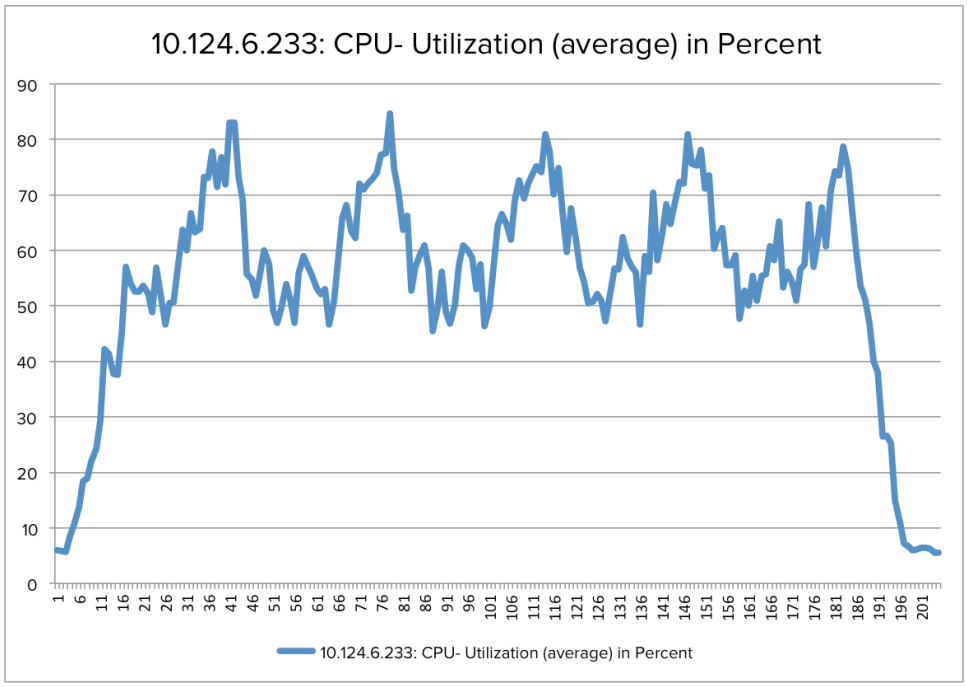

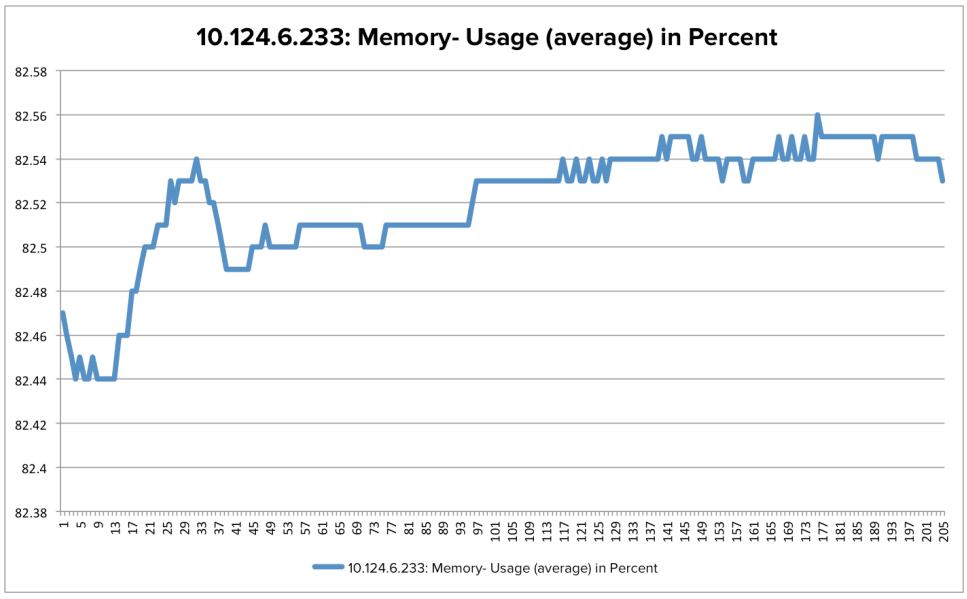

Figure 19.1 Single server scale test CPU/Memory utilization

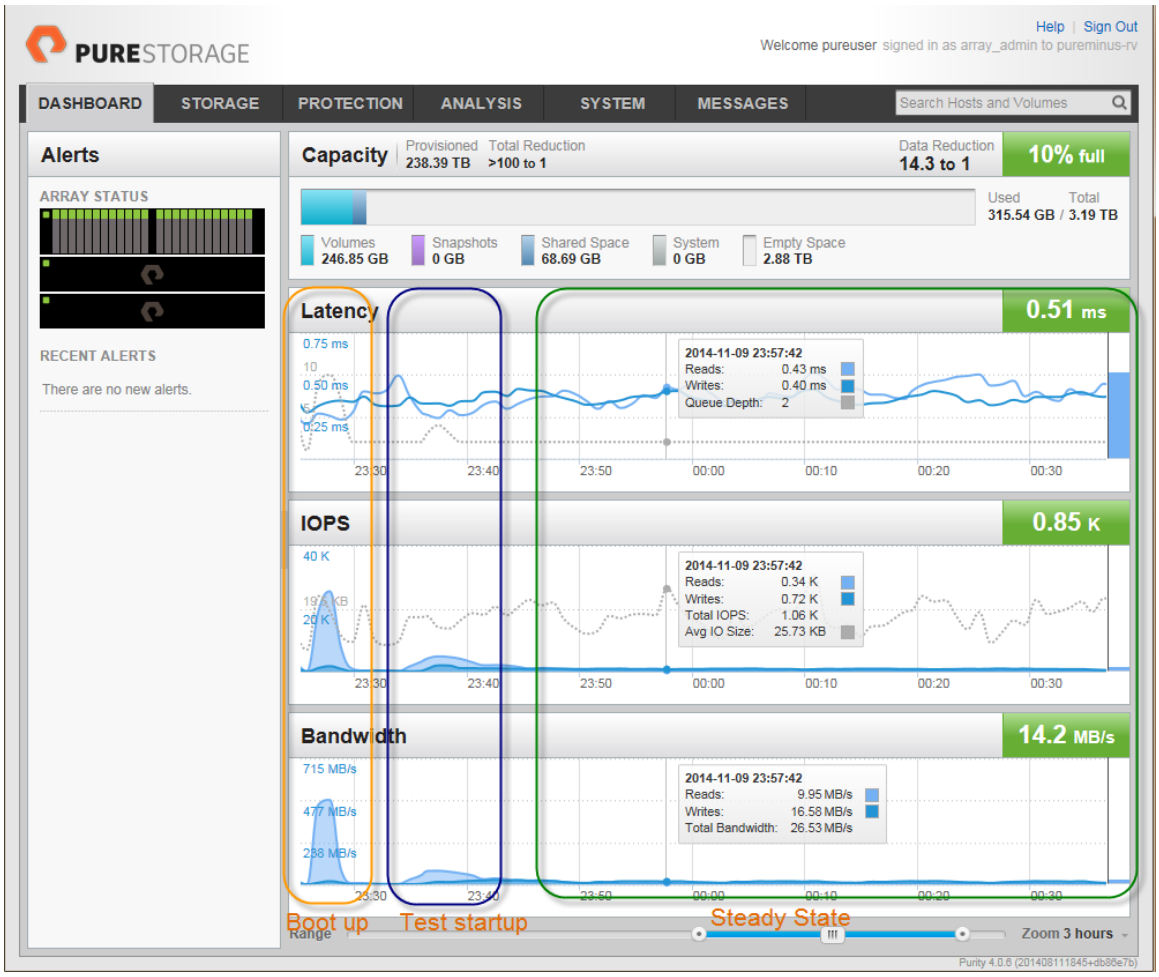

Figure 19.2 Pure Storage resource utilization for one node testing

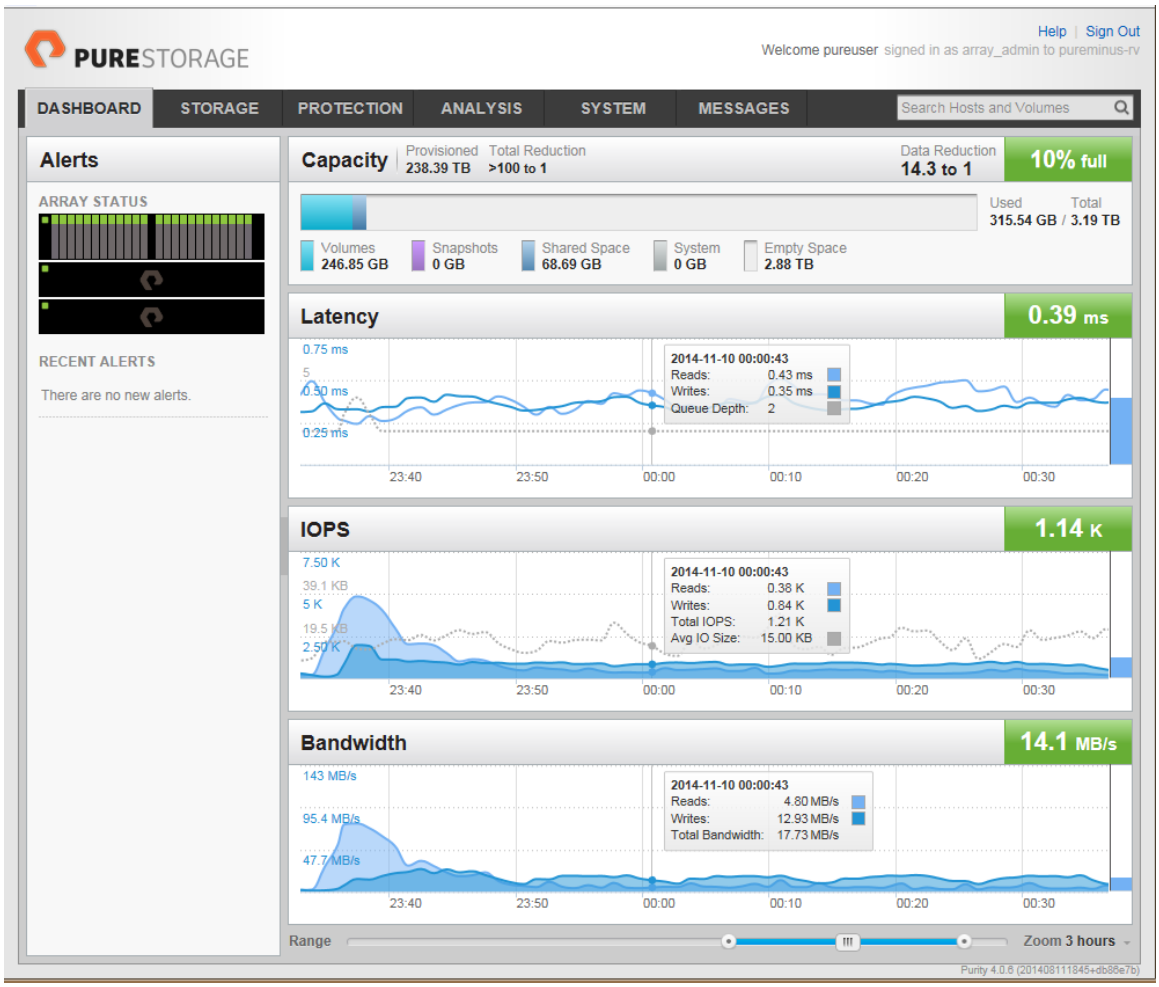

Figure 19.3 Pure Storage resource utilization for one node testing, highlighting the Steady State data

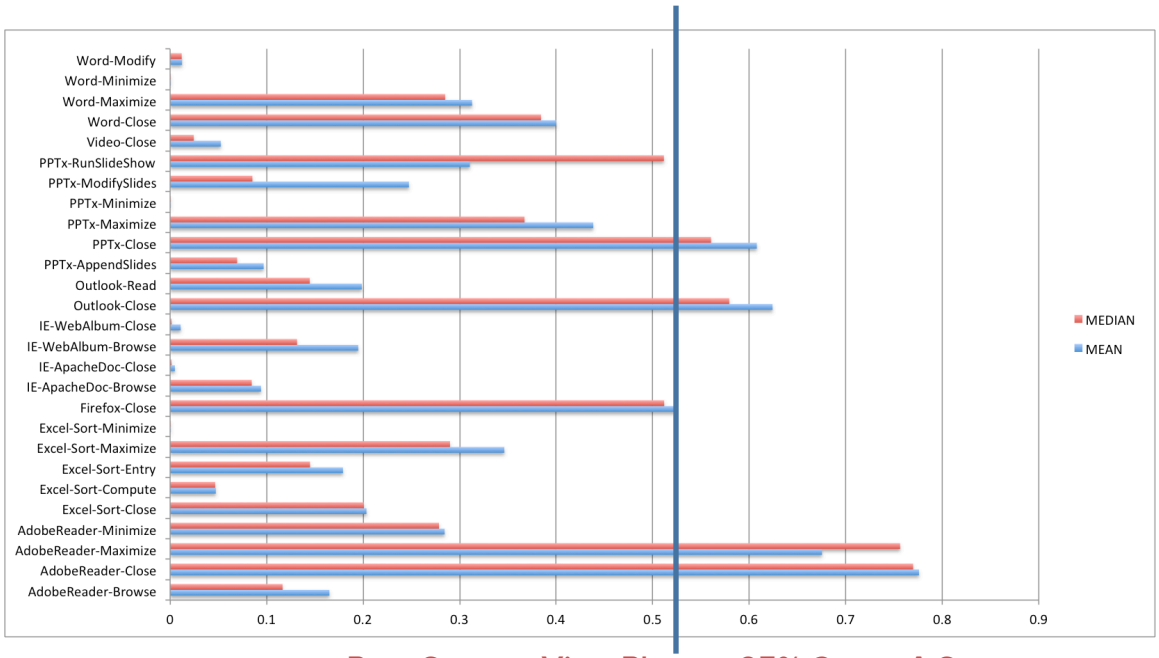

Pure Storage View Planner 95% Group-A Score:
0.52 seconds

Figure 19.4 Single server scale test – VMware View planner time distribution and Score

Two Servers testing

Figure 20 below describes the test setup for a two servers testing. Two Cisco UCS B200-M3 blade servers are deployed and 200 desktops are moved over to them. Using VMware View planner as the workload generator the test data is verified to see if it meets the success criteria and the results are noted.

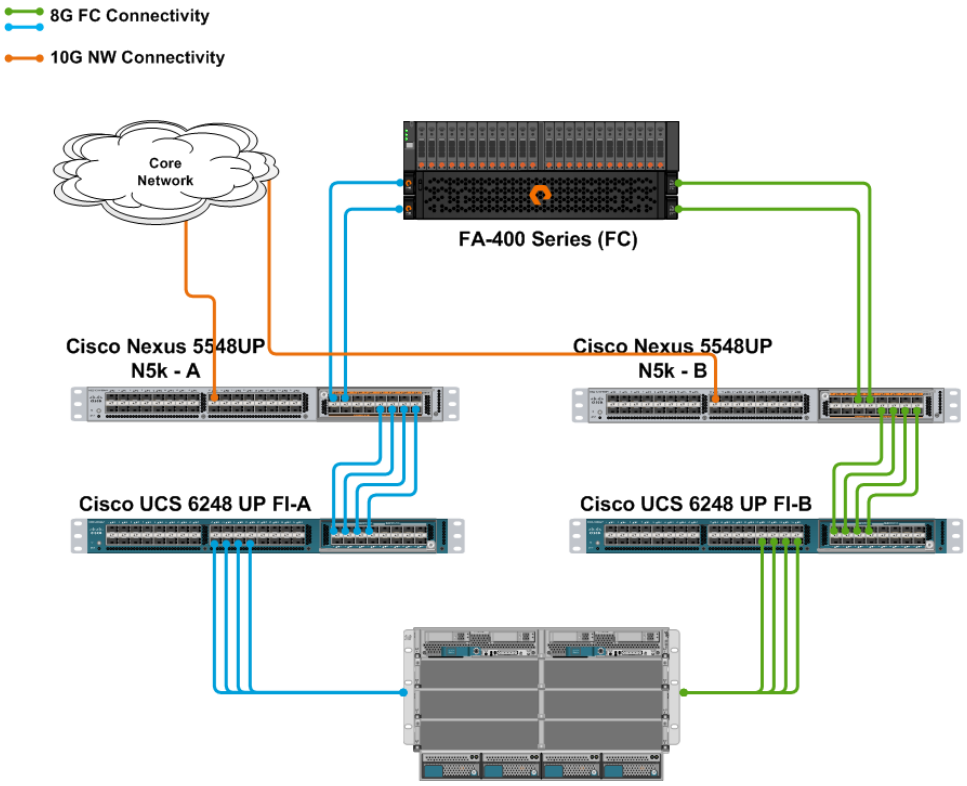

Figure 20. Two servers scale test configuration

Test Results

Running the VMware View planner Benchmark with medium workload on two servers (with all other servers in the Cluster in maintenance mode) we were able to run 200 desktops sessions and the View planner Group-A 95% response time was 0.53 second and Group-B 95% response time was 3.138 seconds. The resource utilization of the ESXi server, storage utilization, and the View planner score is shown below.

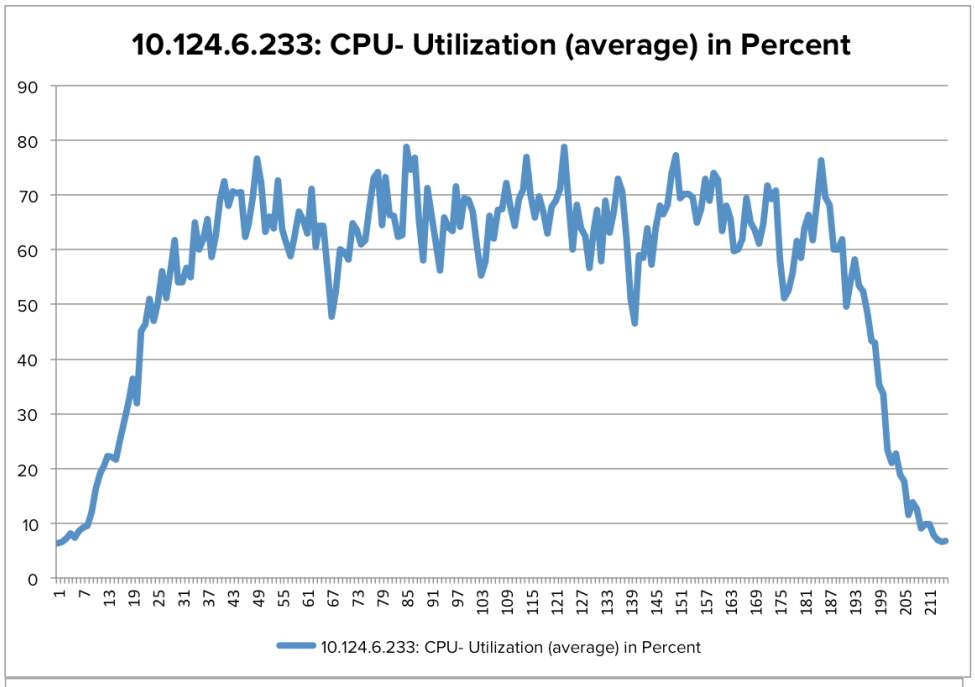
10.124.6.233: CPU- Utilization (average) in Percent
90
80
70
60
50
40
30
20
10
0
1 7 13 19 25 31 37 43 49 55 61 67 73 79 85 91 97 103 109 115 121 127 133 139 145 151 157 163 169 175 181 187 193 199 205 211
10.124.6.233: CPU- Utilization (average) in Percent

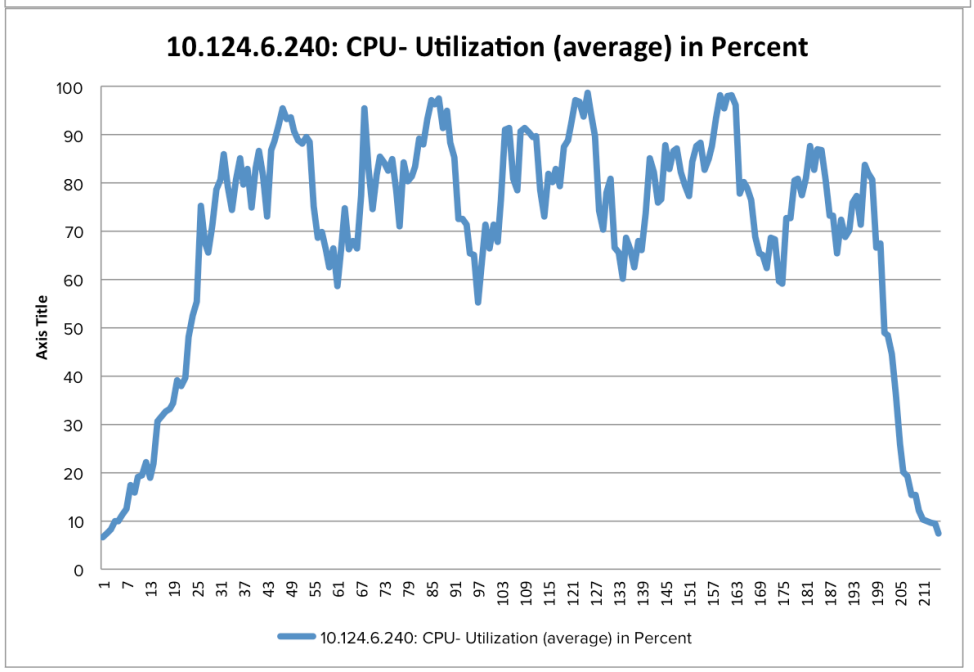
10.124.6.240: CPU- Utilization (average) in Percent
Axis Title
100
90
80
70
60
50
40
30
20
10
0
1 7 13 19 25 31 37 43 49 55 61 67 73 79 85 91 97 103 109 115 121 127 133 139 145 151 157 163 169 175 181 187 193 199 205 211
10.124.6.240: CPU- Utilization (average) in Percent

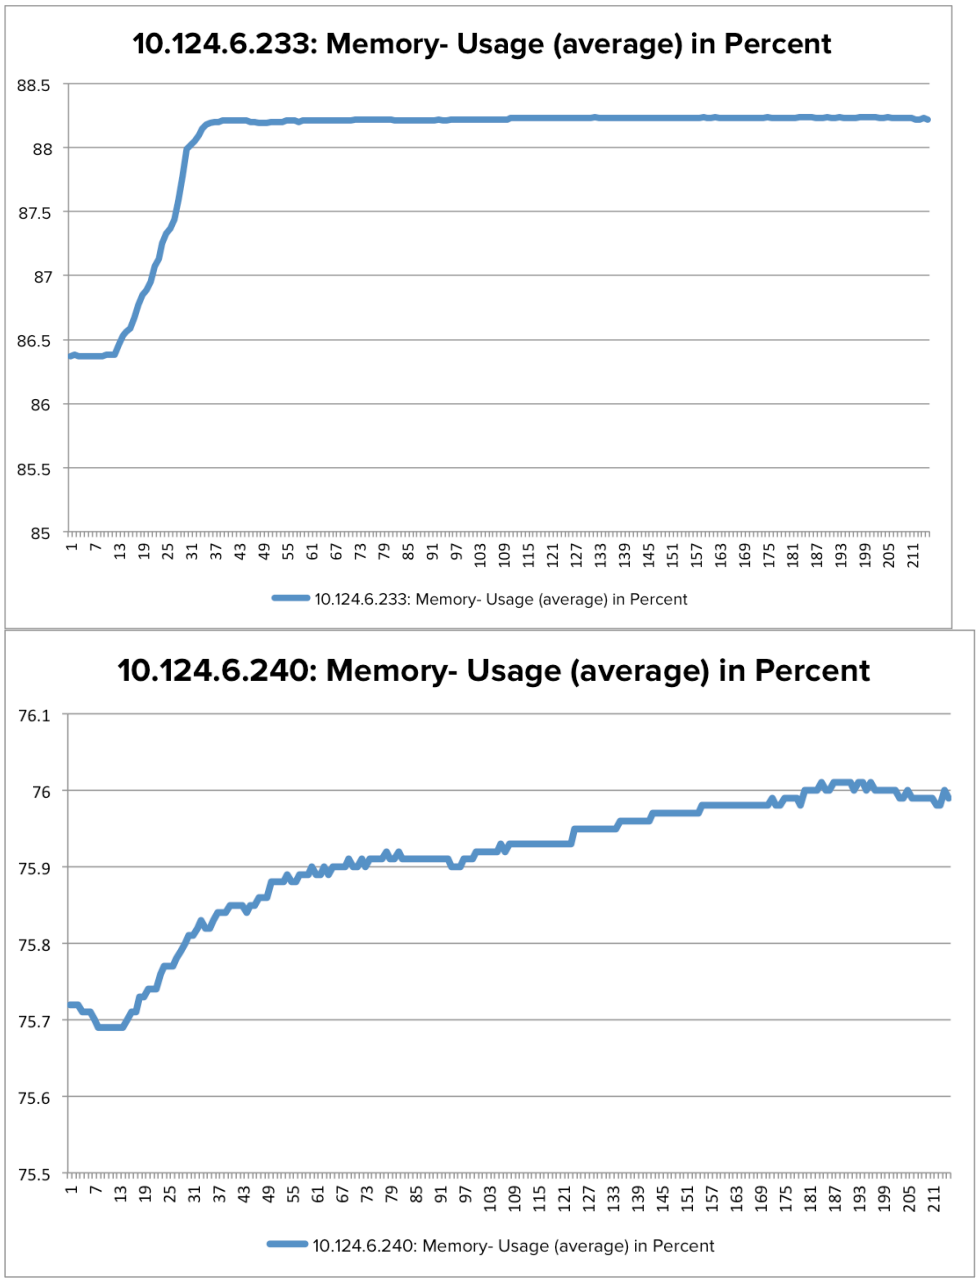

Figure 20.1. CPU / Memory Utilization of servers in two node testing

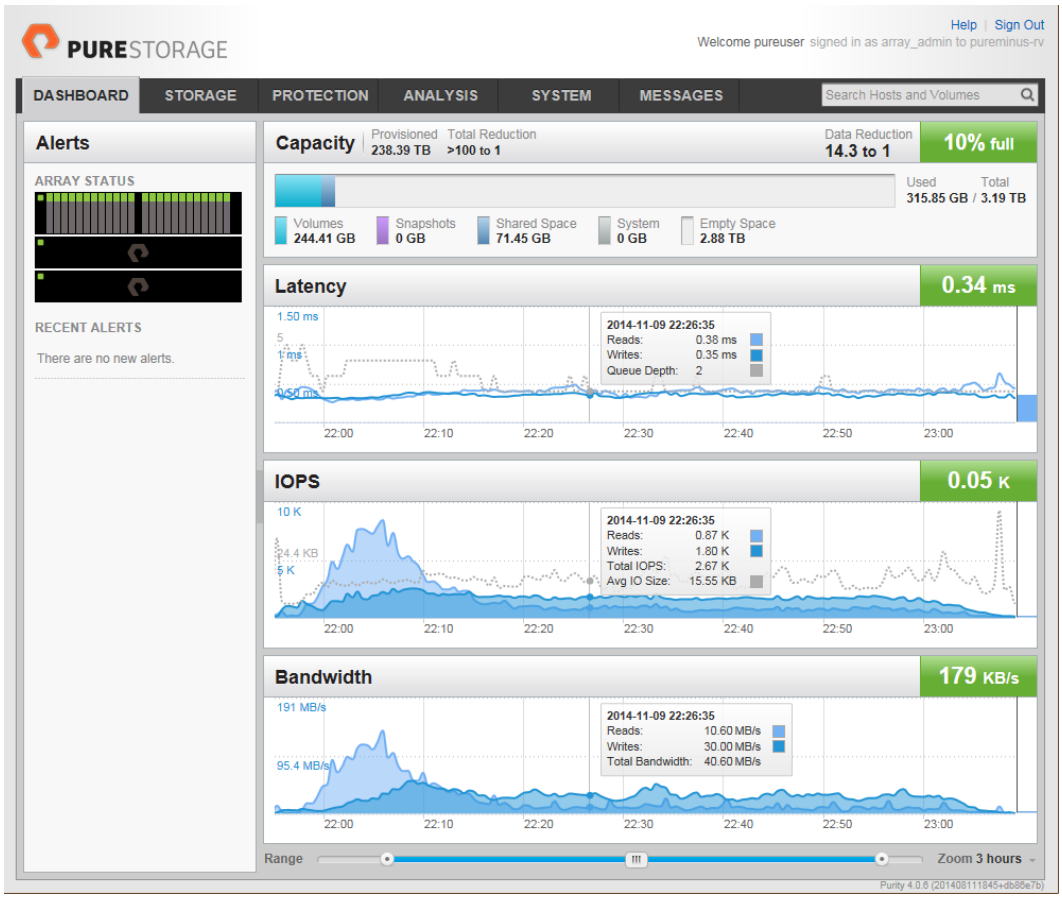

Figure 20.2. Pure Storage resource utilization for two node testing

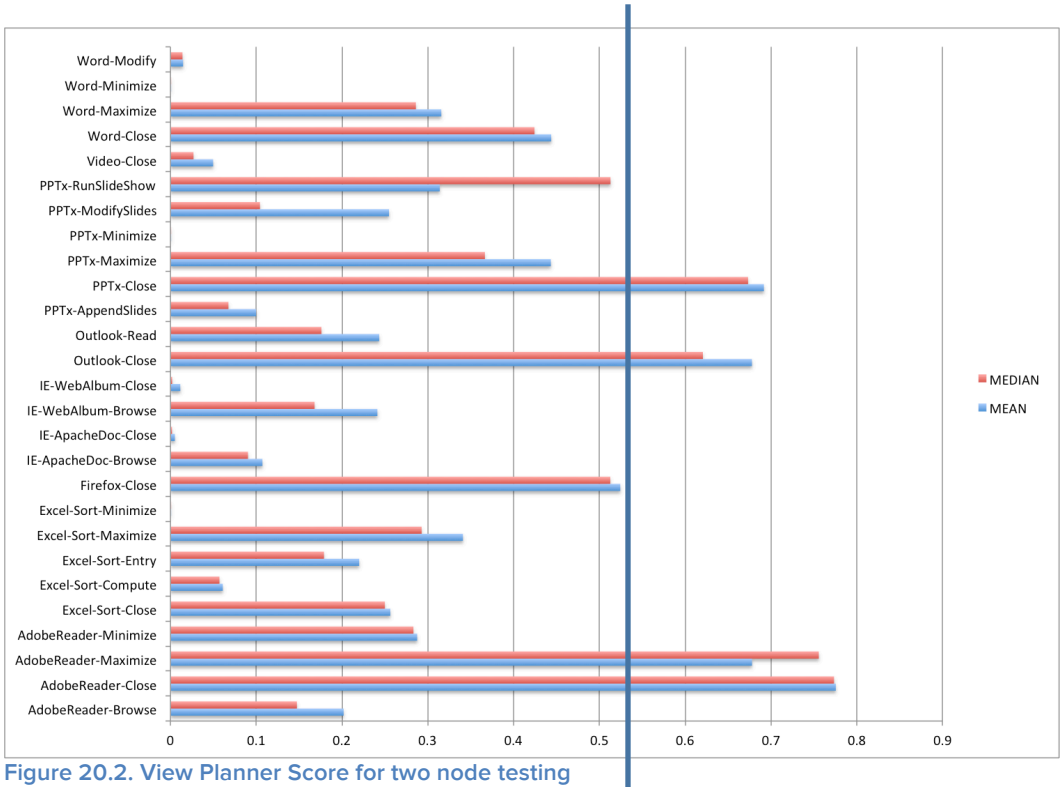

Figure 20.2. View Planner Score for two node testing

Pure Storage View Planner Group-A 95% Score:
0.53 seconds

Four Servers testing

Figure 21 below describes the test setup for a two servers testing. Four Cisco UCS B200-M3 blade servers are deployed and 400 desktops are moved over to them. Using VMware View planner as the workload generator the test data is verified to see if it meets the success criteria and the results are noted.

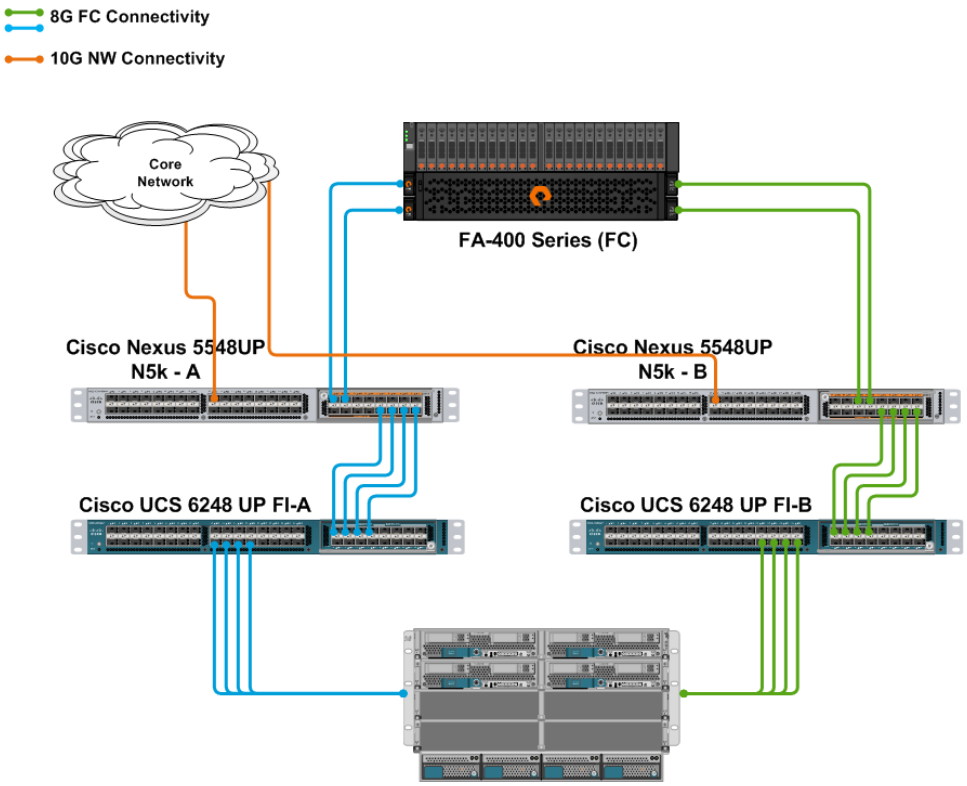

Figure 21. Four servers scale test configuration

Test Results

Running the VMware View planner Benchmark with medium workload on four servers (with all other servers in the Cluster in maintenance mode) we were able to run 400 desktop sessions and the View planner Group-A 95% response time was 0.53 second and Group-B 95% response time was 3.243 seconds. The storage utilization and the View planner score are shown below. We saw similar CPU and memory utilization on the server as seen in one node and two node testing (detailed View planner reports are available on request).

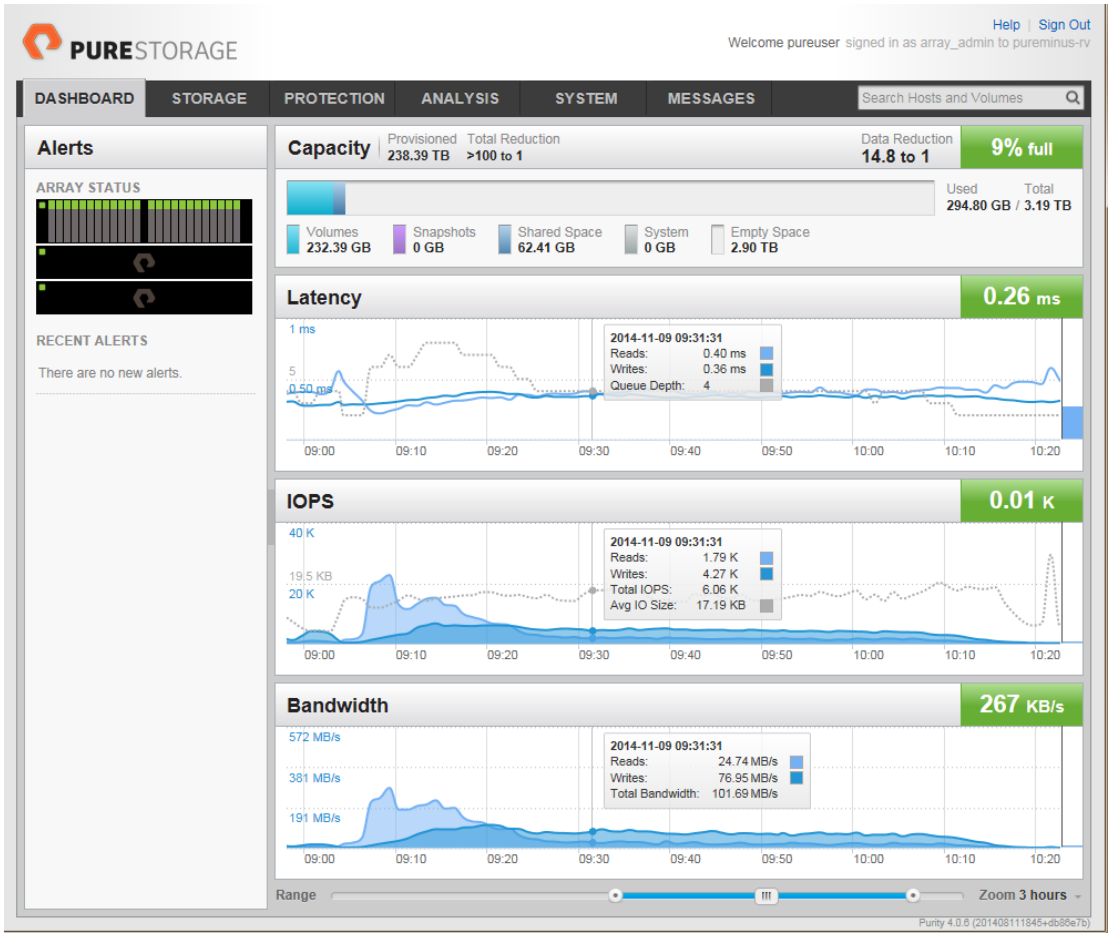

Figure 21.1. Storage Utilization for four node testing

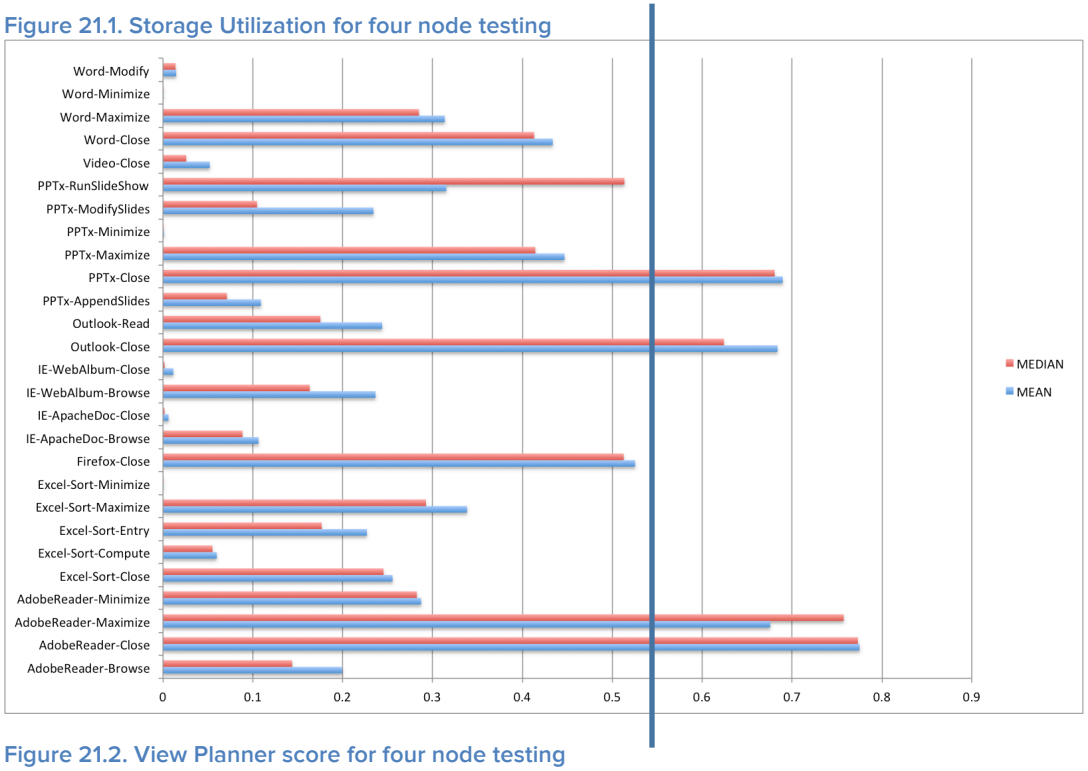

Figure 21.2. View Planner score for four node testing

Pure Storage View Planner Group-A 95% Score:
0.53 seconds

Eight Servers or a Single Chassis testing

Figure 22 below describes the test setup for a two servers testing. Eight Cisco UCS B200-M3 blade servers are deployed and 800 desktops are moved over to them. Using VMware View planner as the workload generator the test data is verified to see if it meets the success criteria and the results are noted.

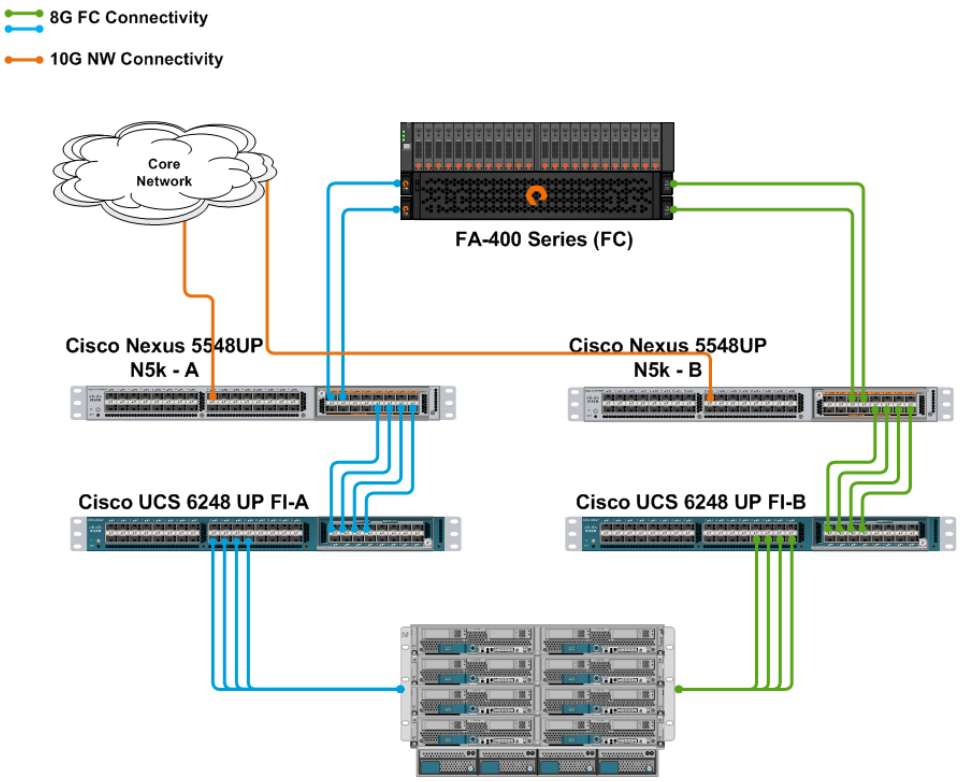

Figure 22. Eight servers or a single chassis scale test configuration

Test Results

Running the VMware View planner Benchmark with medium workload on eight servers we were able to run 800 desktop sessions and the View planner Group-A 95% response time was 0.53 second and Group-B 95% response time was 3.249 seconds. The storage utilization and the View planner score are shown below. We saw similar CPU and memory utilization on the server as seen in one node, two node and four node testing.

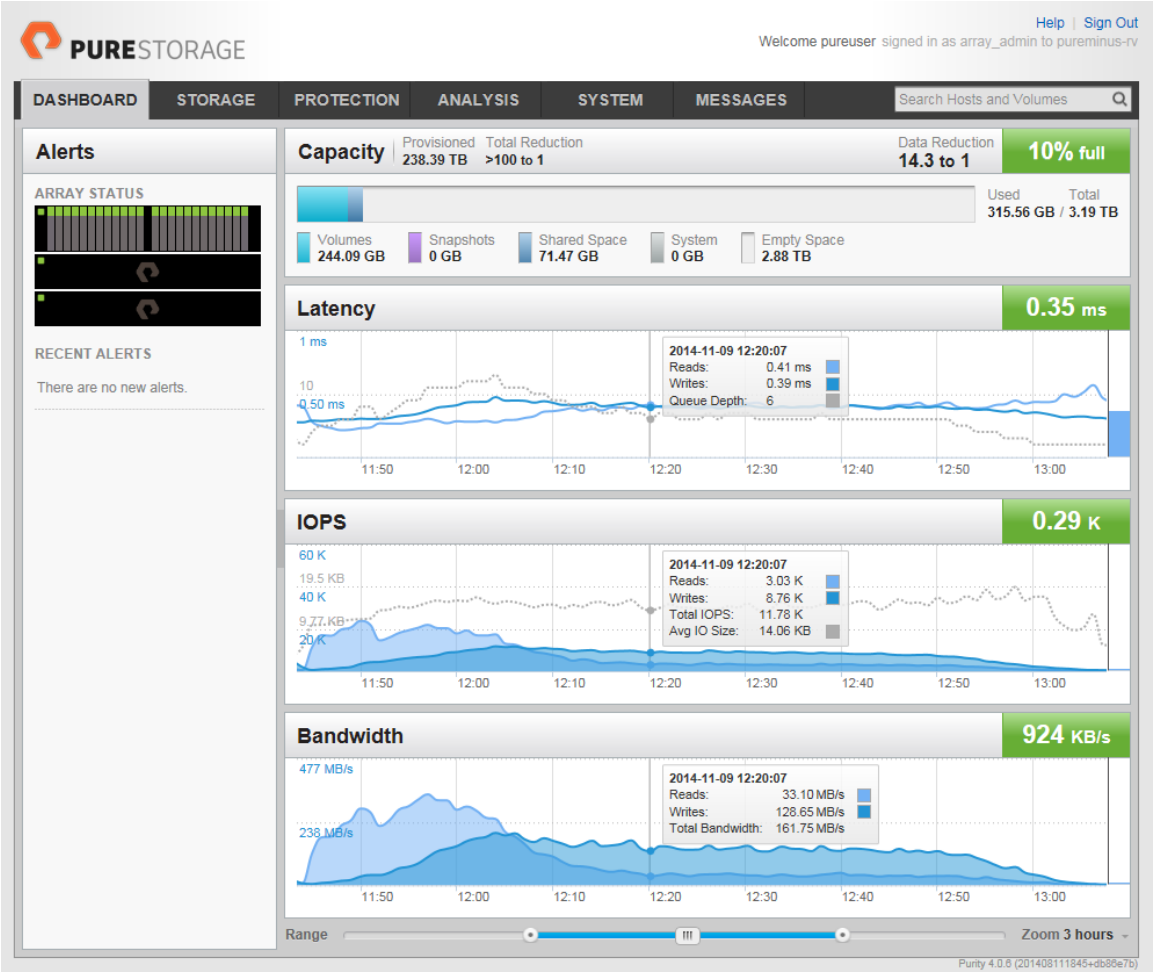

Figure 22.1. Storage utilization for eight node testing

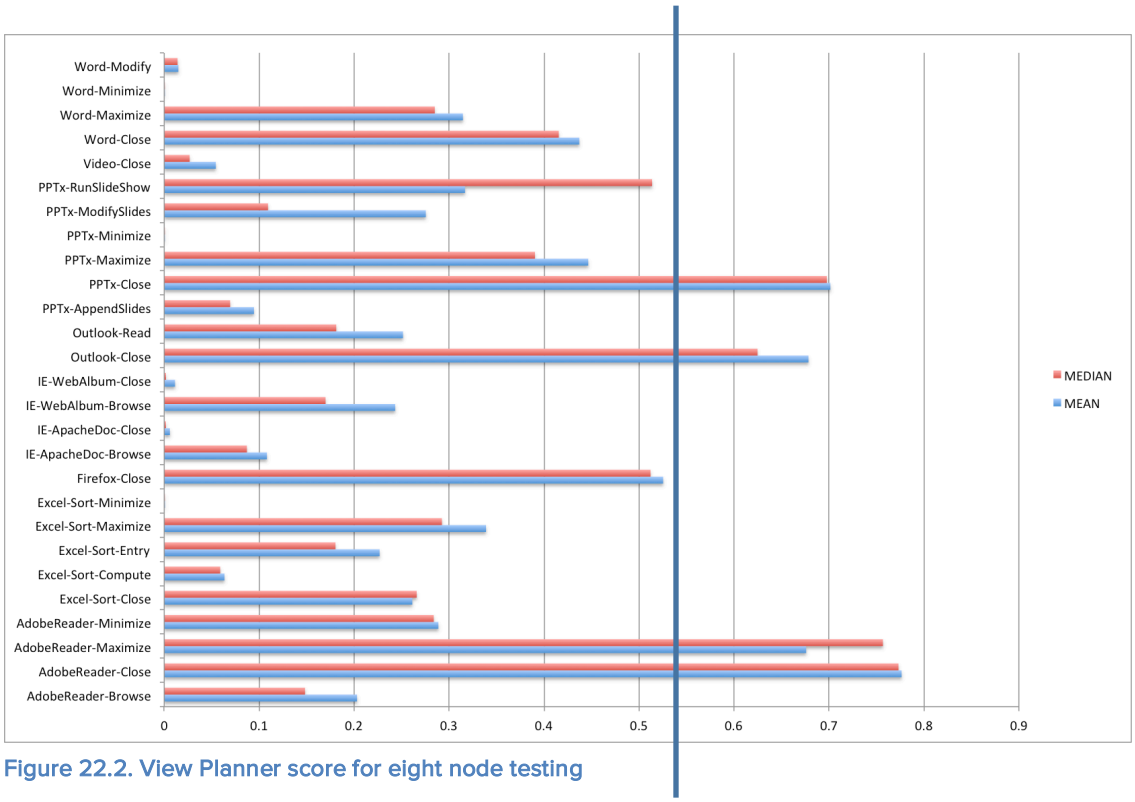

Figure 22.2. View Planner score for eight node testing

Pure Storage View Planner Group-A 95% Score:
0.53 seconds

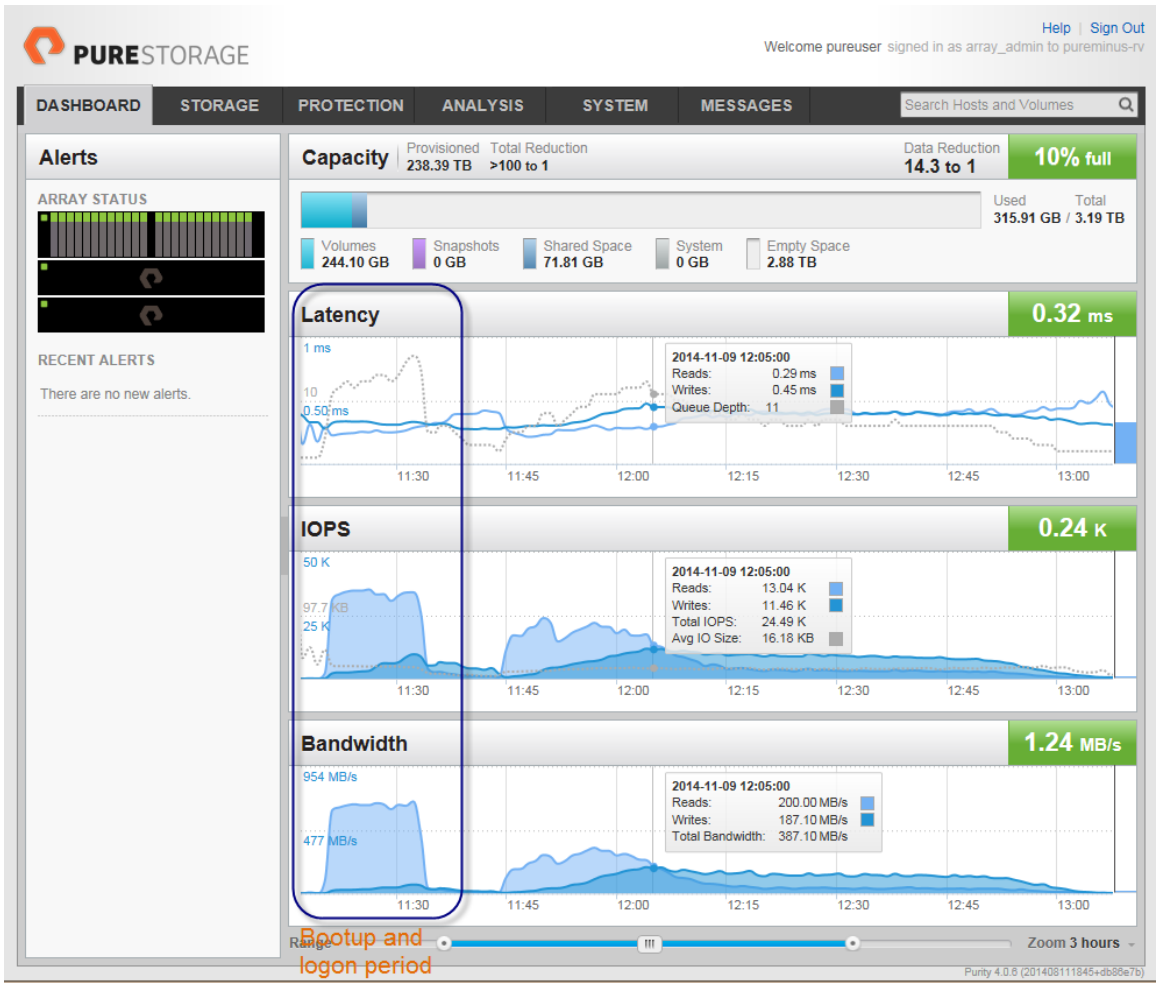

Figure 22.3. Pure Storage dashboard showing the bootup and logon statistics for 800 desktops

Figure 22.3 shows the bootup and logon characterization. Pure Storage FlashArray sustained over 30,000 IOPS with 0.50msec consistent latency and the boot process and logon process finished in less than 15 minutes, which is quite impressive.

Test Summary

Based on the performance data from one, two, four and eight servers testing, we conclude the following:

1. We saw upto 50K IOPS during bootup of 800 desktops during the eight server tests with a consistent < 1msec latency.
2. The Pure Storage FlashArray was consistently showing a 0.5msec latency throughout the tests and an average of 0.35 msec (for both reads and write)
3. The View planner Group-A 95% response time was 0.53 second and Group-B 95% response time was around 3.249 seconds when we scaled the tests from 1 to 8 servers; this demonstrated linear scalability of storage array.

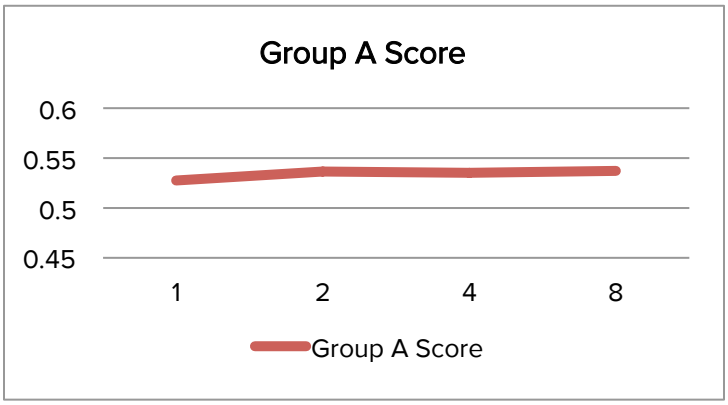

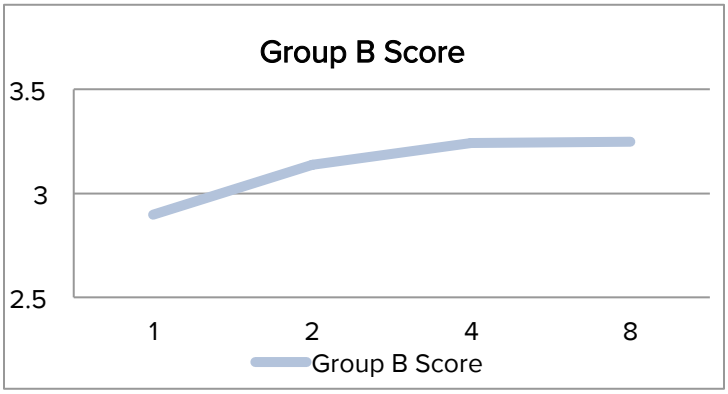

4. Based on the eight servers test with 800 desktops and the low IOP usage, we believe that the FA-405 under test can scale upto 2000 desktops.

5. The data reduction of 14.4-to-1 was constant during and after the test with 1000 desktop image, applications and application data. The FlashArray was 10% full. Figure 23 below shows the data reduction on the entire array.

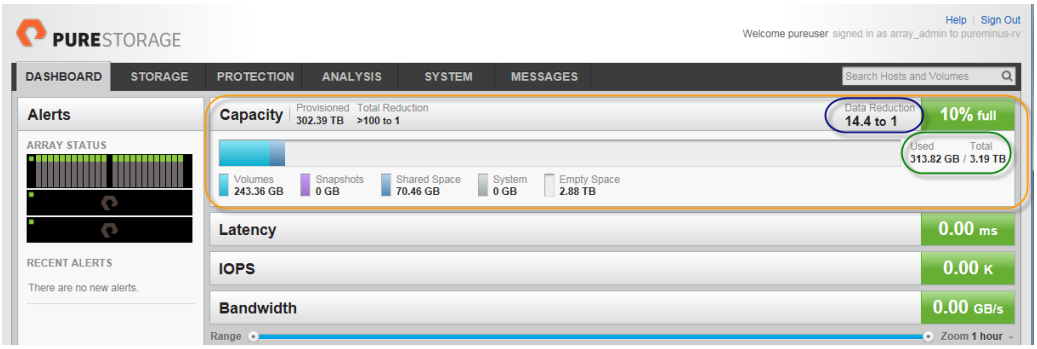

Figure 23. Pure Storage Dashboard showing capacity along with data reduction

Design Considerations and Sizing Guidelines

The space consumption and the IOPS we saw in the 1,000 desktops deployment could easily have been sustained in the smallest FlashArray configuration. As the deployment grows, it is easy to expand capacity by adding more shelves to the array without downtime.

As shown in the Figure 24 below, a pilot can be implemented on a two-controller HA system and ½ drive shelf system. As the deployment passes out of the pilot phase, you can upgrade to a two-controller HA system and full shelf for 1,000 desktops. As your user data grows, additional shelves can be added. Both controllers and shelves can be added without downtime.

Based on the testing we believe that the FA-405 can accommodate up to 2,000 desktops from an IOPS perspective. For a 3,000 desktop deployment, we recommend a fully configured FA-420 with two controllers and two or more drive shelves. The sizing guidelines shown in figure 24 below are approximations based upon the current testing and workload characterization, your actual desktop density may vary depending on how the desktops are configured, whether or not user data is stored in the desktops or the array, and a variety of other factors. Pure Storage recommends a pilot deployment in your user community to fully-understand space and performance requirements.

	FA-405	FA-405	FA-405	FA-420
		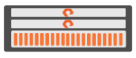	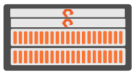	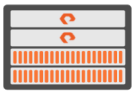
Stage	**POC**	**Go Live**	**Expand**	**Scale-Up**
Users	100-500	Up to 1,000	Up to 2,000	Up to 3,000
Raw Capacity	2.75 TB	5.5 TB	11 TB	11+ TB
Usable VDI Capacity*	20 TB	30 TB	60 TB	60-120 TB
# of UCS Chassis	1	2	3	4

*** Assuming 10X data reduction (Virus scanning & Constant patching)**

Figure 24. Storage sizing for knowledge worker desktops

Adding a new shelf to increase capacity is very straightforward and involves simply connecting SAS cables from the controller to the new shelf that can be done while the array is online. The shelf addition can be done non-disruptively and without any impact to the desktop users (no down time nor any performance impact at all). The Pure Storage FlashArray features stateless controllers, which means all the configuration information is stored on the storage shelves instead of within the controllers themselves. In the event of a controller failure, one can easily swap out a failed controller with a new controller without reconfiguring SAN zoning, which again can be done non-disruptively.

Summary of Findings

- We deployed 800 VMware View based linked clone Windows 7 desktops and ran a realistic load generator with VMware View Planner that simulated 800 users performing common computing tasks, resulting in a best-in-class View Planner Group-A 95% response time score of 0.53 seconds. This score means the majority of the applications (95% of group "A" interactive operations) had a response time of 0.53 second or less, well within the passing score of 1 second. The **View planner Group-B 95% response time** was averaging around **3.13 seconds**.

- Throughout the testing the FlashArray delivered up to 50,000 IOPS and maintained latency under 1 msec, demonstrating the FlashArray's consistent latency and ability to deliver the best all-flash VDI end-user experience at all times. The FlashArray delivers a better desktop experience for end-users than dedicated laptops with SSDs, and doesn't risk the end-user experience by relying on caching as hybrid flash/disk arrays do.

- In total throughout the testing we deployed more than 1,000 desktops (each of 31 GB disk size), together only consuming about 314 GB of physical storage on the FlashArray. This massive data reduction (14.4-to-1) is the result of the high-performance inline data reduction (deduplication and compression) delivered by the FlashArray, which enables using any combination of linked clones or persistent full-clone desktops – both of which reduce to about the same amount of space on the array.

- As tested, the 5.5 TB FlashArray FA-405 delivered best-in-class VDI performance for 800 desktops. Since the FlashArray was significantly under-utilized throughout the testing on both a capacity and performance basis, the array could have supported upto 2,000 desktops.

- Throughout the testing we performed common VDI administrator operations and found a drastic reduction in time for cloning persistent desktops, (re)booting desktops, and other day-to-day virtual desktop operations. Taken together these operational savings deliver substantial efficiency gains for VDI administrators throughout the VDI day.

- The power footprint for the tested FA-405 FlashArray was 8 Amps (110V) which is a fraction of any mechanical disk storage array available in the marketplace. This configuration consumed four rack units (4 RU) in data center space.

- This reference architecture can be treated as a 800 desktop building block. Customers can add more server and infrastructure components to scale the architecture out to 1,000s of desktops. Based on the results, we believe a single FA-405 can support up to 2,000 desktops with any mix of linked clones and/or persistent desktops.

Conclusions

We set out to prove that Pure Storage FlashArray is an ultimate storage for virtual desktop deployment and we achieved the same score while we scaled the desktops from one server to eight servers while running an industry standard desktop workload generator. The View planner score reflect the latency observations on the Pure Storage dashboard that showed 0.5msec latency through out the testing and even during boot storm, login storm and test ramp-up scenarios. Factors like rapid deployment, ease of storage management, lower storage cost, lower power, rack space savings, and lower cooling requirements make the TCO for large scale deployments even more attractive.

About the author

Ravindra “Ravi” Venkat is a Virtualization Solutions Architect at Pure Storage for over three and half years where he strives to be the company’s expert at the intersection of flash and virtualization. Prior to that he held a similar role at Cisco for three plus years where he helped drive the virtualization benefits of Cisco's new servers - Unified Computing System (UCS). He helped build reference architectures and virtualization solutions with Citrix, VMware, and Cisco UCS with various storage partners that are still being used today.

Prior to that he was part of the storage ecosystem engineering team at VMware for three years, and a lead engineer at VERITAS working on storage virtualization, volume management and file system technologies for the prior eight years.

Ravi maintains a blog at http://www.purestorage.com/blog/author/ravi and you can follow him on twitter @ravivenk.

Appendix A

Steps to enable VM disk space reclamation in VMware Horizon View 6.0.

1. Login to VMware Horizon View administrator console
2. Under View configuration select “Servers”
3. Select the VMware vCenter server tab
4. Select the vCenter Server
5. Select Edit
6. In the Edit vCenter Server window, select the storage tab
7. Set the Storage setting – Reclaim VM disk space
8. Select “OK”
9. Note the View Storage Accelerator is left disabled.

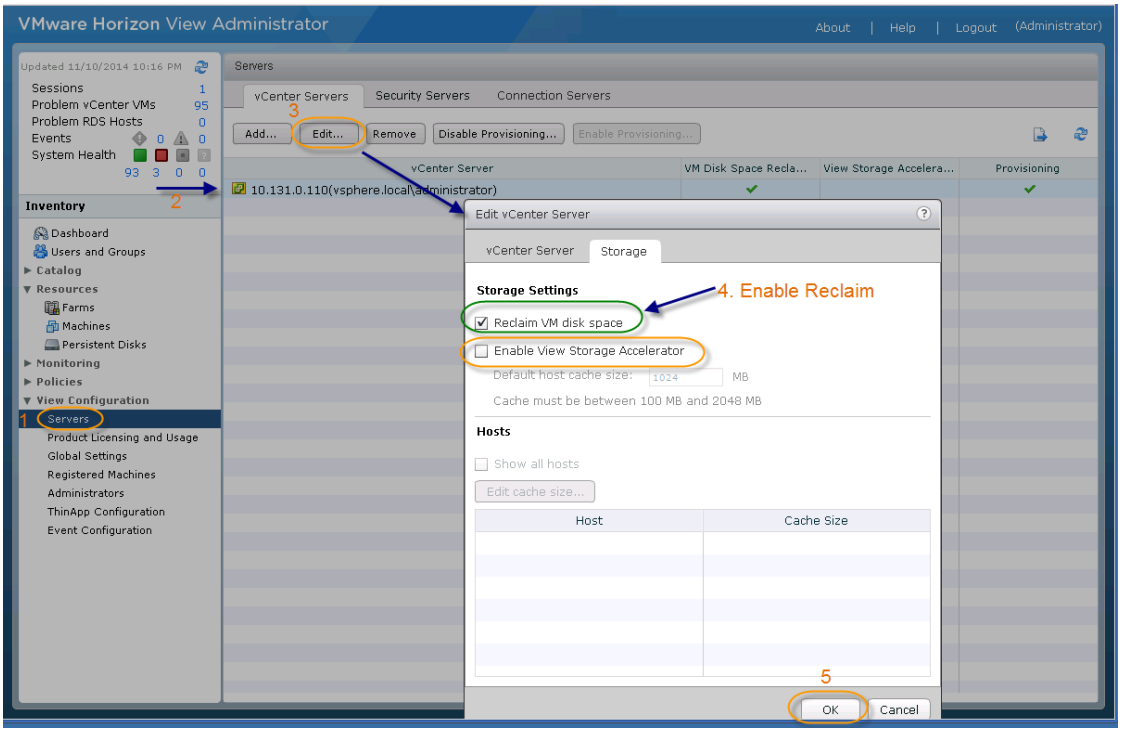

Pure Storage, Inc.
Twitter: @purestorage

650 Castro Street, Suite #400
Mountain View, CA 94041

T: 650-290-6088
F: 650-625-9667

Sales: sales@purestorage.com
Support: support@purestorage.com
Media: pr@purestorage.com
General: info@purestorage.com