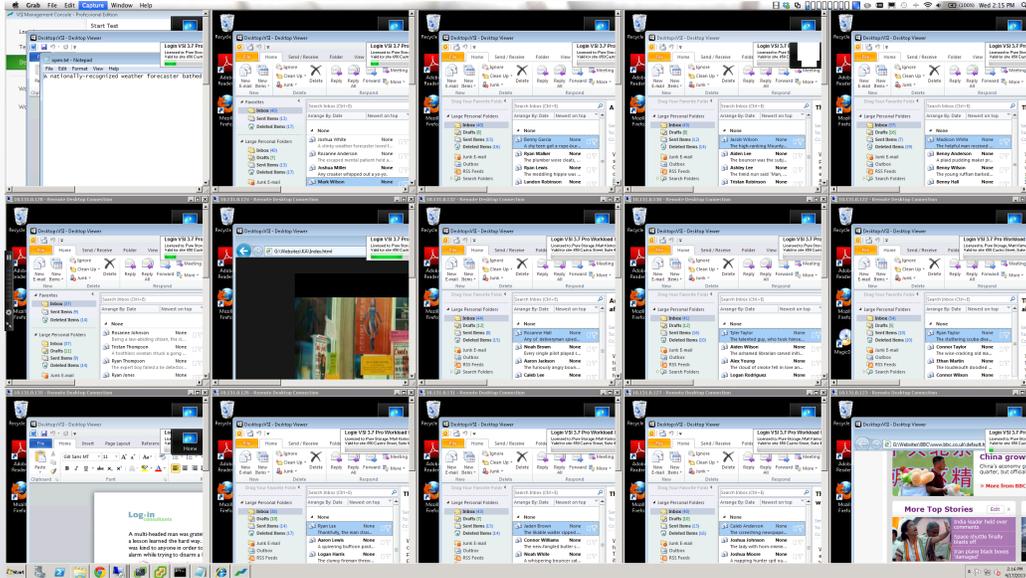




PURESTORAGE



Pure Storage Reference Architecture with Citrix® XenDesktop® 5.6

Overview

This document describes a reference architecture for deploying virtual desktops on the Pure Storage FlashArray using Citrix® XenDesktop® 5.6, Citrix® Provisioning Services 6.1, vSphere 5.1 Hypervisor and Microsoft Windows 7. Pure Storage has validated the reference architecture with Login VSI Pro 3.7 medium workload in its lab – this document presents performance and scalability testing results and offers implementation 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 Citrix XenDesktop based Windows 7 desktops on the FlashArray by deploying 500 virtual desktops using Citrix Provisioning Services (PVS) and running the Login VSI Pro workload to simulate real user interaction and experience in a VDI workload. In addition, we highlight the benefits of the Pure Storage FlashArray including inline data reduction and low latency and show how all-flash storage can dramatically improve both the end-user and administrative experience of VDI compared to traditional disk-based storage.

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 Citrix XenDesktop based virtual desktops on the FlashArray. A working knowledge of VMware vSphere, Citrix XenDesktop, server, storage, network and data center design is assumed but is not a prerequisite to read this document.

Summary of Findings

- We deployed 500 Citrix XenDesktop based Windows 7 desktops and ran a realistic load generator with Login VSI Pro 3.7 that simulated 500 users performing common computing tasks, resulting in a best-in-class score of 720. This score means the majority of the applications had a response time of 720 milliseconds or less, well within the VSI_{max} of 4000 milliseconds.
- Throughout the testing the FlashArray delivered and maintained latency under 1 millisecond, 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 5 GB write cache), together only consuming less than 1 TB of physical storage on the FlashArray. This massive data reduction (10-to-1) is the result of the high-performance inline data reduction (deduplication and compression) delivered by the FlashArray, which enables us to host a massive amount of desktops.
- Based on the current testing, the 11 TB FlashArray, FA-320, has the ability to deliver best-in-class VDI performance at a cost of \$100/desktop. Since the FlashArray was significantly under-utilized throughout the testing on both a capacity and performance basis, the array could have supported 1,000s more desktops, or a smaller array could have been used, either of which would have reduced the \$/desktop cost even further.
- Throughout the testing we performed common VDI administrator operations and found a drastic reduction in time for patching desktops, adding new desktop groups, (re)booting desktops, logon/logoff storms, 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-320 FlashArray was 9 Amps (110V) which is a fraction of any mechanical disk storage array available in the marketplace. This configuration consumed eight rack units (8 RU) in data center space.
- This reference architecture can be treated as a 500 desktop building block. Customers can add more Provisioning Services (PVS) to the PVS farm, and scale their Desktop Controller (DDC), hosting server and infrastructure components to scale the architecture out to 1,000s of desktops. Based on the results, a single FA-320 can support up to 5,000 desktops with any mix of desktop user types.

Introduction

The IT industry has been abuzz over the past several years promoting the idea of VDI: virtualizing and centralizing desktops to enable IT to deliver a more secure, manageable, less costly, and ultimately more mature end-user computing model. While the dream of pervasive deployment of virtual desktop infrastructure has been discussed and tried for literally decades, the recent explosion of x86 virtualization, and the availability of commodity scalable server architectures with increasingly large amounts of CPU power and centralized memory have made the promise of VDI much closer to reality. In fact, sophisticated IT departments are finding that with the right investments in infrastructure, VDI can deliver a client computing model that is both better for the end-user (a truly mobile, multi-device computing experience with better performance than dedicated devices) and better for the IT staff (centralized management, consistent security and policy enforcement, resiliency through device independence, and enablement of “bring your own device” BYOD models).

So if VDI comes with so many potential advantages, why has adoption of VDI been so slow? The reality is that the path to achieving the VDI promise land is a difficult one, and many organizations have abandoned their VDI initiatives outright or in partial stages of deployment. There are many reasons, but most failed deployments boil-down to three key issues:

- **Too expensive:** VDI is often positioned as a technology to reduce desktop cost, but in reality most find that they are unable to achieve the promised ROI due to infrastructure costs. In particular, server, networking, and storage devices are often dramatically more expensive than dedicated desktops/laptops.
- **Poor end-user experience:** if VDI isn't implemented properly, the end result is slow or unavailable desktops that can lead to user frustration and lost productivity.
- **Too difficult to manage:** VDI shifts the desktop administration burden from the end-users to IT staff. While this affords many security and administrative benefits, it also means more work for often burdened IT staff, especially if the VDI environment itself isn't architected correctly.

More often than not, one of the chief contributors to all three of these failure modes is storage. Traditional disk-based storage is optimized for high-capacity, modest performance, and read-heavy workloads – the exact opposite of VDI which is write-heavy, very high performance, and low-capacity. The result is that as performance lags, spindle after spindle of legacy disk storage has to be thrown at VDI, causing a spike in infrastructure costs and a spike in management complexity.

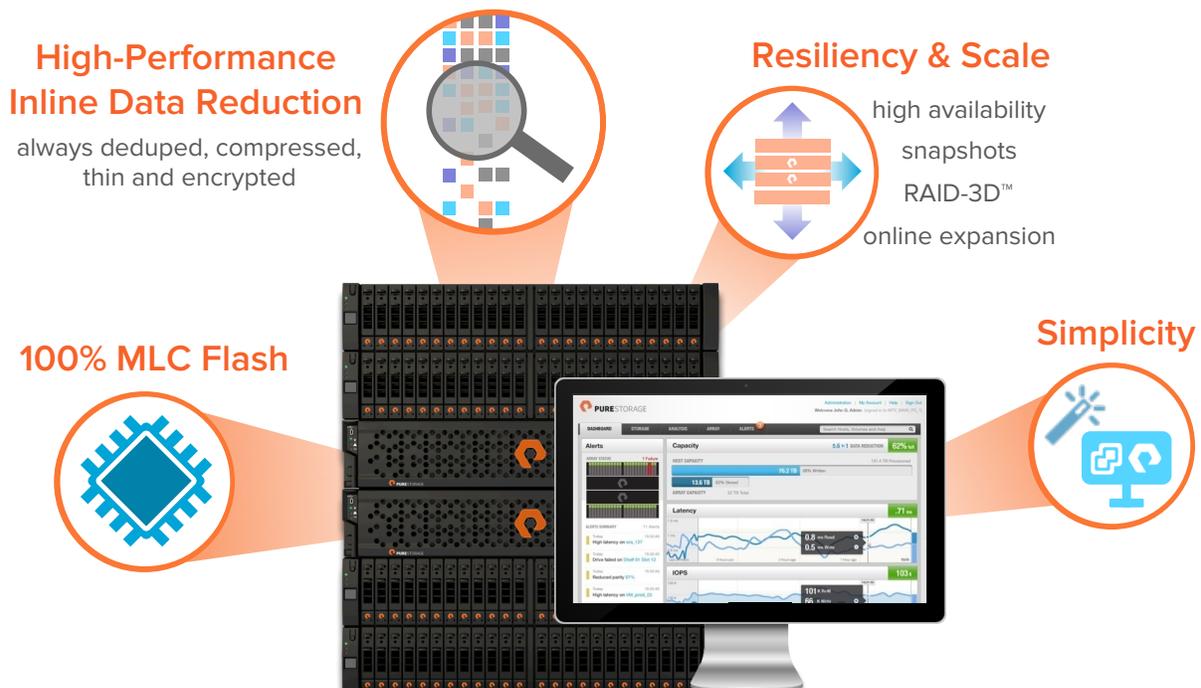
In this reference architecture for virtual desktops we're going to explore how a new, 100%-flash based approach to VDI can help overcome the key VDI failure traps, and help deliver a VDI solution that both end-users and IT administrators will love. We'll start with a high level overview of the Pure Storage FlashArray, then we'll go over the test infrastructure components put together for this work and dive into the details of each component. Finally, we'll discuss the results of the Login VSI Pro load generator and the operational benefits of using the Pure Storage FlashArray for virtual desktop deployment.

The Pure Storage All-Flash Solution for VDI

Introducing Pure Storage

Pure Storage was founded with a simple goal in mind: that 100% flash storage should be made affordable, so that the vast majority of enterprise applications can take advantage of the potential advances that flash memory affords. As such we designed our core product, the Pure Storage FlashArray, from the ground-up for the unique characteristics of flash memory.

The FlashArray's entire architecture was designed to reduce the cost of 100% flash storage. It combines the power of consumer-grade MLC flash memory with inline data reduction technologies (deduplication, compression, thin provisioning) to drive the cost of 100% flash storage to be inline or under the cost of traditional enterprise disk storage. Data reduction technologies are particularly effective in VDI environments, typically providing >5-to-1 reduction for stateless desktops and >10-to-1 reduction for stateful desktops.



It's important to note that unlike some flash appliances, the FlashArray was designed with enterprise-class scale and resiliency in mind. That means a true active/active controller architecture, online capacity expansion, and online non-disruptive code upgrades. The FlashArray also employs a unique form of RAID protection, called RAID-3D™, which is designed to protect against the three failure modes of flash: device failure, bit errors, and performance variability.

Last but not least, the FlashArray is the simplest enterprise storage that you'll ever use. We've designed from the start to remove the layers of complexity of LUN, storage virtualization, RAID, and caching management common in traditional arrays, and have integrated management directly into VMware vSphere's Web Client, making management of a VDI environment seamless.

Reference Architecture Design Principles

The guiding principles for implementing this reference architecture are:

- Create a scalable building block that can be easily replicated at any customer site using a customer's chosen server and networking hardware.
- Implement every infrastructure component in a VM. This ensures easy scale-out of infrastructure components when you go from 500 to 5,000+ virtual desktops.
- Create a design that is resilient, even in the face of failure of any 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.
- Take advantage of inline data reduction and low latency of the Pure Storage FlashArray to push the envelope on desktops-per-server density.
- Avoid tweaks to make the results look better than a normal out-of-box environment.

Solution Overview

Figure 1 shows a topological view of the test environment for our reference architecture. The Citrix XenDesktop 5.6 infrastructure components were placed on a dedicated host. We deployed 1,000 PVS based desktops but due to the high CPU requirement of Login VSI Pro 3.7 medium workload, we were able to run the load on 500 desktops. The infrastructure, virtual machines and desktops were all hosted on a single 11 TB FlashArray FA-320 (although the workload would have easily fit on the smallest 2.75TB FA-320 or FA-310 as well). VMware vSphere and Citrix XenDesktop best practices were used in addition to the requirements as mandated by the Login VSI document [See reference 4].

The tested configuration included:

- One 11 TB Pure Storage FlashArray (FA-320) in HA configuration, including two controllers and two disk shelves:
 - One 40 TB volumes was carved out of the Pure FlashArray to host 1,000 desktops
 - A separate 600 GB volume was used to hold all the infrastructure component
- Eight Intel Xeon x5690 based commodity servers with 192 GB of memory running ESXi 5.1 were used to host the desktops
- Three dedicated Intel Xeon x5690 based commodity servers were used to host all of the infrastructure virtual machines:
 - Active directory, DNS, and DHCP
 - Citrix Provisioning Services (PVS) 6.1

- Citrix Desktop Delivery Controller (DDC)
- VMware Virtual Center Server 5.1
- SQL server for both Virtual Center and PVS database
- A Citrix License server VM which also hosted the Login VSI Share
- Login VSI launcher each running 35 sessions

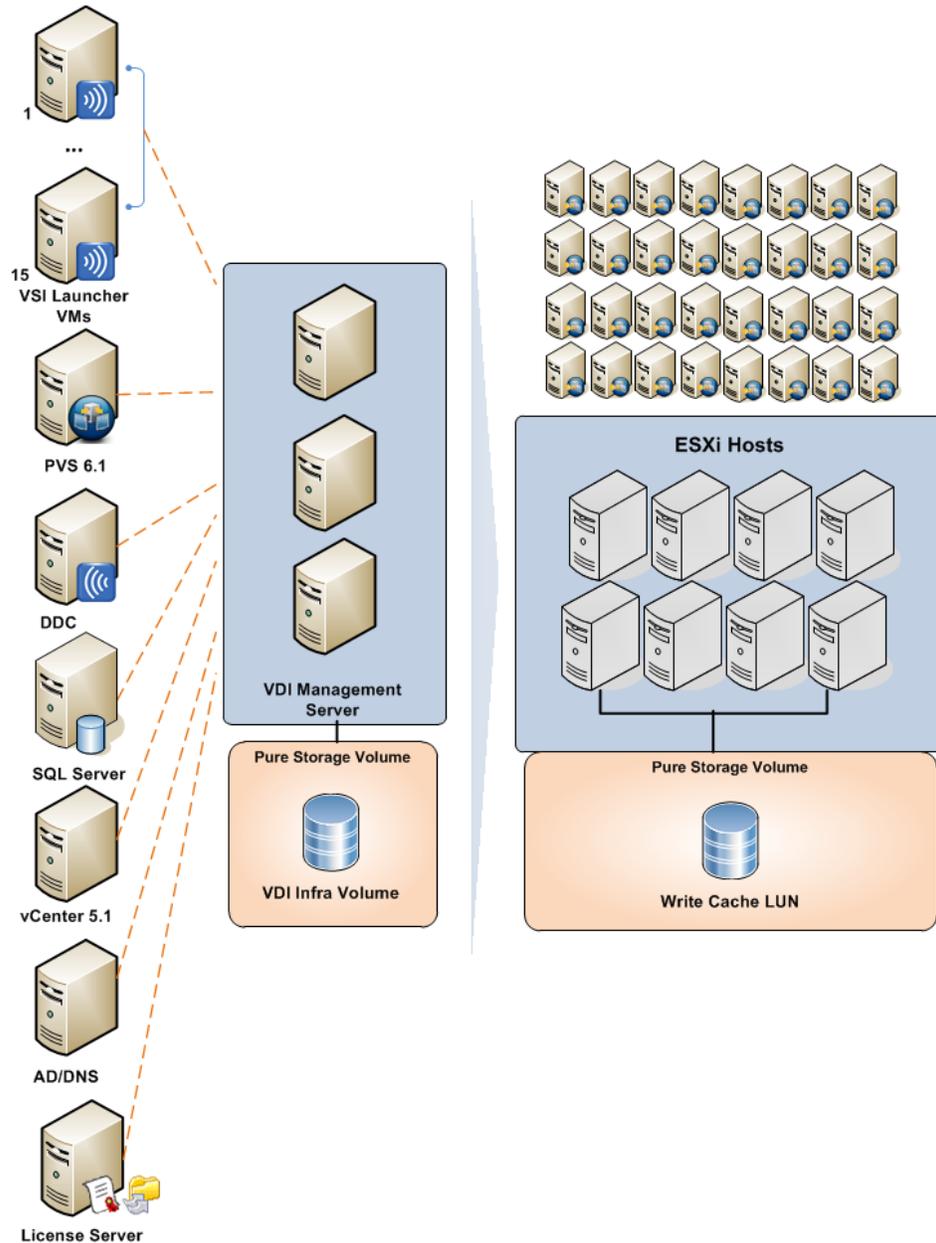


Figure 1: Test Environment overview of Citrix XenDesktop deployment with infrastructure components, ESX hosts and Pure Storage FlashArray volumes.

Reference Architecture Configuration

This section describes the Reference Architecture configuration in brief. Later sections have detailed hardware and software configurations.

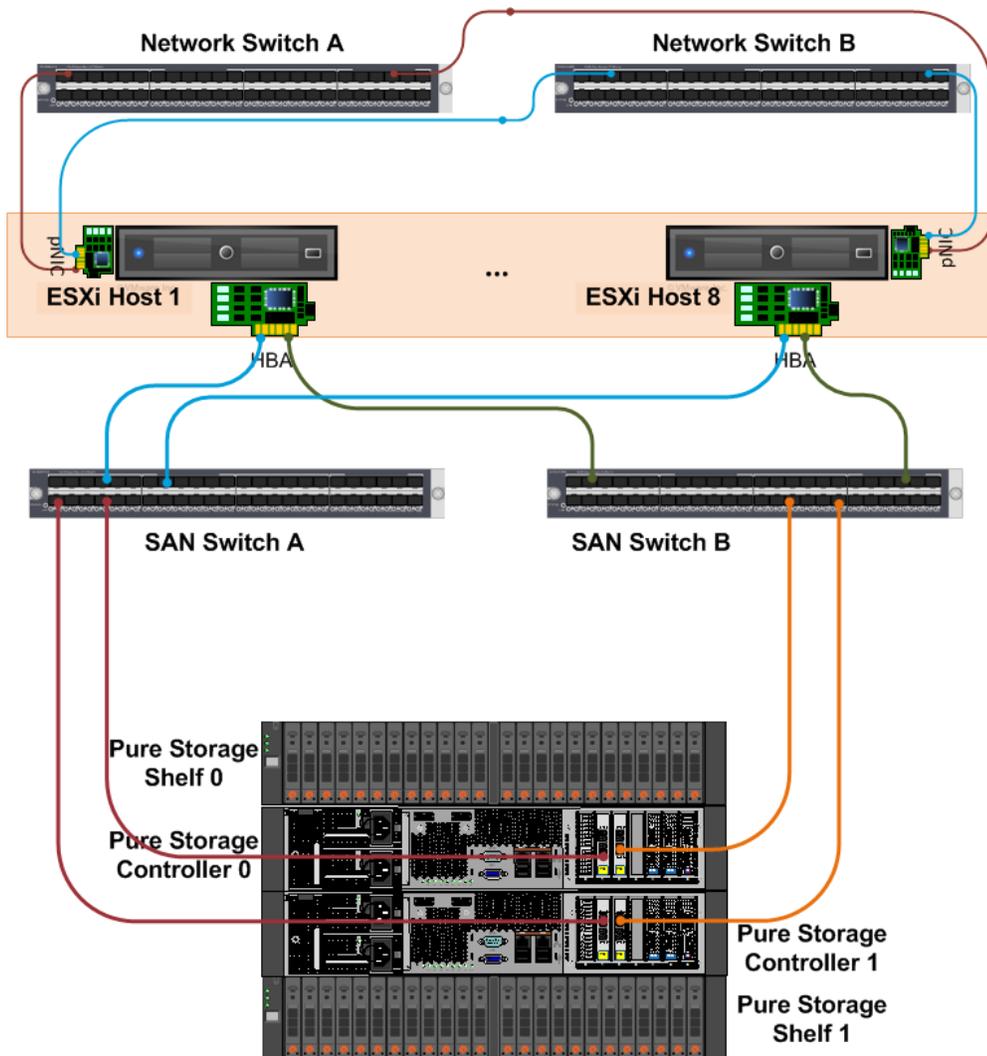


Figure 2: Detailed Reference Architecture Configuration

Figure 2 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 Fiber Channel (FC) 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 with two or more Ethernet ports teamed. The management ports (not shown in the diagram) were connected to a separate management switch.

Hardware Configuration

Pure Storage FlashArray FA-320 configuration

The FlashArray FA-320 configuration has two active/active controllers and two shelves of 5.5 TB of raw flash memory for a total of 11 TB of raw storage. Four Fiber Channel ports were connected to two Cisco MDS 9148 8Gb SAN switches in a highly redundant configuration as shown in Figure 2. Table A below describes the specifications of the FlashArray FA-320.

Component	Description
Controllers	Two active/active controllers which provided highly redundant SAS connectivity (24 Gb) to two shelves and were interconnected for HA via two redundant InfiniBand connections (40 Gb).
Shelves	Two flash memory shelves with 22 SSD drives, 22 X 256 GB or a total raw capacity of 11 TB (10.3 TB).
External Connectivity	Four 8 Gb Fibre Channel ports or four 10 Gb Ethernet ports per controller, total of eight ports for two controllers. As shown in figure 2, only four Fibre Channel ports (two FC ports from each controller) were used for this test.
Management Ports	Two redundant 1 Gb Ethernet management ports per controller. Three management IP addresses are required to configure the array, one for each controller management port and a third one for virtual port IP address for seamless management access.
Power	Dual power supply rated at 450W per controller and 200W per storage shelf or approximately 9 amps of power.
Space	The entire FA-320 system was hosted on eight rack unit (8 RU) space (2 RU for each controller and 2 RU for each shelf).

Table A: Pure Storage FlashArray FA-320 specifications

There was no special tweaking or tuning done on the FlashArray; we do not recommend any special tunable variables as the system is designed to perform out of the box.

LUN Configuration of the FlashArray

One 40 TB thin provisioned volume was configured to host 500 desktops. Because the FlashArray doesn't have any requirement for configuring RAID groups/aggregates, it was a simple two-step task to configure Pure Storage volumes and provision to the server. The task of provisioning the volumes to vSphere cluster was further simplified by creating a host group on Pure Storage FlashArray that provided a one-to-one mapping with the vSphere cluster. Pure Storage provides private volumes that can be used for booting the server from SAN.

A common question when provisioning storage is how many LUNs of what size should be created to support the virtual desktop deployment. Because PVS write cache for each desktops take very little space, we could have either put all the virtual desktops in one big LUN or spread them across several LUNs, we chose to put all the write cache in one big datastore. The FlashArray supports the VMware VAAI ATS primitive which gives you access to multiple VMDKs on a single LUN (Note in vSphere 5.x the maximum size of a LUN is 64 TB). VAAI ATS eliminates serializing of VMFS locks on the LUN, which severely limited VM scalability in past ESX(i) versions. See Appendix A for more details on provisioning Pure Storage.

Now lets look at a sizing exercise, since we are advocating placing the OS image, user data, persona and application data on the same storage, we need to take into account the size of write cache when calculating the LUN size.

Consider a desktop with a 30 GB base image including applications and app data, 5 GB of user data (write cache) and we need to provision "d" desktops:

- We need to provision at least: base image size + (5 * d) GB as the size of the LUN

Regardless of the data reduction, we need to create the LUN with the correct size so that vSphere doesn't run out of storage capacity. Figure 3 below shows the virtual desktop deployment on Pure Storage.

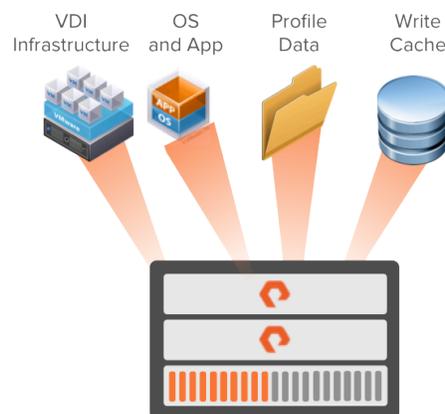


Figure 3: OS image, applications, user data and application data hosted on Pure Storage

Data Reduction with Pure Storage FlashArray

Storage Capacity with 1,000 PVS based Desktops

Figure 4 below shows 1,000 Windows 7 desktops deployed on an existing FlashArray with other workloads running on it. The total physical capacity used was less than 1 TB for the entire 1,000 desktops. Because we used PVS based desktops which stores only the differential data, we achieved an overall data reduction of 24.9-to-1.



Figure 4: Data reduction of 1,000 Windows 7 desktops plus other workloads in the system

In a real world scenario, the data reduction number is more in the order of 10-to-1 as the user data would differ more than in our example.

Unlike traditional storage arrays, we used a common LUN to store the OS image, user data, application data, and persona. We don't see any benefits in separating them on the FlashArray. We do not do data reduction on a volume basis; it is done across the entire array, which is reflected in the shared data in the capacity bar above.

The write cache LUN of 40 TB size was separately created and we saw a 9-to-1 data reduction on the write cache LUN as shown below in figure 5.

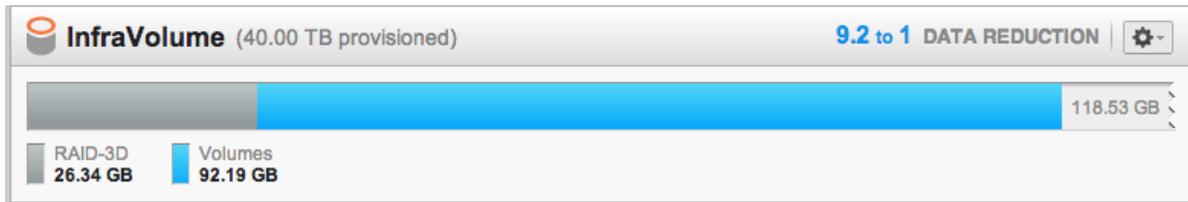
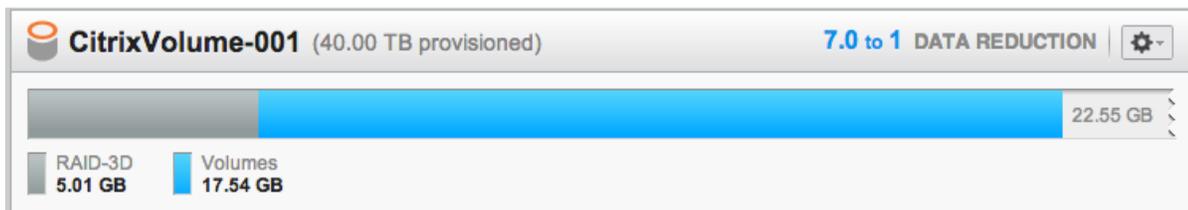


Figure 5: Pure Storage GUI snapshot of the Write Cache Volume (PVS Cache on Server)

Figure below shows Pure Storage GUI snapshot of the Pure LUN used for write cache when "cache on device hard disk" option was used in PVS. Note that the pagefile was automatically stored on the write cache and we saw a data reduction in excess of 7 X.



Server Configuration

Eight identical Intel CPU-based commodity servers were deployed for hosting the virtual desktops. The server's dual HBA ports were connected to two Cisco MDS 9148 SAN switches for upstream connectivity to access the Pure Storage FlashArray LUNs. The server configuration is described in Table B below.

Component	Description
Processor	2 X Intel Xeon X5690 @ 3.47GHz (12 Cores total, 24 Logical CPUs)
Memory	192 GB @ 1333 MHz (16 GB X 12)
HBA	Dual port Qlogic ISP2532-based 8 Gb Fibre Channel PCIe card
NIC	Quad port Intel Corp 82576 1 Gb card
BIOS	Intel Virtualization Tech, Intel AES-NI, Intel VT-D features were enabled
vSphere	ESXi 5.1.0, Build 799733

Table B: Desktop host server configuration

SAN Configuration

Figure 2 shows the SAN switch connectivity with two Cisco 8 Gb MDS 9148 Switch (48 ports). The key point to note is that there is no single point of failure in the configuration. The connectivity is highly resilient in terms of host initiator port or HBA failure, SAN switch failure, a controller port failure, or even array controller failure. The zoning on the Cisco MDS follows best practices i.e a single initiator and single target zoning. Check Appendix B for a sample Cisco MDS zoning. All eight ESXi host dual HBA port World Wide Names (pWWN) were zoned to see the four Pure Storage FlashArray target port World Wide Names. The target ports were picked such that, on a given controller, we had one port from each target Qlogic adapter connected to one switch, the other Qlogic adapter port was connected to a second switch (See Figure 2 for the wiring details). This resulted in ESXi 5.1 hosts to see 8 distinct paths to the Pure Storage FlashArray LUNs (Figure 6 shows vCenter datastore details). Notice the hardware acceleration support and path selection policy set to Round Robin.

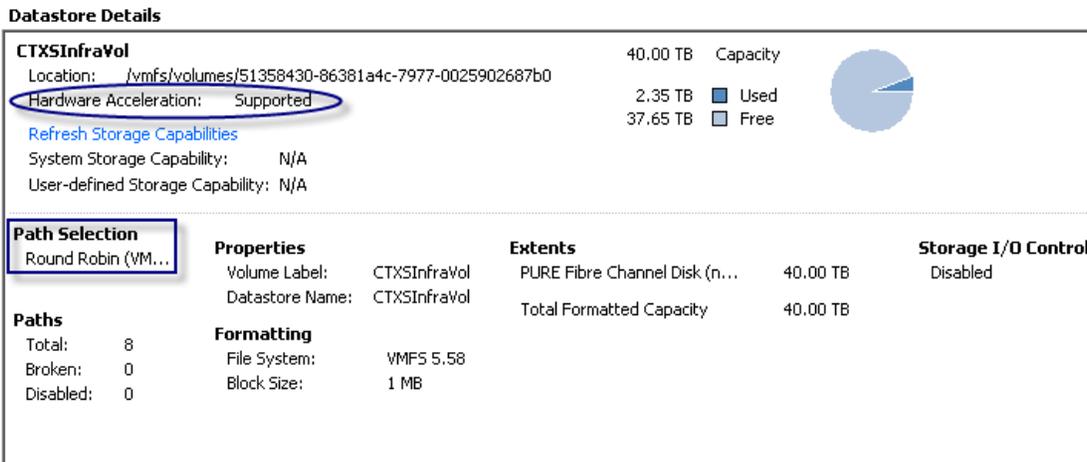


Figure 6: VMware VMFS datastore details

Network Configuration

Figure 7 below illustrates the network design used for the desktop deployment. A virtual machine was setup to run AD/DNS and DHCP services and we used a private domain. As large numbers of desktops were to be deployed, we wanted to setup our own private VLAN (VLAN 131) for desktops to hand out IP addresses to virtual desktops that were spun up. A separate VLAN (VLAN 124) was used for management network including the ESXi hosts on a single Cisco 3750 1Gb Ethernet switch (48 ports).

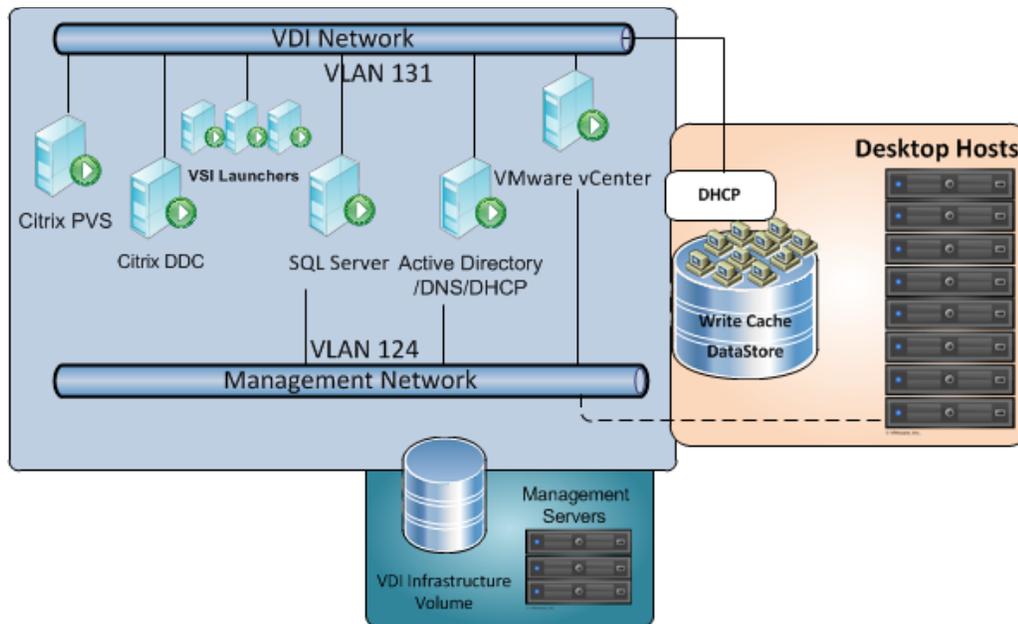


Figure 7: Logical view of the reference architecture showing network configuration

ESX Configuration and Tuning

ESXi 5.1.0, build 799733 was installed on eight identical servers and a separate infrastructure management server. This section talks about the storage, network, and general system configuration followed by specific tuning that was done to get the best performance. We started out with little or no tuning and have narrowed down to a small set of ESXi tuning configurations. Due to the large number of VMs and hosts VMware Management Assistant and vSphere powershell were used extensively and helped us get administrative tasks done efficiently.

Pure Storage FlashArray Best Practices for vSphere 5.1

The FlashArray is a VAAI-compliant, ALUA-based active/active array and doesn't require a special vSphere plugin to make the array work. The default storage array type plugin (SATP), VMW_SATP_ALUA, is automatically selected. However, the default path selection policy (PSP) is "Fixed" path. The PSP should be changed to Round Robin for all the Pure Storage LUNs as all paths to FlashArray is active optimized. This can be done using vCenter, Webclient or ESXi command line (See Appendix C for steps using vCenter). The following ESXi command accomplished this on a per device basis :

```
esxcli storage nmp device set -d naa.6006016055711d00cff95e65664ee011 --psp="VMW_PSP_RR"
```

We set all the Pure Storage LUNs to a Round Robin policy from vMA using the CLI command :

```
for i in `esxcli storage nmp device list | grep PURE|awk '{print $8}'|sed 's/(//g'|sed 's/)//g'`; do esxcli storage nmp device set -d $i --psp=VMW_PSP_RR ; done
```

For our tests, we set the default PSP for VMW_SATP_ALUA as VMW_PSP_RR and every Pure Storage LUN configured got the Round Robin policy. The following command accomplished that:

```
esxcli storage nmp satp set --default-psp="VMW_PSP_RR" --satp="VMW_SATP_ALUA"
```

Figure 8 shows a properly configured Pure Storage LUN with VMware Round Robin PSP.

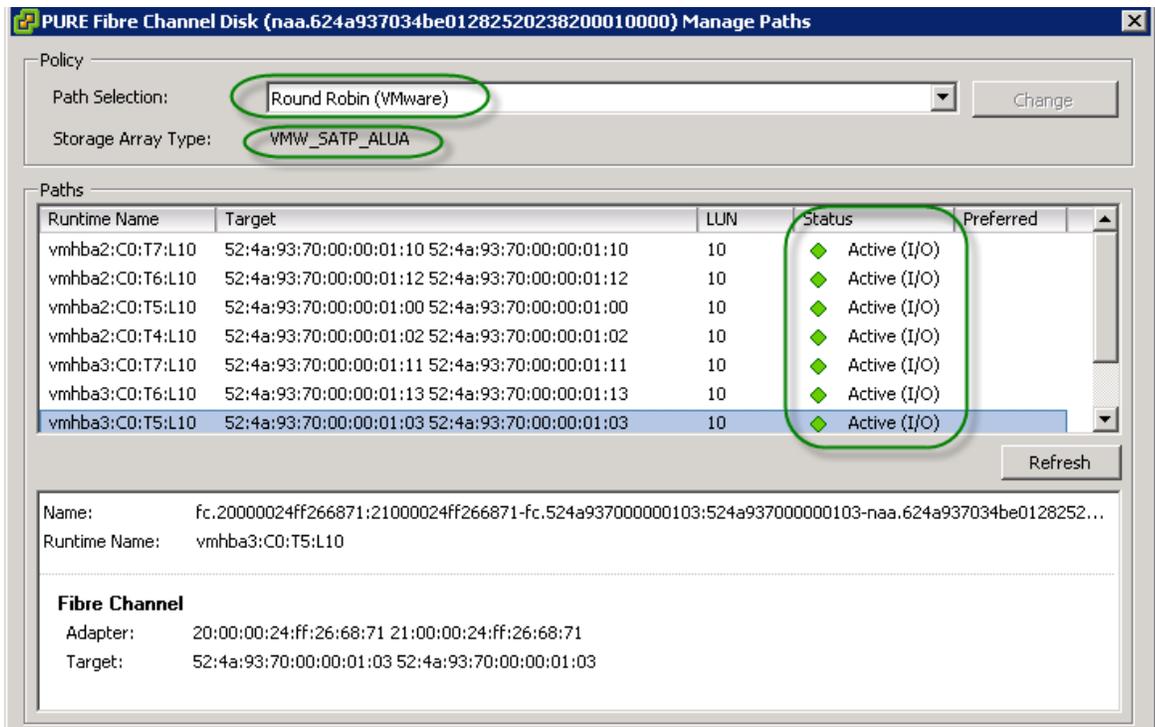


Figure 8: Pure Storage LUN configured with Round Robin path policy

ESXi 5.1 Configuration and Tuning

In this section, we discuss the ESXi 5.1 cluster configuration, network configuration and ESXi tuning for the disk subsystem.

ESXi Cluster Configuration

A datacenter and a cluster with eight hosts were configured with VMware’s High Availability clustering (HA) and Distributed Resource Scheduling (DRS) features. DRS was set to be fully automatic so that the 1,000 desktops would be evenly distributed across the eight hosts. The DRS power management was turned off and the host EVC policy was set to “Intel Westmere.” The BIOS of each host was examined to make sure the Intel VT-d was on, AES-NI instructions were enabled. The HA configuration was setup with VM restart priority as high and isolation policy set to “leave powered on.” Finally, the swap file was stored along with the VM. A resource pool was created to house the desktops with default settings. Due to the one-to-one mapping of the ESX hosts in a cluster to the Pure Storage host group and hosts, all hosts saw all the LUNs except for the private volumes used by each host for boot.

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/active mode and VLAN configurations were done on the upstream Cisco 3750 1 GE switch. The 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 9 and Figure 10 respectively.

Networking

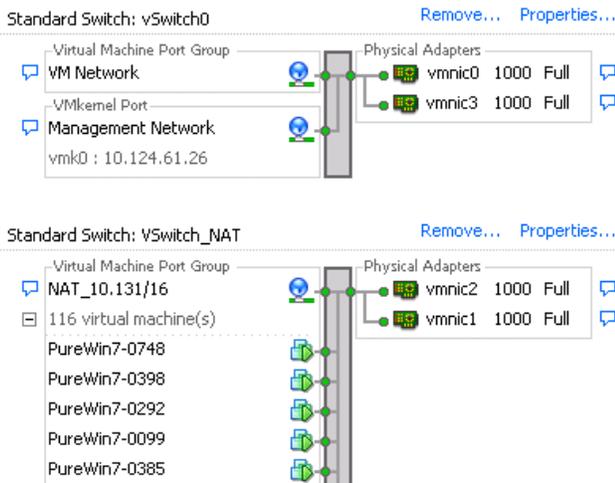


Figure 9: VMware virtual switch configuration

The default 128 ports on a virtual switch were changed to 248, as there was a potential to put more desktops on a single host (host reboot is mandated for this change). The MTU was left at 1500.

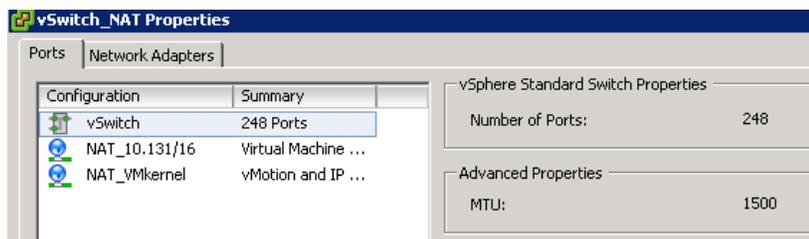


Figure 10: Virtual switch properties showing 248 ports

vSphere System Level Tuning

In order to get the best performance out of the FlashArray some of the default disk parameters in vSphere had to be changed, because the default values are applicable to spindle based arrays that are commonly deployed in the data center. The two disk parameters that were changed to a higher value for this work are Disk.SchedNumReqOutstanding (default value of 32) and Disk.SchedQuantam (default value of 8) to 256 and 64 respectively. The former, DSNRO, will limit the amount of I/Os that will be issued to the LUN. This parameter was tuned to a maximum limit so the FlashArray can service more I/Os. The best treatise on this topic can be found in [See reference 4]. The latter, Disk.SchedQuantum value determines the number of concurrent I/Os that can be serviced from each world (a world is equivalent to a process in VMkernel terms), so we set that value to its maximum value of 64. Figure 11 below shows how to set it using a vCenter on a per host basis.

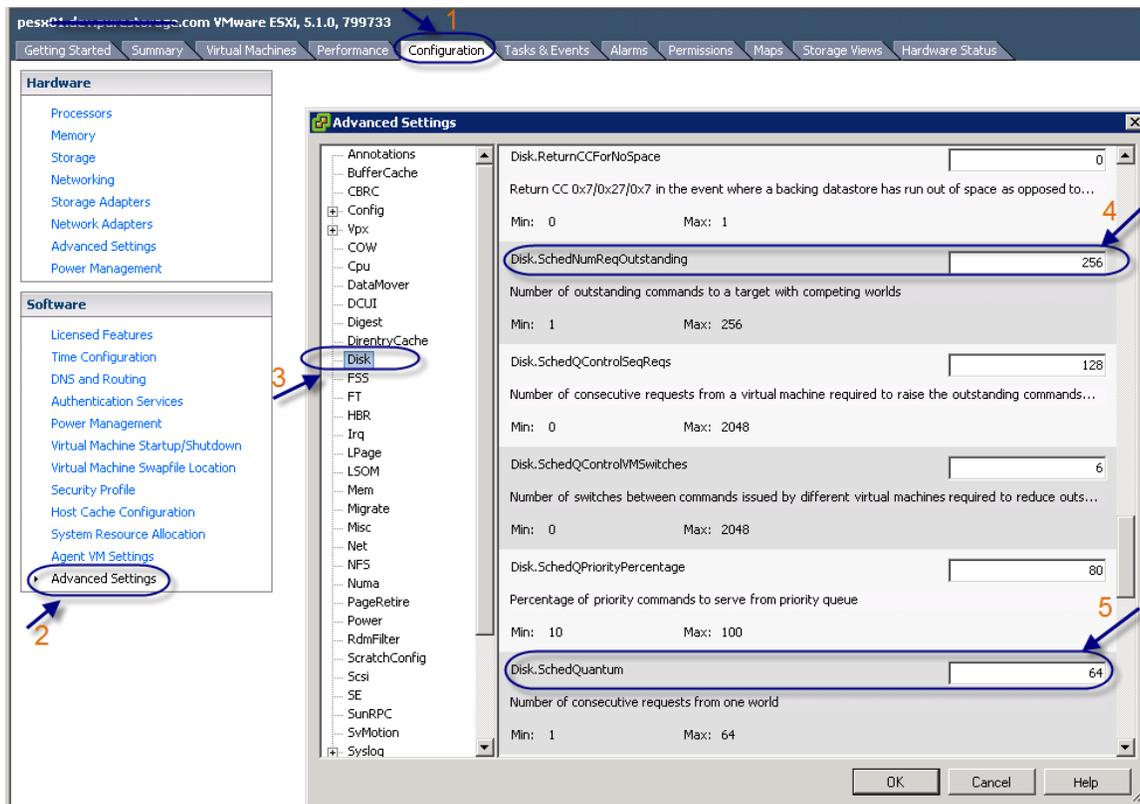


Figure 11: vCenter snapshot of setting Disk tunables

The same can be accomplished by using a vMA appliance and the command line (a script was used to configure these settings):

```
esxcfg-advcfg --set 256 /Disk/SchedNumReqOutstanding
```

```
esxcfg-advcfg --set 64 /Disk/SchedQuantum
```

The qllogic HBA max queue depth was increased to 64 from its default value on all hosts (see VMware KB article [1267](#) for setting this value):

```
# esxcfg-module qla2xxx -g ql2xmaxqdepth

qla2xxx enabled = 1 options = 'ql2xmaxqdepth=64'
```

There was no other tuning that was done to the vSphere server. Note, these settings are global settings on an ESXi server and are recommended if there are no other storage arrays connected to the server.

Management and Infrastructure Virtual Machines

In order to scale the environment it is important to have robust and efficient infrastructure components. As per the design principles, we built the management and infrastructure VMs on a separate management host server running ESXi 5.1.0. Why? Because as we scale this environment, we expect Active directory and DNS/DHCP will have to scale to give the best possible user experience, so we will need dedicated host resources to support that growth.

We created a master Microsoft Windows 2008 R2 template with all the updates and cloned the different infrastructure VMs. The Citrix Provisioning Services (PVS) 6.1 and Citrix Desktop Delivery Controller (DDC) were installed with 10G vmxnet3 network driver and was given 4 vCPU and 8 GB of RAM. The SQL server VM hosted the Microsoft SQL Server 2008 R2 database instance of the vCenter database and the Citrix PVS database. The 15 Login VSI launchers VM were Windows 7 VMs and as per Login VSI recommendations were given 2 vCPU and 8 GB of RAM. The description of each of the infrastructure components is shown in Figure 12. Based on the experiments done, our recommendation is to have a single PVS and DDC VM to deploy 500 desktops. If you want to scale, one has to add more PVS and DDC VMs and build out a PVS and DDC farm.

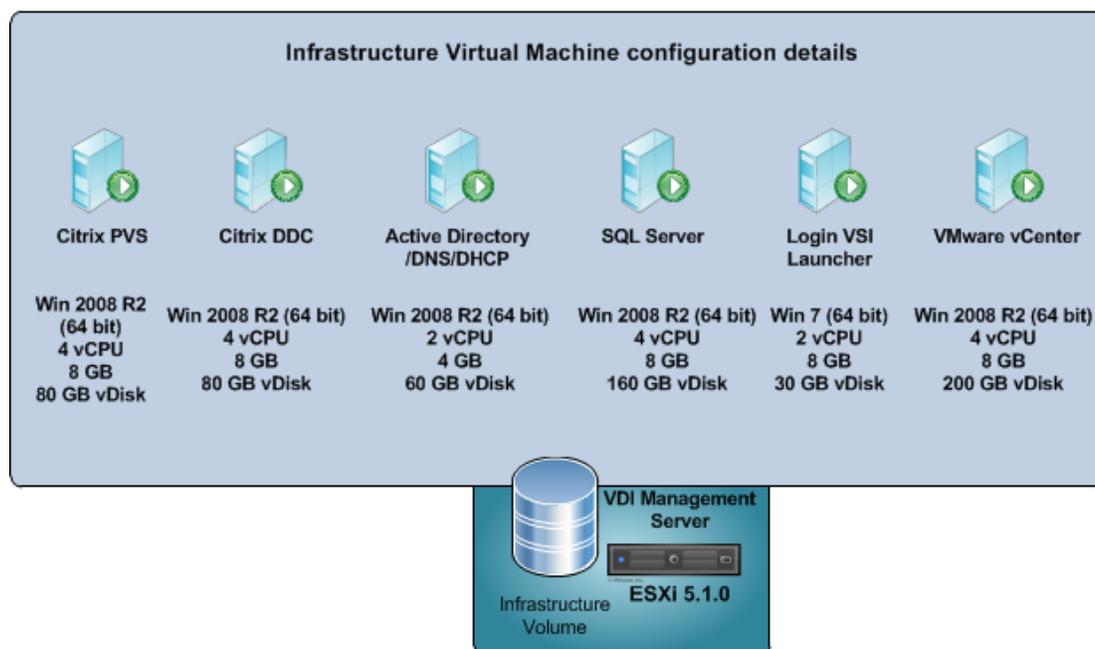


Figure 12: Infrastructure Virtual Machine component detailed description

The management infrastructure host used for the infrastructure VMs was provisioned with a 600 GB LUN; the server configuration is shown in Table C below.

Component	Description
Processor	2 X Intel Xeon E5645 @ 2.40 GHz (12 Cores total, 24 Logical CPUs)
Memory	96GB @ 1333 MHz (8 GB X 12)
HBA	Dual port Qlogic ISP2532-based 8 Gb Fibre Channel PCIe card
NIC	Quad port Intel Corp 82576 1 Gb card
BIOS	AES, Intel VT-D features were enabled
vSphere	ESXi 5.1.0, Build 799733

Table C: Management infrastructure host configuration details

Desktop Software Configuration

Citrix XenDesktop 5.6 Overview

Citrix XenDesktop is a desktop virtualization system that centralizes and delivers Microsoft Windows XP, 7 or 8 virtual desktops to users anywhere. Virtual desktops are dynamically assembled on demand, providing users with pristine, yet personalized, desktops each time they log on. This ensures that performance never degrades, while the high speed delivery protocol provides unparalleled responsiveness over any network. Citrix HDX delivers a high definition user experience for virtual desktops, and offers the best performance over any network connection. The open architecture of XenDesktop offers choice and flexibility of virtualization platform and user device. Unlike other desktop virtualization alternatives, XenDesktop simplifies desktop management by using a single image to deliver personalized desktops to users and enables administrators to manage service levels with built-in desktop performance monitoring. Citrix FlexCast delivery technology enables you to deliver every type of virtual desktop, hosted or local, physical or virtual, each specifically tailored to meet the performance, security and flexibility requirements of each individual user.

Figure 13 below shows remote users working at their local user devices while accessing virtual desktops provided by XenDesktop.

Although the desktops are virtual, running on remote servers, the user experience is equivalent to that of a local Windows desktop. From the user's perspective, logging on to a virtual desktop is the same as logging on to a local desktop. Users enter their credentials once and are connected to their desktops.

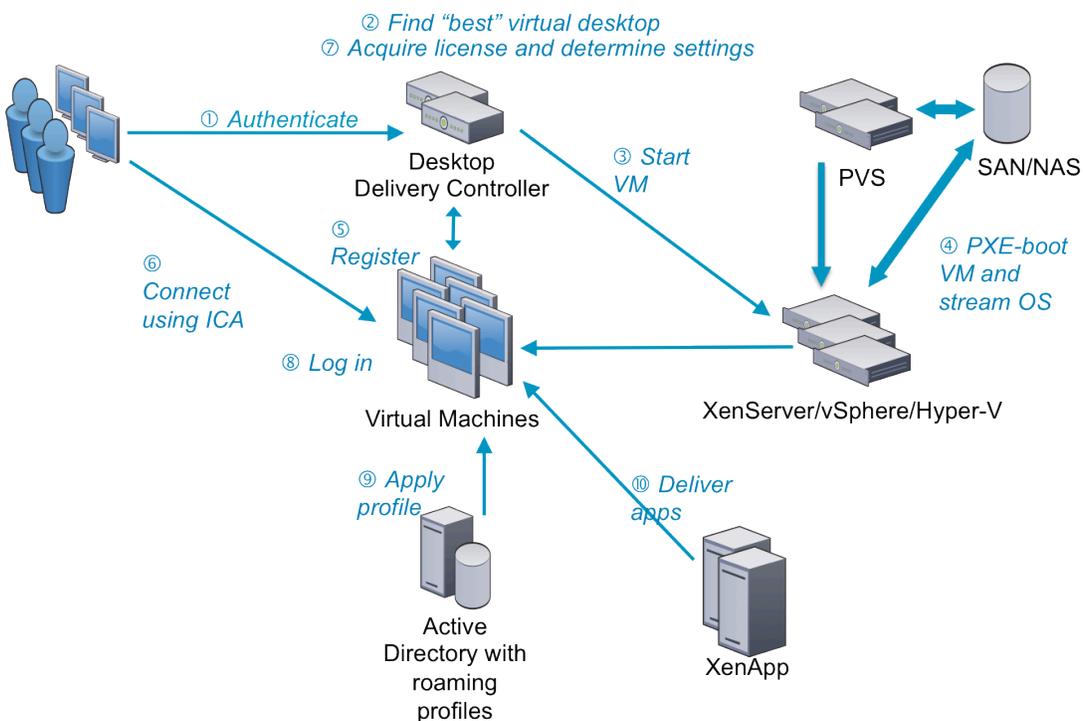


Figure 13: Citrix XenDesktop Architecture Overview

Components of Citrix XenDesktop Deployment

Citrix XenDesktop provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.

The core components of XenDesktop are:

- **Desktop Delivery Controller:** Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops. It controls the state of the desktops, starting and stopping them based on demand and administrative configuration. Desktop Delivery Controller also includes Profile Management, in some editions, to manage user personalization settings in virtualized or physical Windows environments.
- **Virtual Desktop Provisioning powered by Citrix Provisioning Services:** Provisioning Services creates and provisions virtual desktops from a single desktop image on demand, optimizing storage utilization and providing a pristine virtual desktop to each user every time they log on. Desktop provisioning also simplifies desktop images, provides the best flexibility, and offers fewer points of desktop management for both applications and desktops.
- **Virtual Desktop Agent:** Installed on virtual desktops, the agent enables direct ICA (Independent Computing Architecture) connections between the virtual desktop and user devices.
- **Citrix online plug-in:** Installed on user devices, the Citrix online plug-in (formerly "Citrix Desktop Receiver") enables direct ICA connections from user devices to virtual desktops.
- **Citrix XenApp:** You can use XenApp to benefit from the efficiencies associated with application streaming and virtualization. XenApp provides a better-than-installed application experience for both users and administrators. Applications start up faster, the user experience is dramatically improved, and application management costs are significantly lowered.

Additional XenDesktop components provide the following features:

- **Secure delivery:** When users connect from outside the corporate firewall, XenDesktop can use Citrix Access Gateway technology to secure these connections with SSL. Access Gateway is a SSL VPN appliance that is deployed in the demilitarized zone (DMZ) to provide a single secure point of access through the corporate firewall.
- **WAN optimization:** In XenDesktop deployments where virtual desktops are delivered to users at remote locations such as branch offices, Citrix Branch Repeater (formerly WANScaler) technology can be employed to optimize performance through Quality of Service management. Branch Repeater can prioritize different parts of the user experience so that, for example, the user experience does not degrade in the branch location when a large file or print job is sent over the network. HDX IntelliCache with Branch Repeater provides tokenized compression and data de-duplication, dramatically reducing bandwidth requirements and improving performance.
- **Monitoring:** Citrix EdgeSight for Virtual Desktops allows you to monitor individual virtual desktops. EdgeSight can be used not only to analyze and troubleshoot issues, but also to warn administrators in advance of problems that may arise in the future.

- **Password Manager:** Password Manager provides single sign-on access regardless of how or where users connect, and it enables users to reset their own Windows password or unlock their account. Note that Password Manager is not fully supported on Windows 7 and the Hot Desktop feature, which allows users to share workstations securely and efficiently, is not supported in XenDesktop.

For more complete information please see reference [8] for Citrix documentation on XenDesktop 5.6.

Citrix Provisioning Services Configuration

Provisioning Services creates and provisions virtual desktops from a single desktop image on demand, optimizing storage utilization and providing a pristine virtual desktop to each user every time they log on. Desktop provisioning also simplifies desktop images, provides the best flexibility, and offers fewer points of desktop management for both applications and desktops.

The target device downloads the boot file from a Provisioning Server and then the target device boots. Based on the device boot configuration settings, the appropriate vDisk is located, then mounted on the Provisioning Server. The software on that vDisk is streamed to the target device as needed. To the target device, it appears like a regular hard drive to the system.

Instead of immediately pulling all the vDisk contents down to the target device (as done with traditional or imaging deployment solutions), the data is brought across the network in real-time, as needed. This approach allows a target device to get a completely new operating system and set of software in the time it takes to reboot, without requiring a visit to a workstation. This approach dramatically decreases the amount of network bandwidth required by traditional disk imaging tools; making it possible to support a larger number of target devices on your network without impacting overall network performance.

In order to make the DHCP point to the PVS server the following changes were made to the scope options 066 and 067 :

 066 Boot Server Host Name	Standard	CitrixPVS.vdi.puredemo.com	None
 067 Bootfile Name	Standard	ARDBP32.BIN	None

In our setup, Citrix Provisioning Services (PVS) 6.1 was installed on a Windows 2008 R2 VM with 4 vCPU and 8 GB of memory. We used a SQL server to manage the database as we were trying to deploy a large number of desktops instead of SQL Express.

One of the most important decisions Citrix XenDesktop admins have to make is the placement of the write cache. There are varieties of options to choose from, including:

- Cache on device hard drive
- Cache on device hard drive Persisted (Experimental)
- Cache in device RAM
- Cache on a server disk
- Cache on server persisted

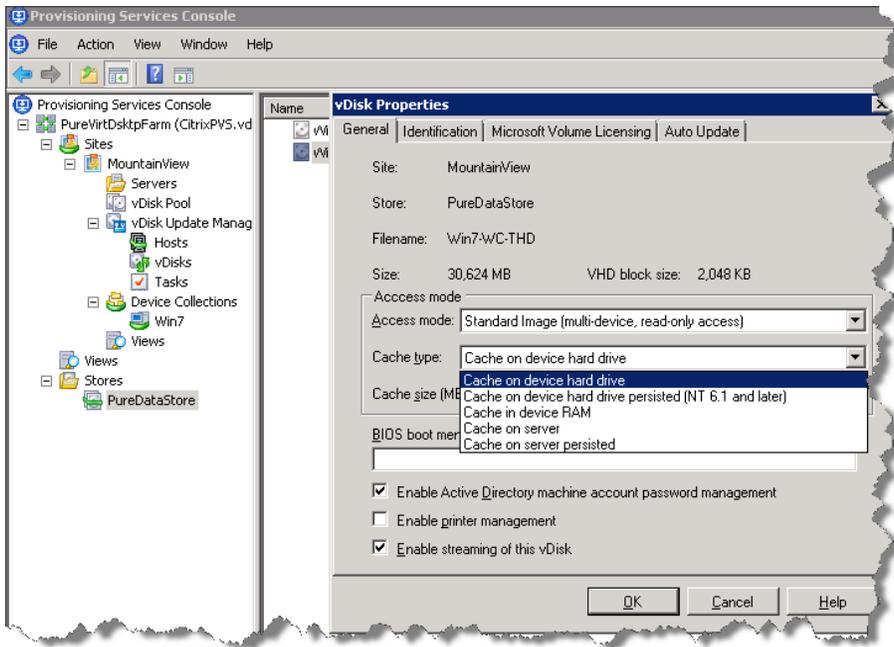


Figure 14: Citrix Provisioning Services 6.1 options for Cache type

There are extensive blogs and Citrix knowledge base articles which talk about how to size the write cache and describe the pros and cons of each approach (See Reference [5] and [6] for more detailed information on write cache sizing and placement considerations).

Storage Provisioning for Citrix PVS on Pure Storage

In this paper we chose two options of PVS write caching:

- Cache on device hard drive
- Cache on a server disk

We ran the Login VSI Pro 3.7 medium workload on both setups.

One of the advantages of using the Pure Storage FlashArray in a virtual environment is that it takes out commonly seen I/O bottleneck issues regardless of whether the Pure LUN is exposed as a RDM or a VMFS based datastore (reference [2] talks more on this topic). As is apparent from the results section, contrary to the beliefs and recommendations, the cache on server disk option works quite well and we got a very good Login VSI result.

- One 40TB LUN was configured for hosting the write cache when we chose the “Cache on device hard disk” option. The pagefile will be automatically placed on this write cache when the Citrix target driver is installed.
- A 2 TB volume was carved out on a Pure Storage LUN for the write cache (note the disk has to be MBR and 2 TB is the max size) when the “cache on server” option was used.
- A 600 GB LUN was carved out to host the infrastructure VMs including the Login VSI launcher.

The only tweak done to PVS was to boot more VMs and shutdown more desktops in the server-advanced configuration.

For DDC, there was no special tweaking expect we turned off the shutdown after logout feature using the PowerCLI command :

```
> Set-BrokerDesktopGroup -Name "Pure Windows 7 Login VSI group" -ShutdownDesktopsAfterUse $False
```

Desktop Operating System - Microsoft Windows 7 Configuration

Table D describes the configuration of the desktops.

Component	Description
Desktop	Windows 7 Professional 64-bit
Hardware version	8
vCPU	1
Memory	2 GB
vNIC adapter	1 X vmxnet 3 adapter
vStorage adapter	1 X LSI Logic SAS adapter
Virtual disks	30 GB base image + 5 GB Write Cache
XenDesktop	Virtual Desktop Agent 5.6
Installed Application	MS Office 2010, Adobe Reader 9, Internet Explorer 9

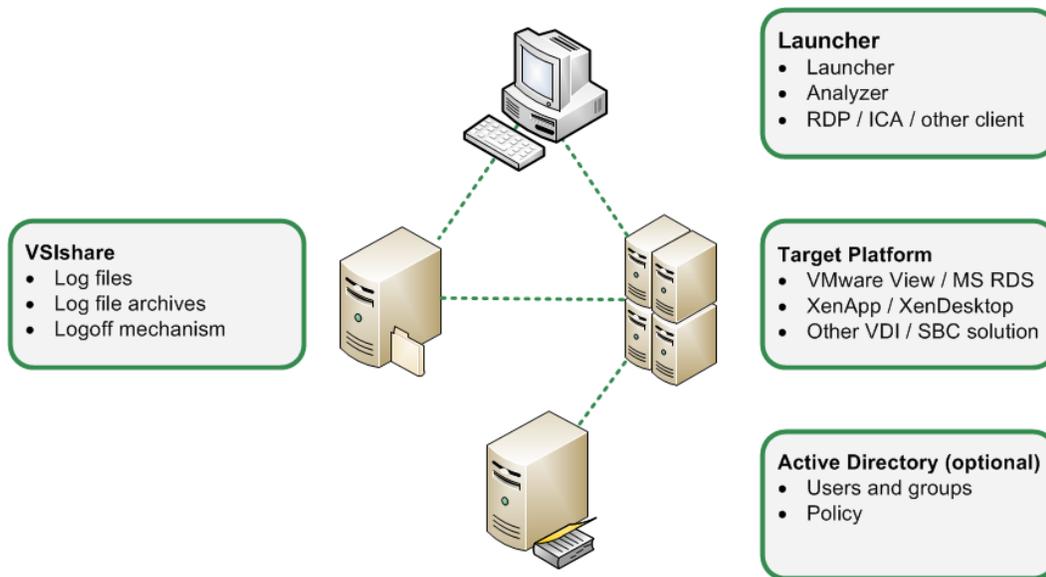
Table D: Windows 7 virtual desktop configuration summary

In addition to the XenDesktop VDA we installed the Login VSI target device components and the applications needed for the test runs.

Software Testing Tool – Login VSI Pro 3.7

Login VSI is 100% vendor independent and is used to test virtual desktop environments like Citrix XenDesktop and XenApp, Microsoft VDI and Remote Desktop Services, VMware View or any other Windows based SBC or VDI solution.

Login VSI is used for testing and benchmarking by all major hardware and software vendors, and is recommended by both leading IT-analysts and the technical community. Login VSI is vendor independent and works with standardized user workloads, therefore conclusions that are based on Login VSI test data are objective, verifiable, and replicable. This is the reason why Pure Storage used Login VSI for our testing.



Source: <http://www.loginvsi.com/documentation/v3>

Figure 15: Architecture Overview of Login VSI Test Infrastructure

Figure 15 above shows the architecture of the Login VSI test infrastructure. It is comprised of four components :

- Active Directory (AD) for user accounts and standard policies.
- A file share for central configuration and logging (VSIShare).
- Launcher workstations (master and slaves) to initiate the sessions.
- Target Platform (VDI) where the workload scripts are installed and performed.

The launcher VMs have Citrix receiver installed on them and connect to the target desktops to initiate a client session and start the load. A maximum of 35 sessions is supported by each launcher VM, so based on the number of sessions one has to spin up as many launchers to run the workload. Login VSI management console can be run on one of the launchers and it acts as a primary control point to coordinate all other launcher session launching activities. All operations are done via the console, including configuring and starting the workload based on configuration parameters.

There are many workload settings but we chose to run the “medium” workload to simulate knowledge worker desktops.

Medium Workload

The medium workload is the only workload available in [Login VSI Express](#) and also available in Login VSI Pro. The medium workload is the default workload in Login VSI.

- This workload simulates a medium knowledge working using Office, IE and PDF.
- Once a session has been started the medium workload will repeat every 12 minutes.
- During each loop the response time is measured every 2 minutes.
- The medium workload opens up to 5 apps simultaneously.
- The type rate is 160ms for each character.
- Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007/2010, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com and Lonelyplanet.com.
- Word 2007/2010, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007/2010, a very large randomized sheet is opened.
- PowerPoint 2007/2010, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

Figure 16 below shows the execution sequence of each application.

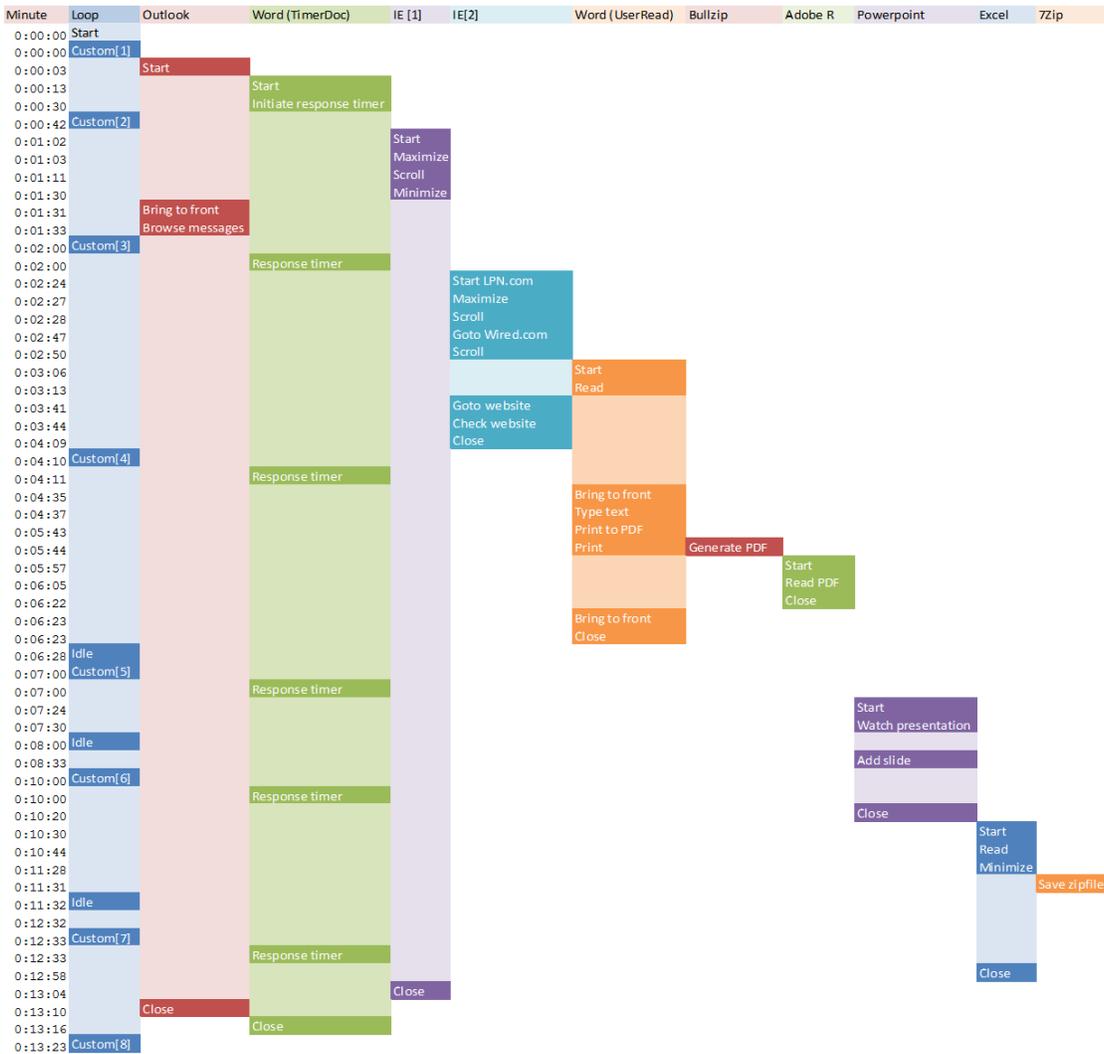


Figure 16: Login VSI Medium Workload Workflow

More information on setting up Login VSI and example videos can be found at - <http://www.loginvsi.com/documentation/v3>

Success Criteria and VSImax

Typically, the desktop workload is scripted in a 12-14 minute loop when a simulated Login VSI user is logged on. After the loop is finished it will restart automatically. Within each loop the response times of seven specific operations are measured in a regular interval: six times within each loop. The response times of these seven operations are used to establish **VSImax**.

The seven operations from which the response times are measured are:

- Copy new doc from the document pool in the home drive.
- Starting Microsoft Word with a document.

- Starting the “File Open” dialogue.
- Starting “Notepad.”
- Starting the “Print” dialogue.
- Starting the “Search and Replace” dialogue.
- Compress the document into a zip file with 7-zip command line.

These measured operations with Login VSI do hit considerably different subsystems such as CPU (user and kernel), Memory, Disk, the OS in general, the application itself, print, GDI, etc. These operations are specifically short by nature. When such operations are consistently long: the system is saturated because of excessive queuing on any kind of resource. As a result, the average response times will then escalate. This effect is clearly visible to end-users. When such operations consistently consume multiple seconds the user will regard the system as slow and unresponsive (See <http://www.loginvsi.com/documentation/v3/analyzing-results/calculating-vsimax> for more detailed information on VSImax calculations). As recommended we used the VSImax dynamic scores in our tests. Refer to Appendix D for a description of test configuration file used for simulating 500 users in our testing.

Test Results

Login VSI Pro 3.7 medium workload was run on 500 desktops with 15 Login VSI launchers, with each launcher (parallel mode) starting a session every 45 seconds. The launcher test configuration file can be found in Appendix D. After a successful test run we collected the Login VSI VSImax data and it is shown in figure 16 below. 99.9 % of the desktops were under 1.5 second mark, well within the 4000 milliseconds of VSImax value in both cases.

In summary we achieved:

- A score of 720 for Cache on device Hard Disk option
- A score of 738 for Cache on Server option

The DDC, PVS perfmon logs (see Appendix E, F, G, H for various perfmon statistics) were collected during the run as well as esxtop data on one of the ESXi server in the environment (the desktops were equally balanced on all hosts). In the write cache on server case, we saw a slightly higher CPU utilization on the PVS VM (which had 4 vCPUs, 8 GB memory).

PVS Write Cache on Server Results

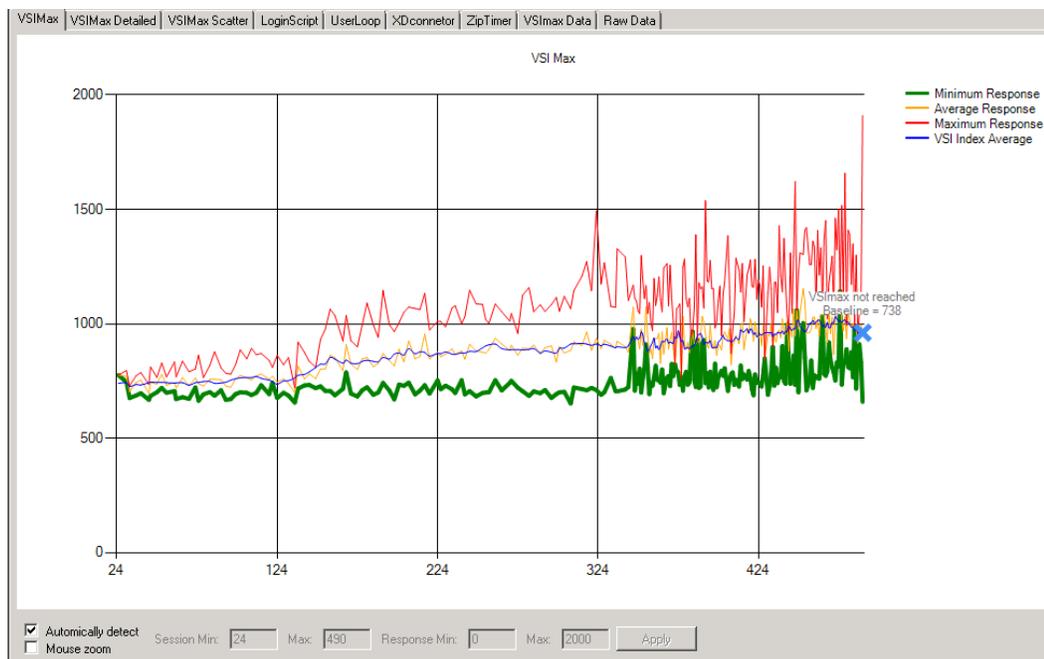


Figure 17: Login VSI VSImax score with 500 Citrix desktops (99% of desktops are < 1000 msec)

Figure 17 above shows the VSImax baseline at 738 with no hung sessions and the write cache on PVS server. Figure 18 below shows the corresponding Pure Storage GUI dashboard during the test run. Notice the latency during the entire duration of the tests was within 1 millisecond. The maximum CPU utilization on the servers was 78% and the memory used was 70% as shown in Figure 17 and 18 respectively. We saw no ballooning and no swapping as we had 192 GB of memory on each virtual desktop server host.

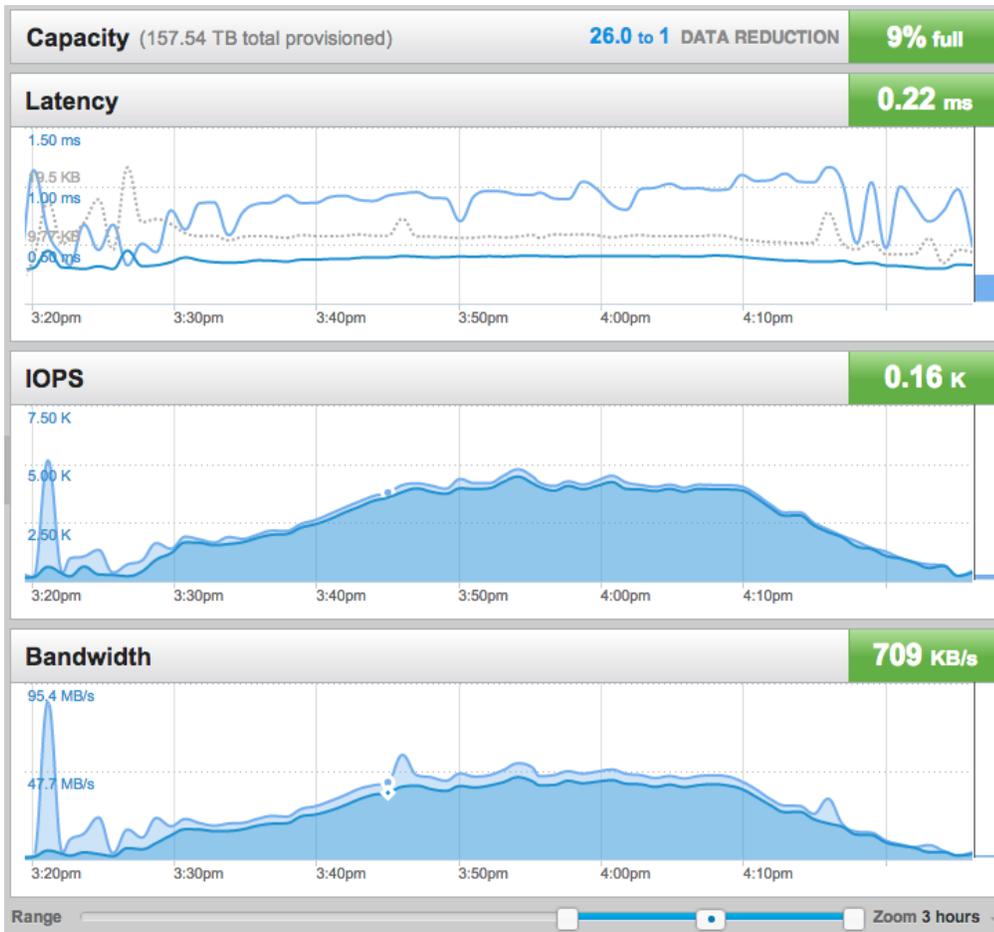


Figure 18: Pure Storage Dashboard view during Login VSI test run

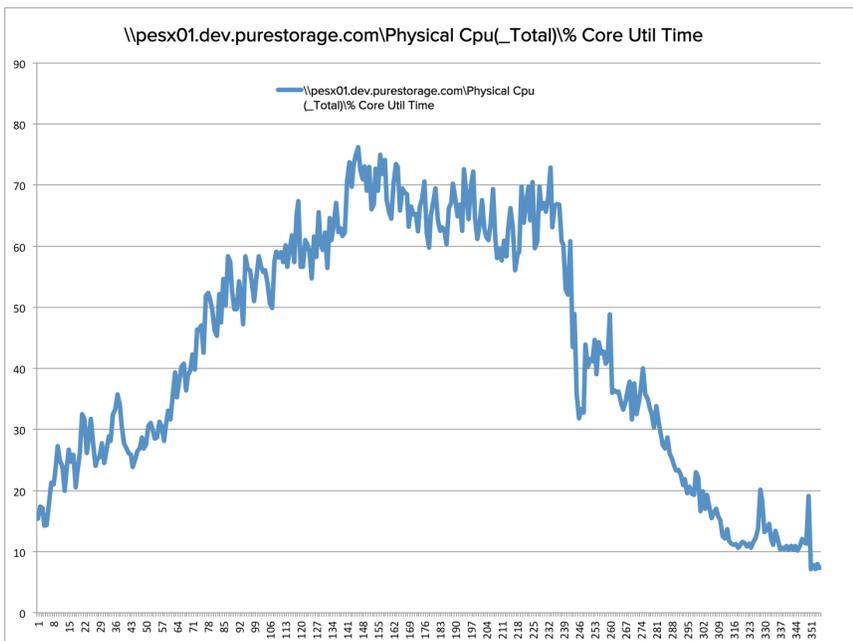


Figure 19: ESXi host performance data of a single host: CPU utilization

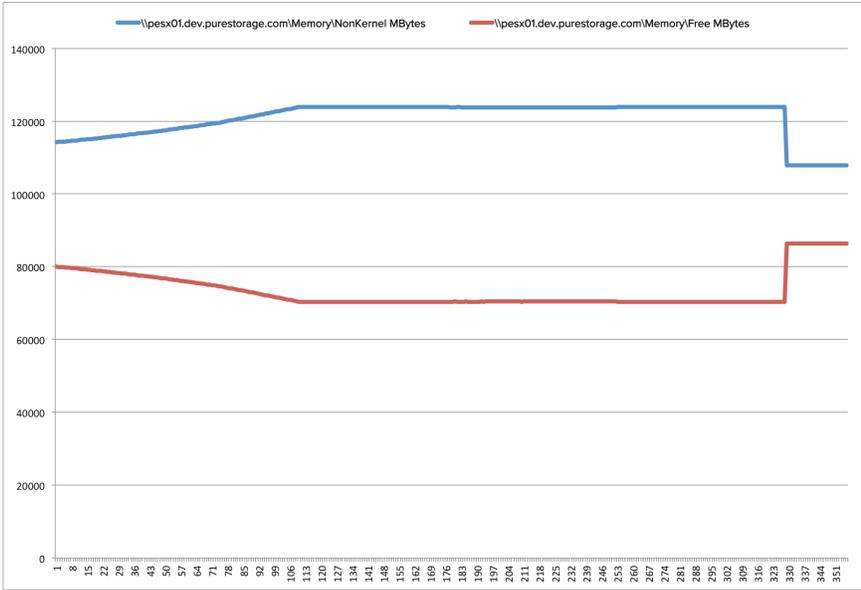


Figure 20: ESXi performance data of a single host: memory utilization (192 GB total memory)

It is important that we characterize the load on the Citrix Provisioning Services and Citrix Desktop Delivery Controller. We collected perfmon data on both the PVS and DDC VMs and it is shown in Appendix E and F respectively. The key point is with 4vCPU and 8 GB of memory, the PVS VM was able to host the write cache of 500 target devices while sustaining 50% max CPU and 40% max memory utilization.

PVS Write Cache on Hard Drive Results

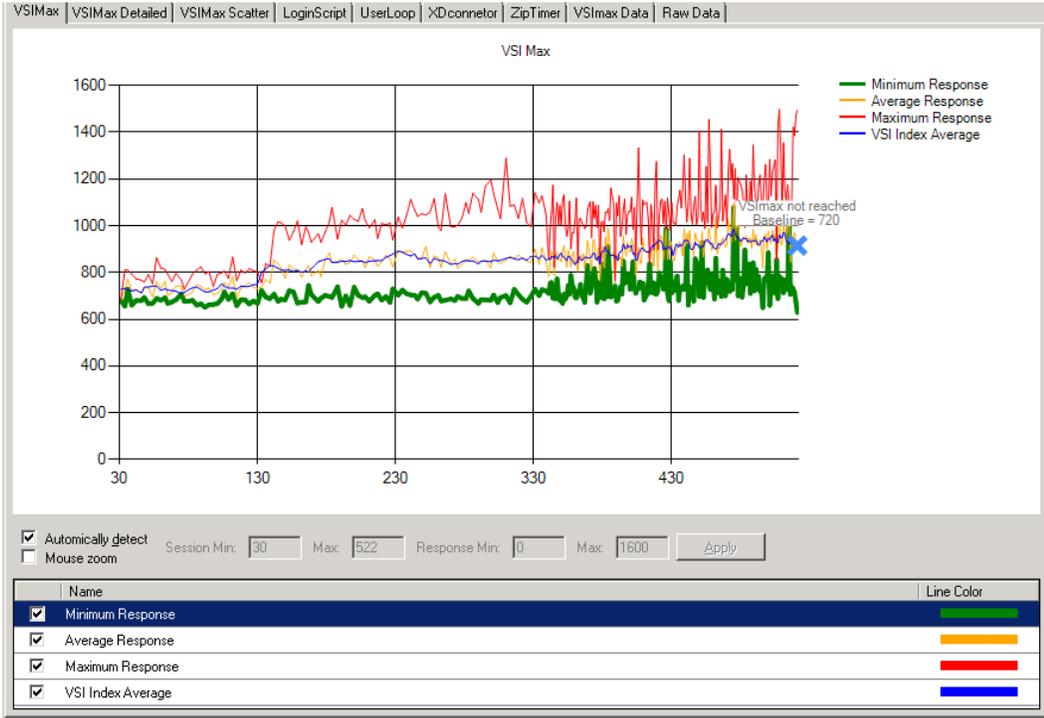


Figure 21: Login VSI VSImax score with 500 Citrix desktops (99% of desktops are < 1500 msec)

Figure 21 shows the VSImax value not reached at 720 with no hung sessions and all 500+ desktops running. Figure 22 below shows the corresponding Pure Storage GUI dashboard during the test run. Notice the latency during the entire duration of the tests was within 1 millisecond. The maximum CPU utilization on the servers was 64% and the memory used was 62% as shown in Figure 23 and 24 respectively.

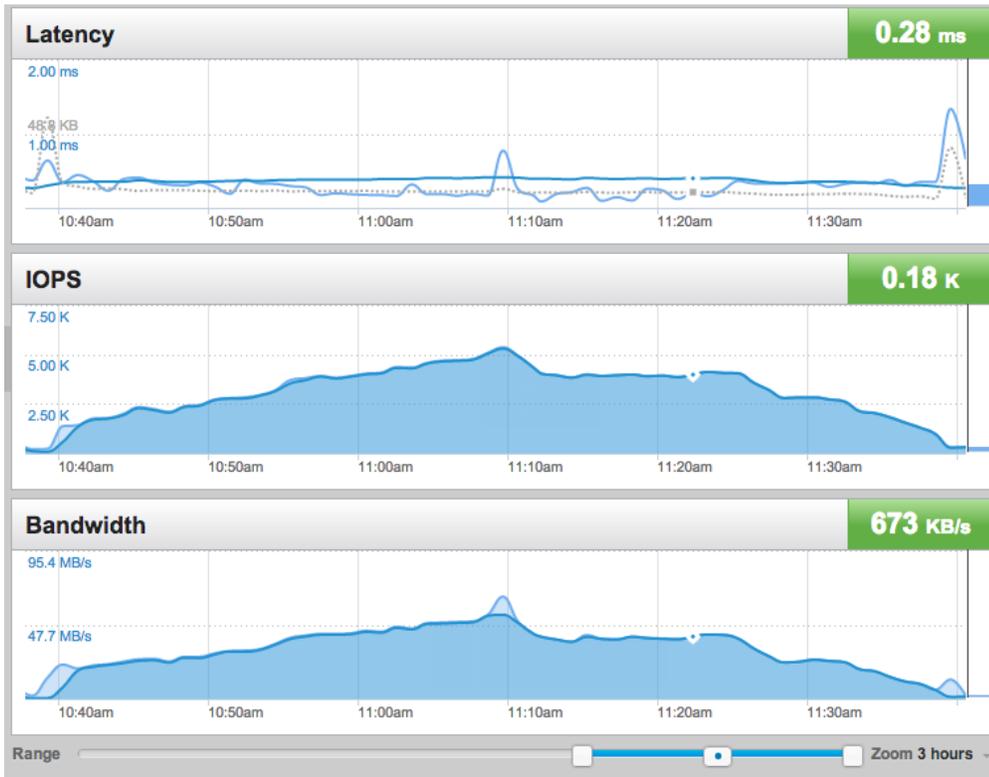


Figure 22: Pure Storage Dashboard view during Login VSI test run

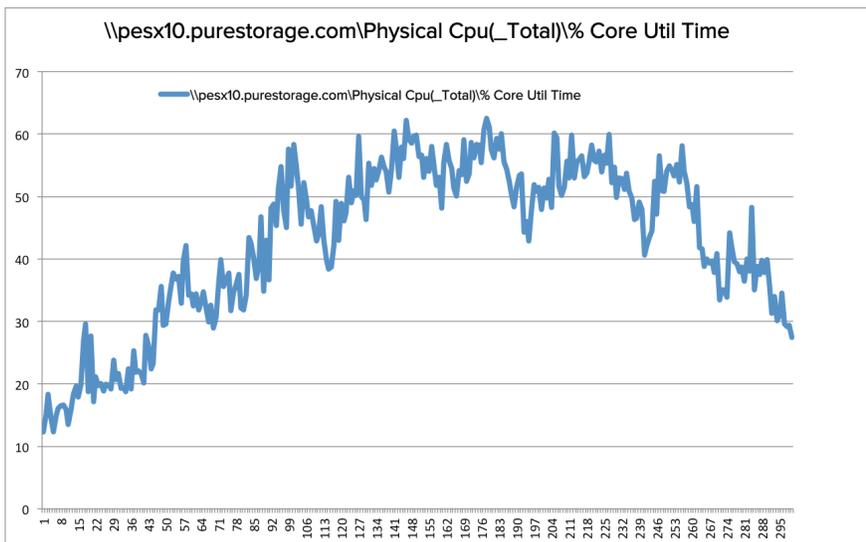


Figure 23: ESXi host performance data of a single host: CPU utilization

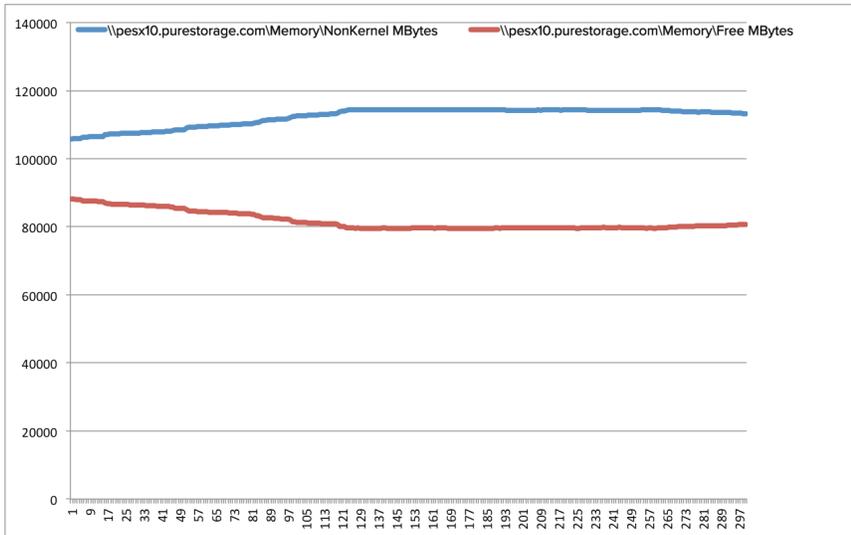


Figure 24: ESXi host performance data of a single host: Memory utilization

SUMMARY OF RESULTS

- Login VSI showed all the desktops well within the 1000 msec mark with average VSI_{max} score of 720
- 500 PVS based desktops created in less than one hour
- Deploying desktops, updating PVS desktop images including merging happens extremely fast
- 500 desktops booted in 30 minutes and sustained sub-1msec latency
- Data reduction (deduplication and compression) in excess of 10-to-1 across all desktops

Benefits of Pure Storage FlashArray for Citrix Deployment

The Pure Storage FlashArray fundamentally changes the storage experience from a virtualized data center perspective [see Reference 2 for more information]. Virtual desktop deployment with Citrix XenDesktop is a great use case for the FlashArray, due to its very low latency and efficient use of data reduction to increase storage efficiency. We have demonstrated the high throughput and low latency aspect resulting in sub second application response times with Login VSI Pro medium workload profile. Furthermore the data reduction on the write cache will help customers to scale to 1000s of desktops without worrying about storage capacity. Great performance and a higher IOPS coupled with remarkable space reduction makes the Pure Storage FlashArray an ideal storage for Citrix XenDesktop deployment.

Ease of Storage Capacity Management and Storage Provisioning

The Pure Storage GUI and CLI were designed with one goal in mind: simplify storage management. The GUI dashboard is a central location for understanding capacity, IOPS/Latency/Bandwidth and system status (see Figure 18). Most commonly used operations are simple two step operations, including creation of a LUN, increasing LUN size, or masking a LUN to a host. The LUN connection map will show if the host FC ports are zoned correctly and has redundant connectivity to the FlashArray (see Figure 25 below). This feature is included for all platforms with no agents in the OS.

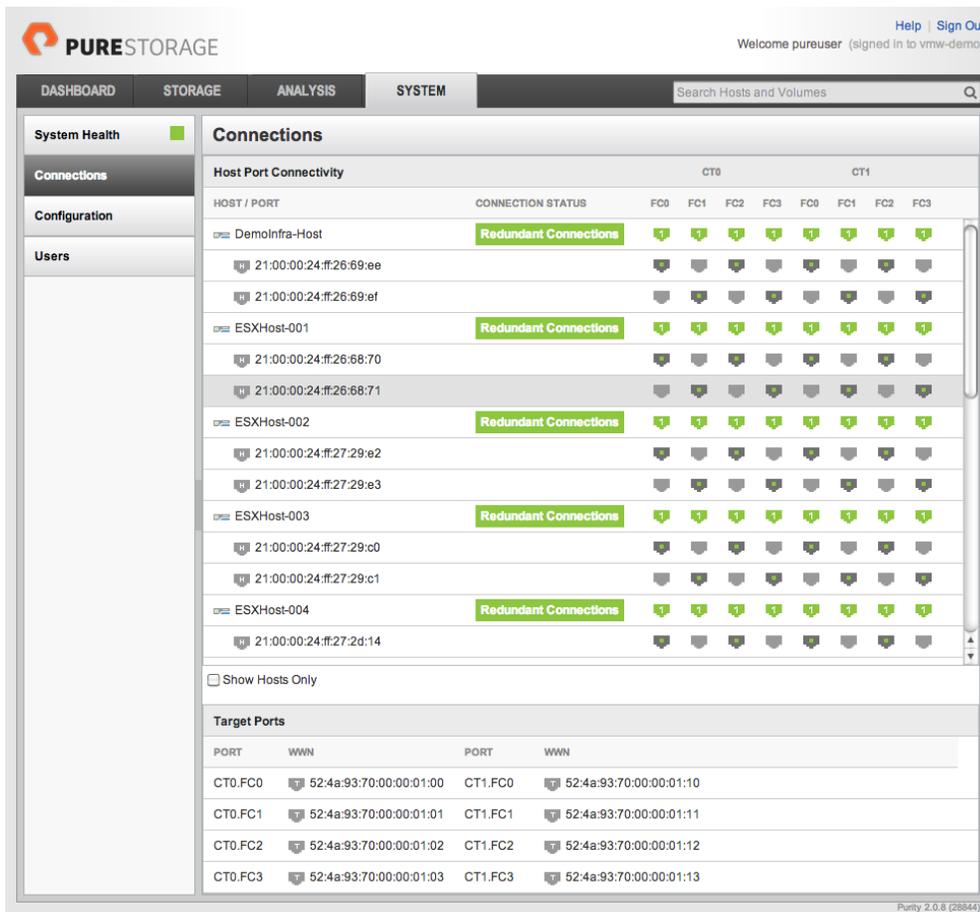


Figure 25: Connections map for hosts connected to Pure Storage FlashArray

Common operations like creating virtual machine clones from templates, Storage vMotion, VMFS datastore creation, VM snapshots, and general infrastructure deployment are all tremendously accelerated compared to mechanical disk based storage arrays.

Storage administrators have adapted to myriads of painstaking ordeals to provision storage and it is refreshing to see that those practices can be put to rest with the Pure Storage management approach. In addition to ease of storage capacity management and provisioning we found several benefits that help in rapid deployment and adaption of virtual desktops. These benefits of the FlashArray for virtual desktops are broadly classified into three sections and we describe them in detail in the next subsection.

- VDI operational aspects – Desktop group maintenance (patching desktops), efficiencies while booting or rebooting desktops.
- Data center power, cooling and rack space savings.
- Lower cost per desktop.

Benefits of Pure Storage FlashArray in Common Virtual Desktop Operations

Patching Desktops

Pushing out patches to desktops is a common operation and hosting the desktop in the data center promises to make this process much more efficient. But in traditional deployments, this task consumes a lot of time and results in very high backend disk array IOPS. Desktop admins perform this task during non-peak hours, weekends and in small batches. This not only makes pushing out patches more efficient, but also keeps the organization free from malware, virus infections, worms and other common bugs that plague desktops due to a lack of timely updates. The efficiency of the organization improves and the software applications are always up-to-date. With Citrix PVS, patching involves updating a single image and updates are pushed out after users log back in.

Creation of new Desktops

Desktop administrators can create new desktops using XenDesktop Setup Wizard on PVS, desktop groups and catalogs anytime of the day, as the IOPS and latency are not a limiter on the FlashArray. Adding new desktop groups, creation of catalogs, and desktop assignments are extremely fast.

Reduced Desktop Boot Times

When using mechanical disk storage arrays, constant Desktop Group maintenance activities can lead to unpredictable spikes in storage performance. This is a real killer in virtual desktop deployment projects. When Desktop admins spin up a new Desktop Group, the desktops power up and create a boot storm, this has the adverse effect of hindering the active users. The same is true when users login and logout; creating login and logoff storms respectively.

We simulated the worst-case scenario of powering on 500 virtual desktops and measuring the backend IOPS. Figure 26 below shows the Pure Storage GUI dashboard for this activity. We were able to boot the 500 desktops in less than 10 minutes and shutdown 500 desktops on less than 10 minutes while maintaining less than 1 millisecond latency. Yes, the FlashArray can withstand heavy loads like boot storms and still deliver sub-millisecond latency.



Figure 26: Booting 500 VMs in less than 10 minutes with < 1 msec latency

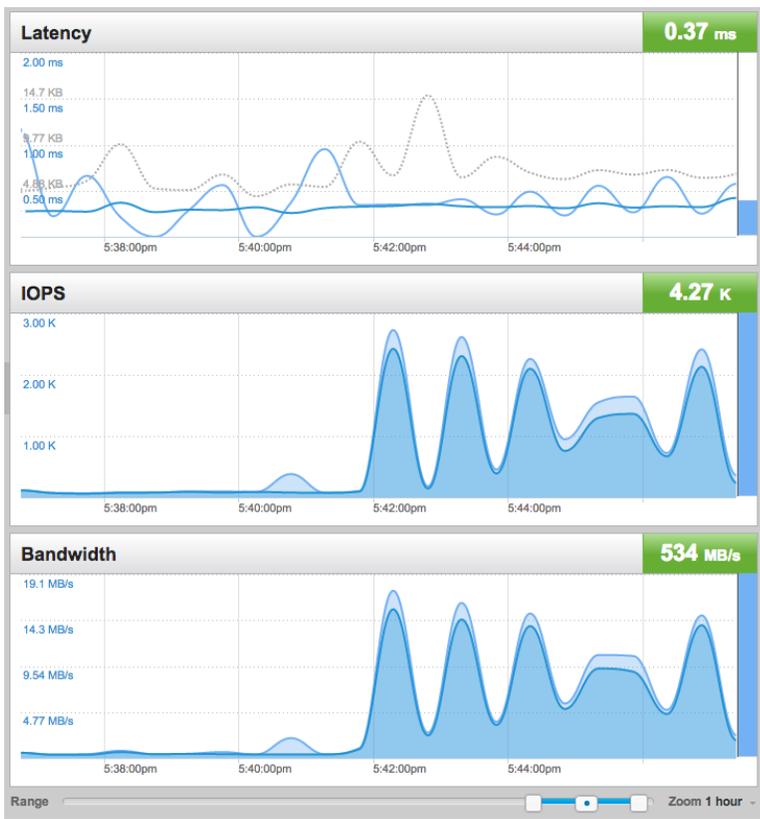


Figure 27: Shutting down 500 VMs in less than 10 minutes with < 1 msec latency

Data Center Efficiencies

As data centers get denser and denser, the lack of rack space and challenges around power and cooling (availability and rising cost) are increasingly prevalent problems. In comparison to a spindle based storage system, the Pure Storage FlashArray has no moving parts except for the fans. The FlashArray's smallest configuration can fit in 4 Rack Units (RU) space in a standard data center rack. This section talks in detail about the additional data center efficiency advantages the FlashArray brings. As data center power consumption goes through the roof, power becomes a huge factor in storage total cost of ownership calculations. The overall VDI operational budget should consider datacenter power and space, as previously a desktop's power and cooling impact was coming out of the facilities budget and not from the datacenter budget.

Power and Cooling

A fully loaded FA-320 with two controllers and two shelves uses less than 10 Amps of power (110V AC). The FA-310 with one controller and one shelf consume about a half of that (i.e. 5 Amps of power). The SSDs used in the FlashArray are low power consuming devices and dissipate very little heat, which in turn reduces cooling overhead.

Rack Space Savings

The FA-310, 4RU box can deliver up to 200K IOPS and with latencies less than 1 msec. A fully loaded FA-320 occupies only 8 RU of rack space and delivers same results complete with high availability. Rack space is a highly priced commodity in the data center, the advantages of having a low foot print will help in scaling the number of desktops per rack units of storage.

Lower Cost per Desktop

The number one cost in any virtual desktop deployment is storage. Scaling from a pilot of a few hundred desktops to large-scale production requires a lot more capacity and IOPS, which is readily available on the FlashArray. Throughout the various phases of this project, we deployed more than 2,000 virtual desktops on a single FlashArray. Based on the test results, you can easily put in excess of 5,000 desktops on the FlashArray. Additional desktops do not consume additional storage capacity (except the user data) as we are seeing a considerable write cache data reduction.

In the next section we talk more about the different configurations of the FlashArray that can be procured for your VDI deployment. The different form factors are designed to host certain number of desktops and cost varies based on the FlashArray configuration deployed.

Sizing Guidelines

The space consumption and the IOPS we saw in the 500 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 Figure 28 below, a pilot can be implemented on a single controller and ½ drive shelf system. As the deployment passes out of the pilot phase, you can upgrade to a two-controller HA system and ½ drive shelf for 1,000 desktops. As your user data grows, additional shelves can be added. Both controllers and shelves can be added without downtime.

If more desktops are needed, customers can expand to a full shelf to accommodate up to 2,000 desktops. For a 5,000 desktop deployment or larger, we recommend a fully-configured FA-320 with two controllers and two drive shelves. The sizing guidelines below are approximations based upon best practices, 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.

Adding a new shelf to increase capacity is very straightforward and involves simply connecting SAS cables from the controller to the new shelf. This is a non-disruptive operation and can be done while the array is online. 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.

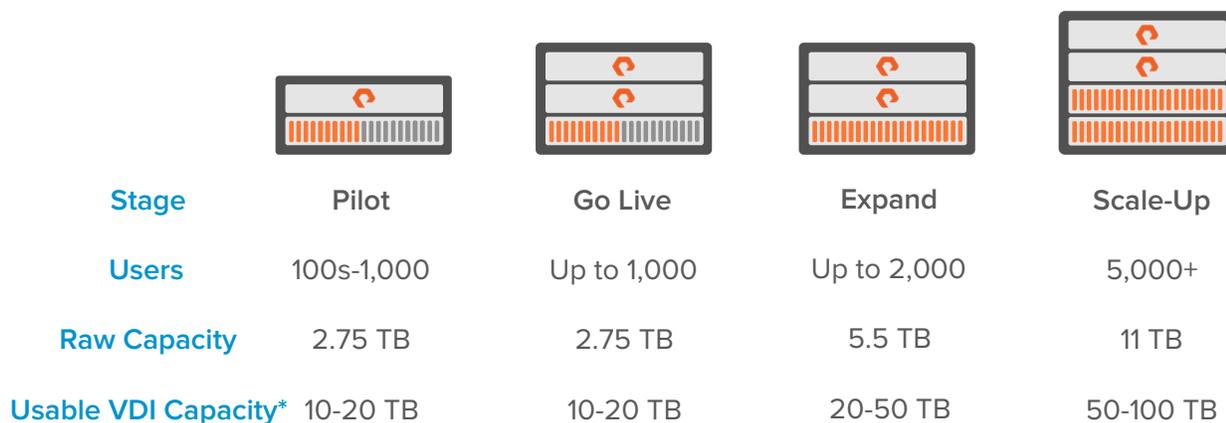


Figure 28: Pure Storage Virtual Desktop Sizing

Conclusion

Virtual desktop deployment is an ideal use case for the Pure Storage FlashArray; we achieved unprecedented results while running an industry-standard desktop workload generator. The sub one second response time on Login VSI benchmark is a testimony of the FlashArray's ability to deliver an unmatched VDI end-user experience. Beyond user experience, the FlashArray demonstrated additional VDI administrative and operational benefits, including rapid desktop provisioning, ease of storage management, lower storage cost, lower power, rack space savings, and lower cooling requirements. The FlashArray's integrated data reduction delivered >10-to-1 reduction of the VDI workload, enabling the use of desktops, and delivering all-flash VDI storage for less than \$100/desktop in most configurations. Furthermore, we expect the FlashArray can scale up to 5,000+ virtual desktops with proper infrastructure in place.

Now that Pure Storage has broken the price barrier for VDI on 100% flash storage, why risk your VDI deployment on disk?

About the Author



Ravindra “Ravi” Venkat is a Virtualization Solutions Architect at Pure Storage 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 two 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 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.

For all his contribution to the Virtualization community, Ravi was awarded the vExpert for 2013.

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

References

1. Benchmarking Citrix XenDesktop using Login VSI:
<http://support.citrix.com/servlet/KbServlet/download/24130-102-673656/XenDesktop%20with%20Login%20Consultants%20VSI%20V3%20-%20Guide.pdf>
2. Pure Storage FlashArray – Virtualization Benefits: Three part blog article:
<http://www.purestorage.com/blog/say-goodbye-to-vm-alignment-issues-and-poor-performance-with-pure-storage/>
3. DSNRO, the story: <http://www.yellow-bricks.com/2011/06/23/disk-schednumreqoutstanding-the-story/>
4. Login VSI documentation at - <http://www.loginvsi.com/documentation/v3>
5. PVS write cache: <http://blogs.citrix.com/2011/10/06/pvs-write-cache-sizing-considerations/>
6. Citrix CTX article for adding persistent write cache: <http://support.citrix.com/article/CTX125590>
7. Citrix PVS 6.1 eDocs - <http://support.citrix.com/proddocs/topic/provisioning-61/pvs-provisioning-61.html>
8. General Citrix XenDesktop 5.6 documentation - <http://support.citrix.com/proddocs/topic/xendesktop-rho/cds-landing-page-xd56.html>

About Login VSI

Login Virtual Session Indexer ([Login VSI](#)) is the de-facto industry standard benchmarking tool to test the performance and scalability of centralized Windows desktop environments like Server Based Computing (SBC) and Virtual Desktop Infrastructures (VDI).

SBC and VDI oriented vendor organizations committed to enhance end-user experience in the most efficient way, use Login VSI as an objective benchmark to test, benchmark and improve the performance and scalability of their solutions. VSImax provides the proof (vendor independent, industry standard and easy to understand) to innovative technology vendors to demonstrate the power and scalability (gains) of their solutions.

Login VSI based test results are published in technical white papers (for an overview visit www.loginvsi.com/white-papers) and presented in conferences. Login VSI is also used by end-user organizations, system integrators, hosting providers and testing companies.

For more information about Login VSI, or a free test license, visit www.loginvsi.com

APPENDIX A

Pure Storage LUN provisioning

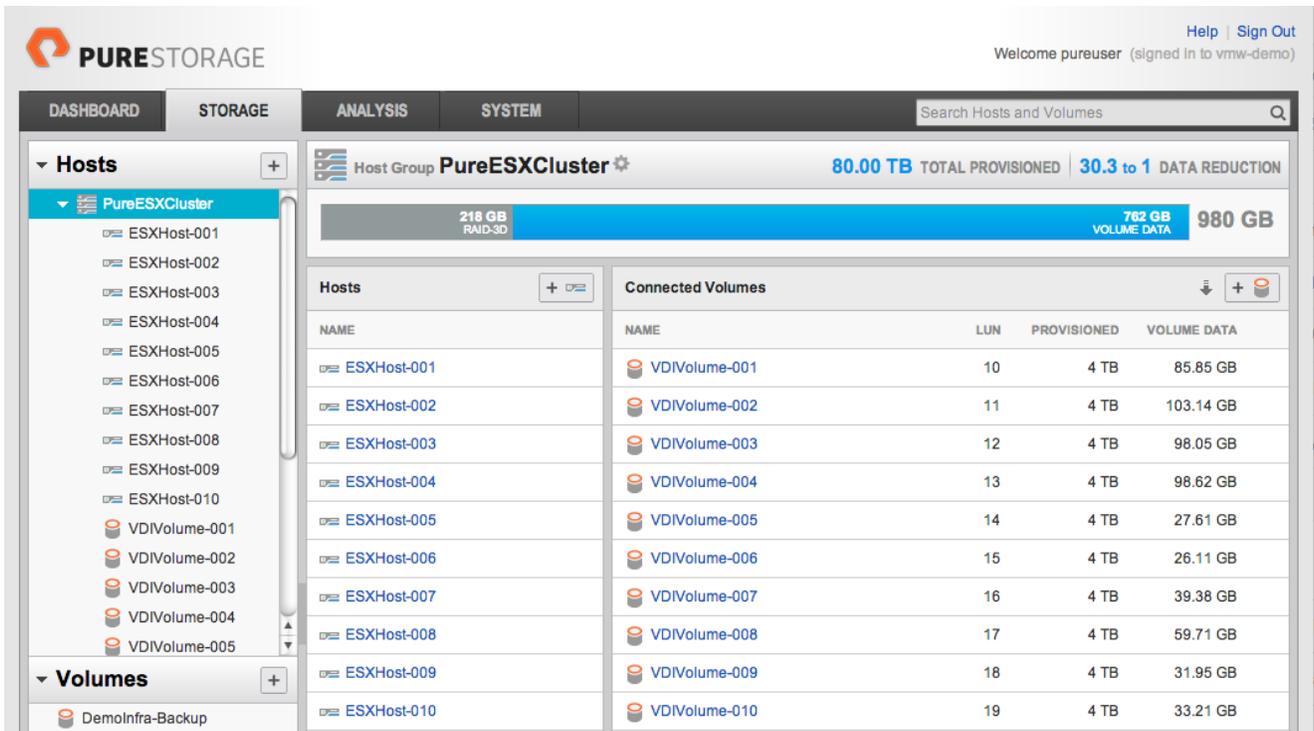
The following example creates a 40 TB volume, VDIVolume-001 and a host called ESXHost-001. A hostgroup called PureESXCluster is created with ESXHost-001 and the volume, VDIVolume-001 is connected to the hostgroup, PureESXCluster :

```
purevol create VDIVolume-001 -size 40t  
  
purehost create -wwnlist 21:00:00:00:ab:cd:ef:00,21:00:00:00:ab:cd:ef:01 ESXHost-001  
  
purehgroup create -hostlist ESXHost-001 PureESXCluster  
  
purehost connect -vol VDIVolume-001 PureESXCluster
```

New hosts are created using step 2 and

```
purehgroup setattr -addhostlist HOSTLIST HGROUP is used new hosts to the host group.
```

The figure below shows the Pure Host Group and LUN configuration.



The Pure Storage GUI can accomplish this using similar operation.

APPENDIX B

Cisco MDS zoning Example

Single initiator and single target zoning example script to configure a single port of the Pure Storage FlashArray port to all initiator HBA ports.

Example:

```
# conf t

(config) # zoneset name pure-esx-vdi-cluster-zoneset vsan 100

(config-zoneset) # zone name zone_pureArray_Port1_hpesx2_vmhba1

(config-zone) # member pwwn 21:00:00:24:ff:23:27:aa

(config-zone) # member pwwn 21:00:00:24:ff:32:87:32

(config-zone) # exit

(config-zoneset) # zone name zone_pureArray_Port1_hpesx2_vmhba2

(config-zone) # member pwwn 21:00:00:24:ff:23:27:aa

(config-zone) # member pwwn 21:00:00:24:ff:27:29:e6

(config-zone) # exit

(config-zoneset) # zone name zone_pureArray_Port1_hpesx2_vmhba3

(config-zone) # member pwwn 21:00:00:24:ff:23:27:aa

(config-zone) # member pwwn 21:00:00:24:ff:32:87:26

(config-zone) # exit

(config-zoneset) # zone name zone_pureArray_Port1_hpesx2_vmhba4

(config-zone) # member pwwn 21:00:00:24:ff:23:27:aa

(config-zone) # member pwwn 21:00:00:24:ff:27:2d:04

(config-zone) # exit
```

APPENDIX C

Setting up Round-Robin PSP on a Pure LUN:

The screenshot shows the vCenter configuration interface for a VM named 'pes01.dev.purestorage.com'. The 'Configuration' tab is active, showing hardware and software settings. A table of datastores is visible, listing various Win7VDI volumes. A 'PURE Fibre Channel Disk' management dialog is open, showing a 'Round Robin (VMware)' storage array type and a list of paths. A 'Win7VDI Volume-001 Properties' dialog is also open, showing volume details and a 'Manage Paths' button. Numbered callouts (1-9) and orange arrows indicate the sequence of steps for setting up Round-Robin PSP on a Pure LUN.

Identification	Status	Device	Drive Type	Capacity	Free	Type	Last Update	Alarm Actions	SR
Win2k8_vol_2	Normal	PURE Fibre Chan...	SSD	64.00 TB	63.83 TB	VMFS5	11/29/2012 5:36:09 PM	Enabled	Dis
Win2k8_vol_3	Normal	PURE Fibre Chan...	SSD	30.00 TB	29.90 TB	VMFS5	11/29/2012 5:36:18 PM	Enabled	Dis
Win2k8_vol_4	Normal	PURE Fibre Chan...	SSD	40.00 TB	39.90 TB	VMFS5	11/29/2012 5:36:18 PM	Enabled	Dis
Win7VDIVolume-...	Normal	PURE Fibre Chan...	SSD	4.00 TB	1.67 TB	VMFS5	11/29/2012 6:00:55 PM	Enabled	Dis
Win7VDIVolume-...	Normal	PURE Fibre Chan...	SSD	4.00 TB	1.64 TB	VMFS5	11/29/2012 6:00:55 PM	Enabled	Dis
Win7VDIVolume-...	Normal	PURE Fibre Chan...	SSD	4.00 TB	1.64 TB	VMFS5	11/29/2012 6:00:55 PM	Enabled	Dis
Win7VDIVolume-...	Normal	PURE Fibre Chan...	SSD	4.00 TB	2.52 TB	VMFS5	11/29/2012 6:00:55 PM	Enabled	Dis

Runtime Name	Target	LUN	Status	Preferred
vmhba2:CO:17:1:10	S2:4a:93:70:00:01:10	S2:4a:93:70:00:01:10	Active (I/O)	
vmhba2:CO:15:1:10	S2:4a:93:70:00:01:12	S2:4a:93:70:00:01:12	Active (I/O)	
vmhba2:CO:15:1:10	S2:4a:93:70:00:01:00	S2:4a:93:70:00:01:00	Active (I/O)	
vmhba2:CO:14:1:10	S2:4a:93:70:00:01:02	S2:4a:93:70:00:01:02	Active (I/O)	
vmhba3:CO:17:1:10	S2:4a:93:70:00:01:11	S2:4a:93:70:00:01:11	Active (I/O)	
vmhba3:CO:16:1:10	S2:4a:93:70:00:01:13	S2:4a:93:70:00:01:13	Active (I/O)	
vmhba3:CO:15:1:10	S2:4a:93:70:00:01:03	S2:4a:93:70:00:01:03	Active (I/O)	

APPENDIX D

Test configuration file (VSIlauncher.ini file)

```
[Launcher]
Servername=
Username=
Password=enc::enc
Domain=
CommandPassword=
ConnectionType="Custom with CSV file"
ConnectionNumber=User
CCL=\\CitrixReadyVM\VSIShare\CTXConnector\CTXConnector.exe -s
"http://citrixddc/Citrix/DesktopWeb/auth/login.aspx" -u %CSV_User% -p xxxxxxxx -d
"vdi.puredemo.com" -r "DesktopsVSI" -t1 5 -t2 10
CSV=\\CitrixReadyVM\VSIShare>LoginVSI_Users.csv
Launchmode=Parallel
PreTestScript=
PostTestScript=
ParallelDelay=10
ParallelTimeframe=1600
InitialStartNumber=1
NumberOfSessions=525
SequentialInterval=30
Fancy_number=1
Autologoff=1
LogoffTimeOut=900
CreateProfile=0
UseLocalLauncher=1
```

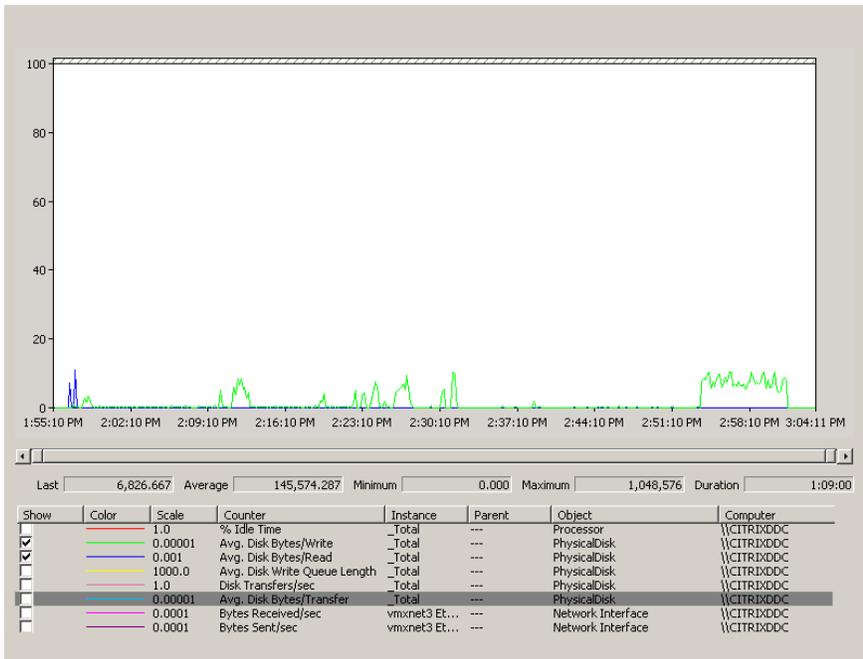
APPENDIX E

DDC Resource Utilization Statistics (PVS Write Cache on Server)

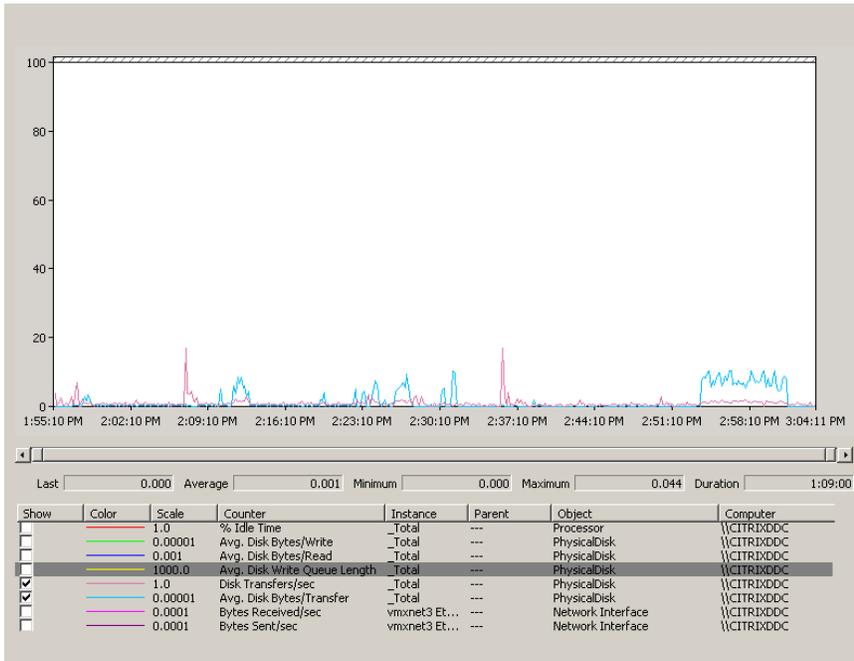
DDC Perfmon CPU Statistics



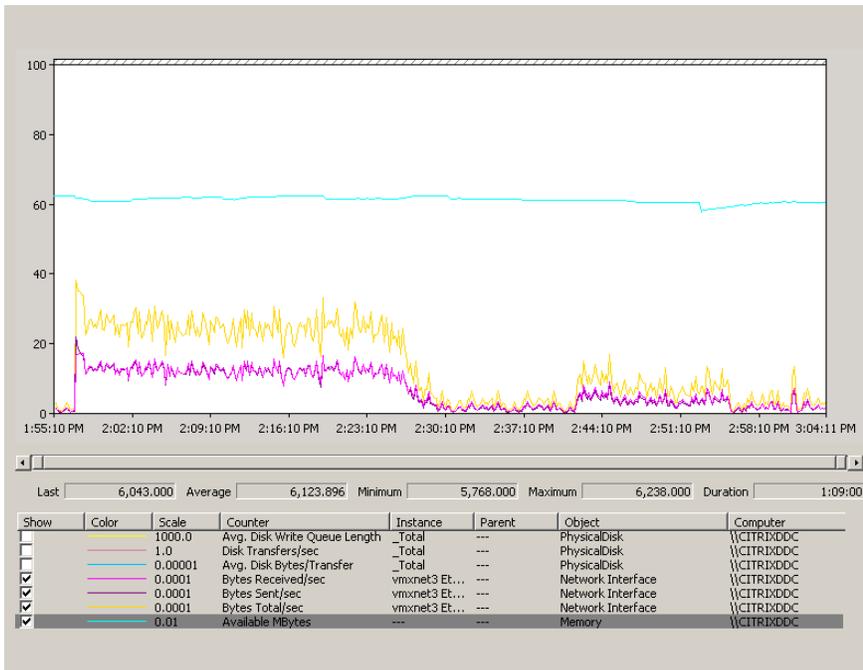
DDC Perfmon Disk Statistics



DDC Perfmon Disk Data Transfer Statistics



DDC Perfmon Network and Memory Statistics



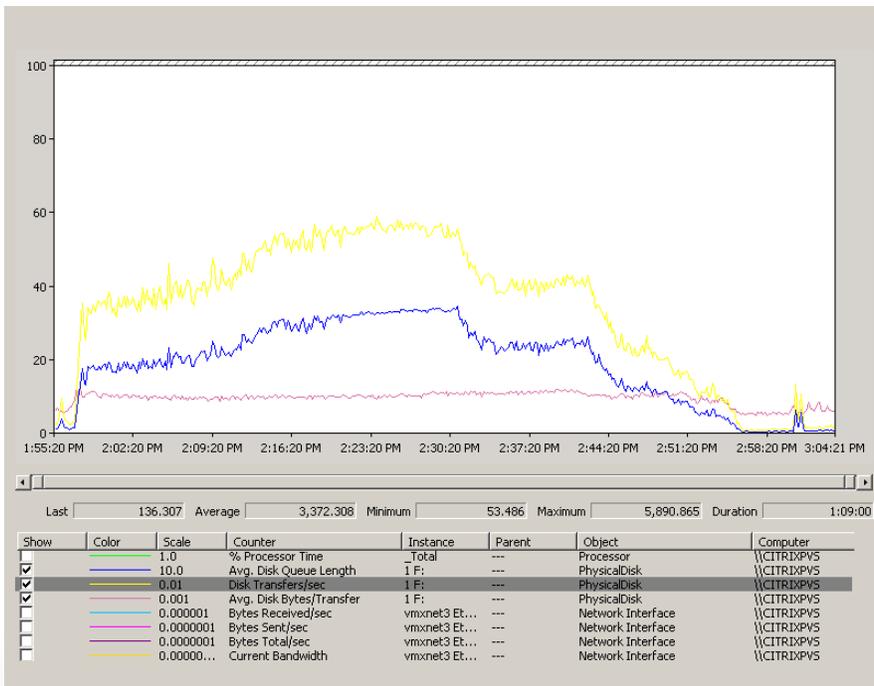
APPENDIX F

PVS Resource Utilization Statistics (PVS Write Cache on Server)

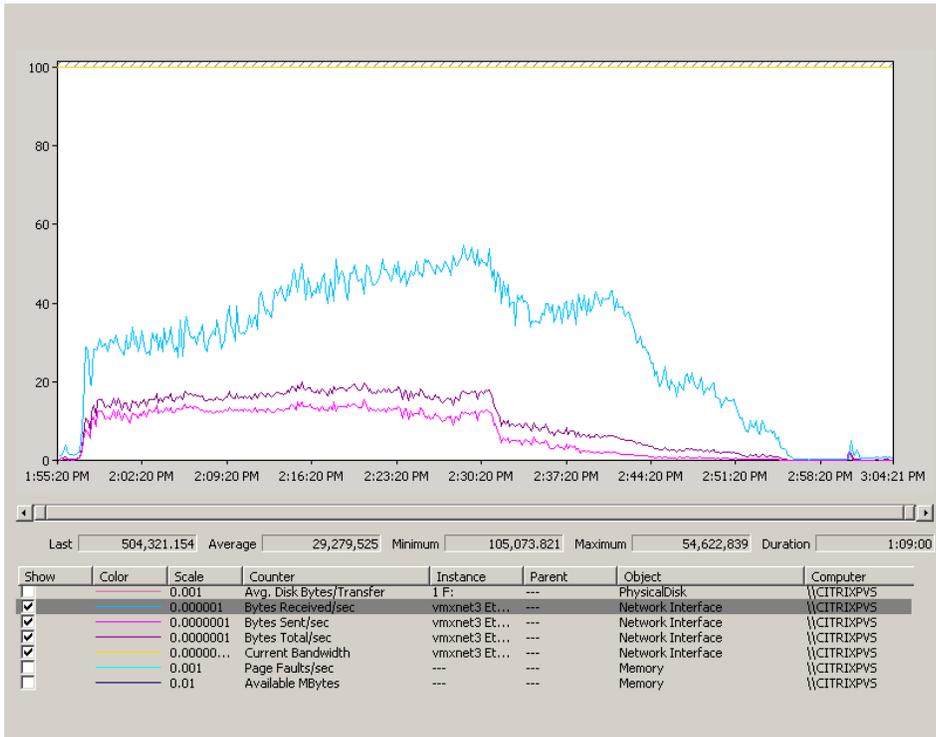
PVS Perfmon CPU Statistics



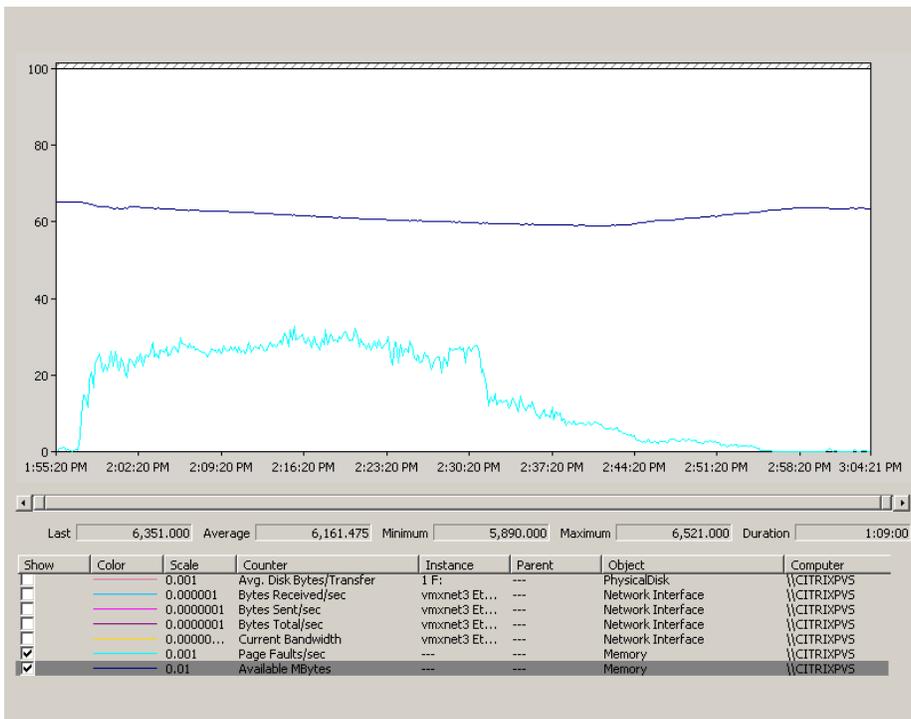
PVS Perfmon Disk Statistics



PVS Perfmon Network Statistics



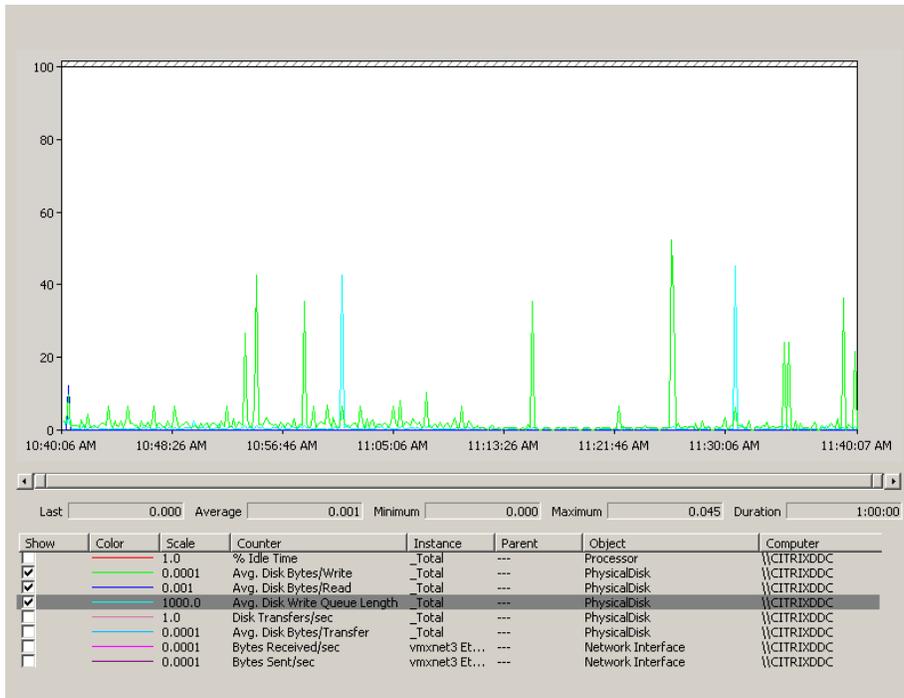
PVS Perfmon Memory Statistics



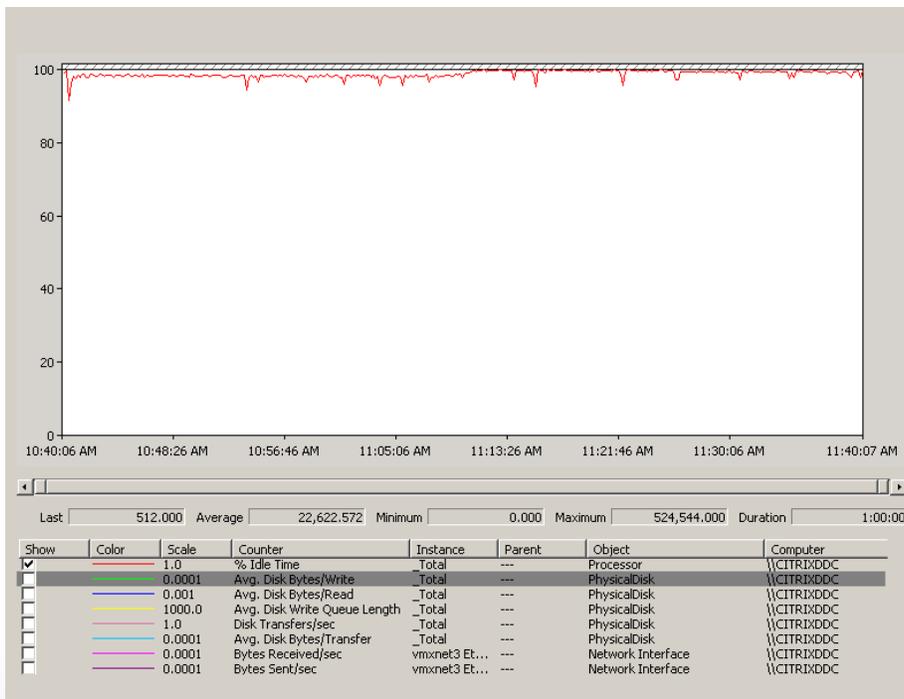
APPENDIX G

DDC Resource Utilization Statistics - PVS Write Cache on Hard Disk

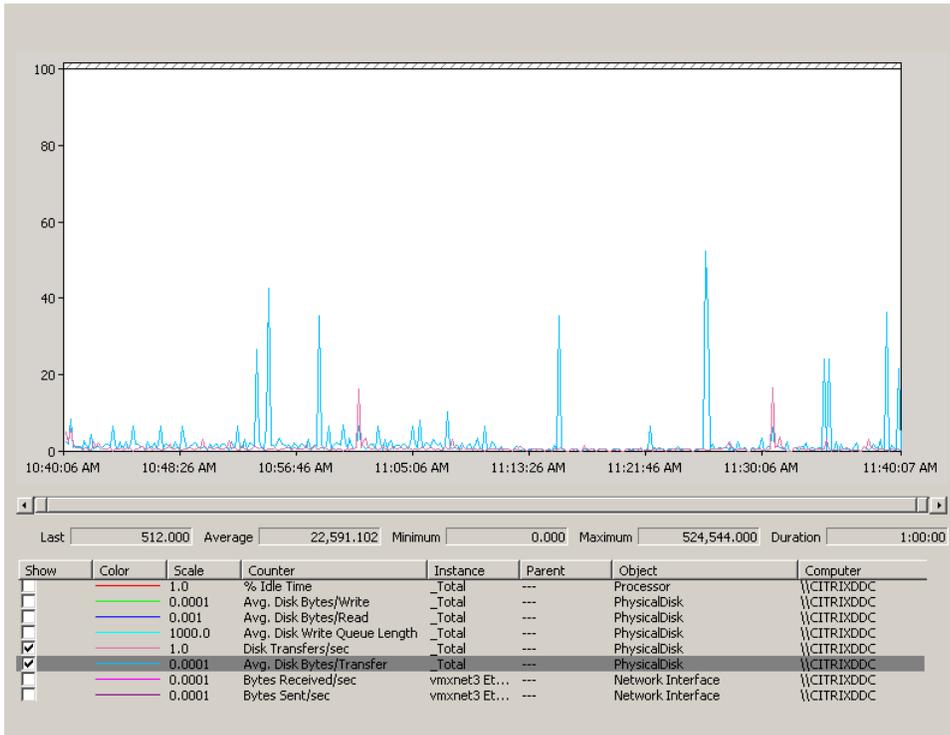
DDC Perfmon CPU Statistics



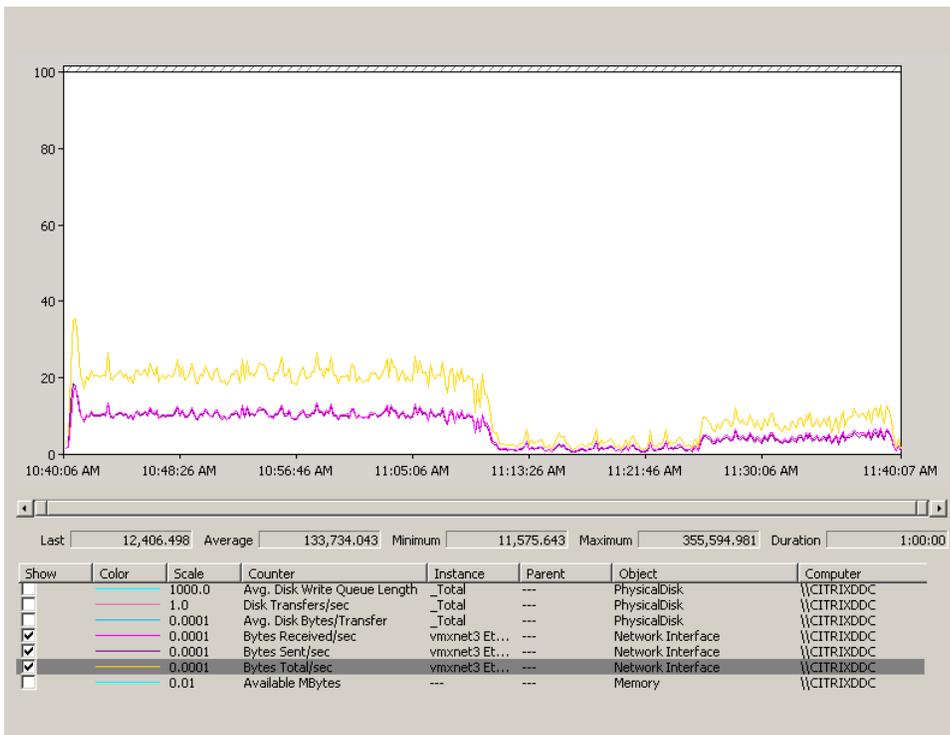
DDC Perfmon Disk statistics



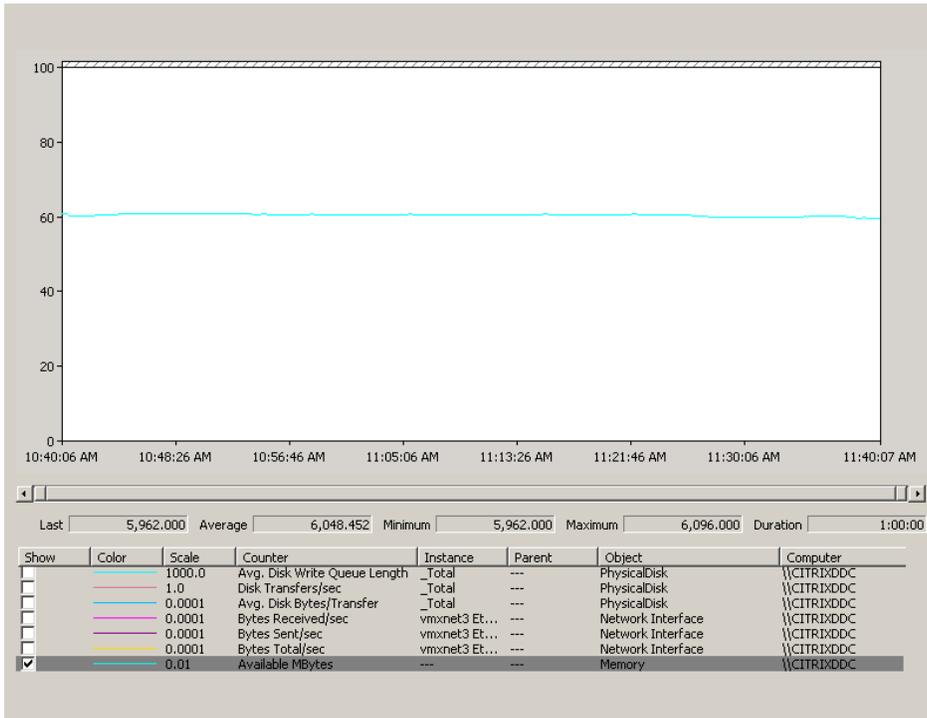
DDC Perfmon Disk transfer Statistics



DDC Perfmon Network Statistics



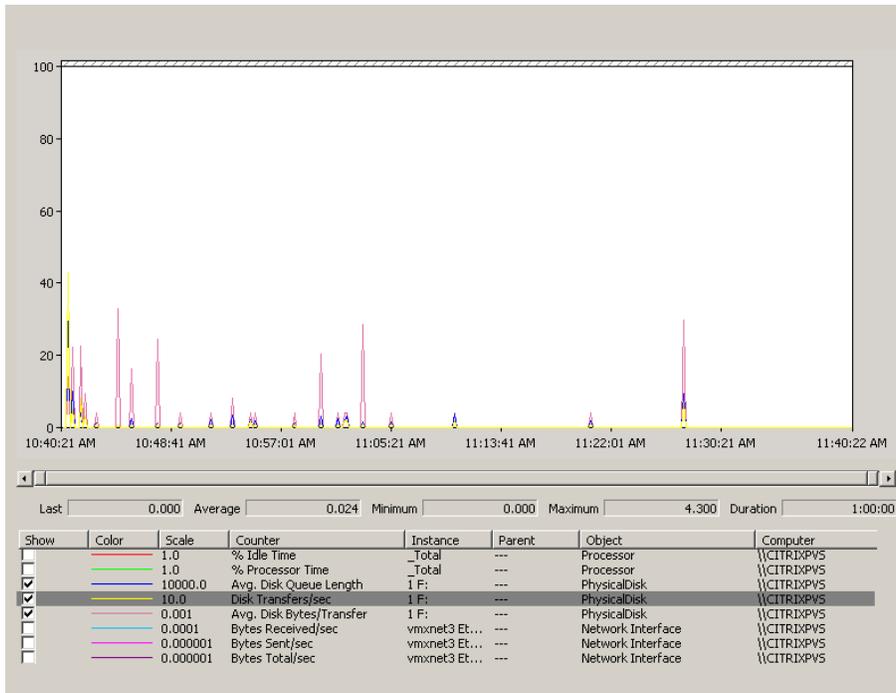
DDC Perfmon Memory consumption Statistics



APPENDIX H

PVS Resource Utilization Statistics - PVS Write Cache on Hard Disk

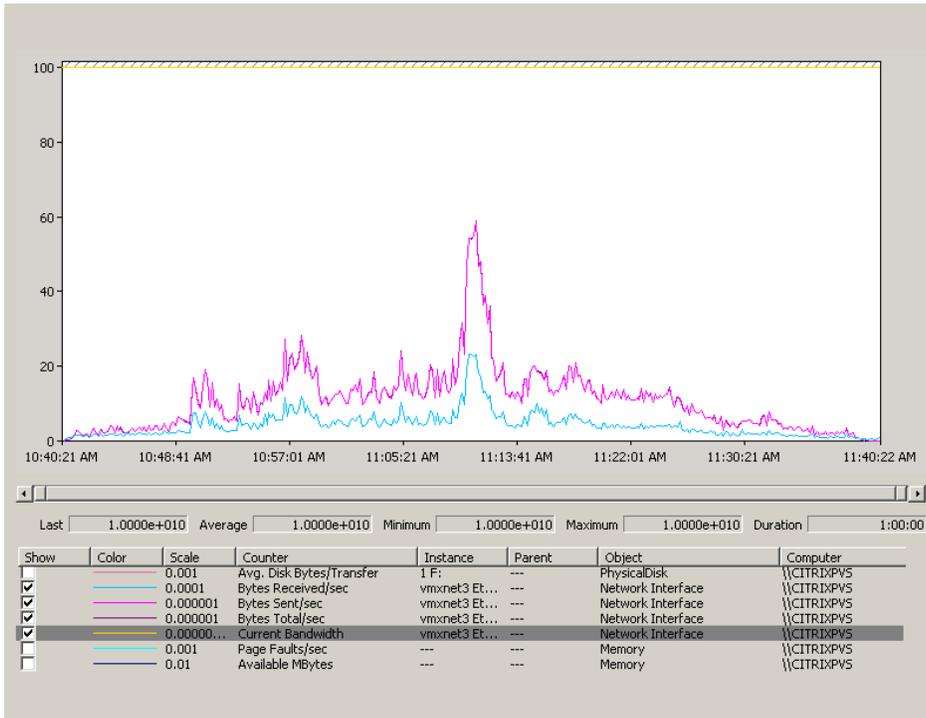
PVS CPU Perfmon Statistics



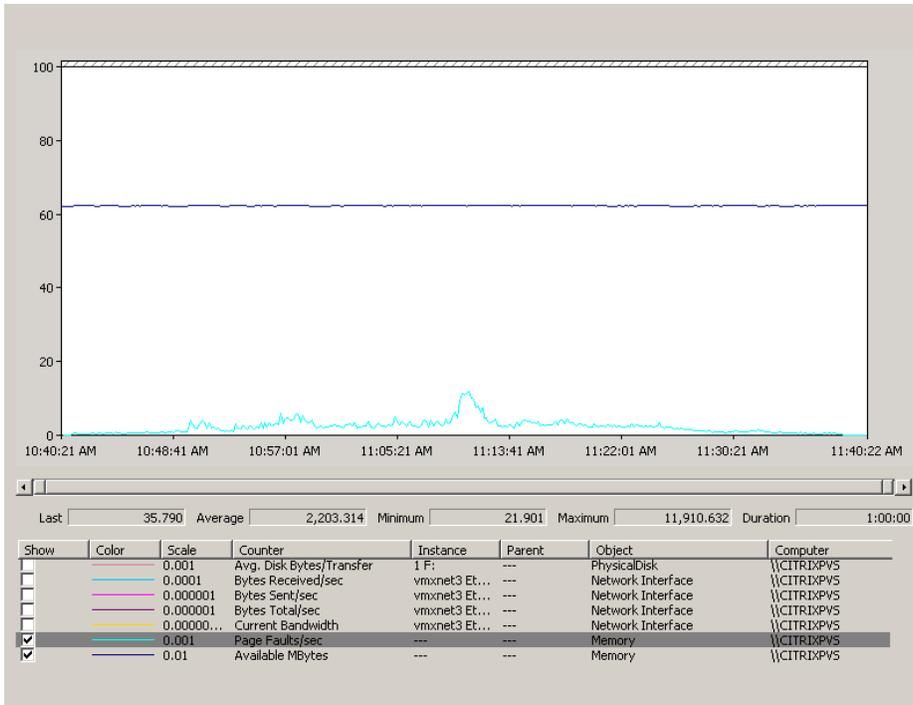
PVS Disk Transfer Perfmon Statistics



PVS Network Perfmon Statistics



PVS Memory Perfmon Statistics





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