

Pure Storage Reference Architecture for VMware[®] View[™]

purestorage.com

Overview

This document describes a reference architecture for deploying virtual desktops on the Pure Storage FlashArray using VMware[®] View[™] 5.0, View Composer 2.7, vSphere 5.1 hypervisor and Microsoft Windows 7. Pure Storage has validated the reference architecture with VMware's View Planner 2.1 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 VMware View based Windows 7 desktops on the FlashArray by deploying 1,000 virtual desktops in both linked clone and full clone persistent desktop configurations and running the VMware View Planner 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 VMware View based virtual desktops on the FlashArray. A working knowledge of VMware vSphere, VMware View, server, storage, network and data center design is assumed but is not a prerequisite to read this document.



Summary of Findings

- We deployed 1,000 VMware View based linked clone Windows 7 desktops and ran a realistic load generator with VMware View Planner that simulated 1,000 users performing common computing tasks, resulting in a best-in-class score of 0.52 seconds. This score means the majority of the applications (95% of group "A" interactive operations) had a response time of 0.52 second or less, well within the passing score of 1.5 seconds.
- We then repeated the test using 1,000 persistent full-clone desktops, achieving the same View Planner score and showing that **users can confidently use any combination of linked clone or full clone persistent desktops** on the FlashArray both perform the same.
- Throughout the testing the FlashArray delivered up to 50,000 IOPS and maintained latency
 under 1.1ms, 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 2,000 desktops, including both 1,000 linked clones and 1,000 persistent desktops (each of 31 GB disk size), together only consuming about 1.1 TB of physical storage on the FlashArray. This massive data reduction (>20-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 11TB FlashArray FA-320 delivered best-in-class VDI performance at a cost of \$100/desktop for 2,000 desktops. 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 recomposing desktops, 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-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 1,000 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-320 can support up to 5,000
 desktops with any mix of linked clones and/or persistent desktops.



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 indeed 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 deployments. The reasons why are many, 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, expensive 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 followed by the test infrastructure components that was put together for this work and dive into the details of each component. Finally, we'll discuss the results of the VMware View Planner load generator and the operational benefits of using 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, and 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 enterpriseclass 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 1,000 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 VMware View infrastructure components were placed on a dedicated host. We tested 1,000 linked clone desktops, 1,000 full-clone persistent desktops, and various mixtures of the two. The infrastructure, virtual machines and desktops were all hosted on a single 11TB 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 VMware View best practices were used in addition to the stringent requirements as mandated by the View Planner guideline document [See reference 1].

The tested configuration included:

- One 11TB Pure Storage FlashArray (FA-320) in HA configuration, including two controllers and two disk shelves:
 - Ten x 4 TB volumes were carved out of the Pure FlashArray to host 2,000 desktops (1,000 linked clones + 1,000 persistent desktops)
 - A separate 600 GB volume was used to hold all the infrastructure components
- Eight Intel Xeon x5690 based commodity servers with 192 GB of memory running ESXi 5.1 were used to host the desktops
- One dedicated server was used to host the all of the infrastructure virtual machines:
 - Active directory, DNS, and DHCP
 - View Connection server



- Virtual Center server
- SQL server for both virtual center and View event database
- VMware View Planner 2.1 appliance



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



Reference Architecture Configuration

This section describes the 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 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.

Hardware Configuration

Pure Storage FlashArray FA-320 configuration

The FlashArray FA-320 configuration comprised of two active/active controllers and two shelves of 5.5TB of raw flash memory for a total of 11TB of raw storage. Four Fibre 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 (24Gb) to two shelves and were interconnected for HA via two redundant InfiniBand connections (40Gb)
Shelves	Two flash memory shelves with 22 SSD drives, 22 X 256 GB or a total raw capacity of 11TB (10.3 TiB).
External Connectivity	Four 8Gb 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

Ten thin provisioned volumes of 4 TB each were configured to host 1,000 linked clone desktops and 1,000 persistent 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 ten 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. For boot from SAN, Pure storage provides private volumes that can be used for boot the LUN.

A common question when provisioning storage is how many LUNs of what size should be created to support the virtual desktop deployment. Because linked clone desktops take very little space, we could have either put all the virtual desktops in one big LUN or spread them across several LUNs. The FlashArray supports the VMware VAAI ATS primitive which gives you access of 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 used to severely limit VM scalability in previous ESX versions. See Appendix A for more details on provisioning Pure Storage.

Since we are advocating placing the OS image, user data, persona and application data on the same storage, we need to take in account the size of those drives when calculating the LUN size.

Consider a desktop with a 30 GB base image including applications and app data, 20GB of user data and we need to provision "d" desktops:

- We need to provision 50 * d as the size of the LUN --or--
- One could distribute the "d" desktops across "n" LUNs with (50 * d) / n desktops on each 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 4 below shows the virtual desktop deployment on Pure Storage.



Figure 4: OS image, applications, user data and application data hosted on Pure Storage



Data Reduction with Pure Storage FlashArray

Storage Capacity with 1,000 Linked Clone Desktops

Figure 5 below shows 1,000 Windows 7 linked clone desktops deployed on a brand new FlashArray. The total physical capacity used was 66.1 GB for the entire 1,000 desktops. Because linked clones only store the differential data, we achieve a 25 to 1 data reduction. In essence, 40TB of provisioned storage actually consumed about 66 GB of space on the flash memory.

Capacity	40.00 TB TOTAL PROVISIONED	25.6 to 1 DATA REDUCTION	1% full
66.1 GB TOTAL USED			10.3 TB

Figure 5: Data reduction of 1,000 Windows 7 linked clone desktops

Storage Capacity with 2,000 Linked Clone and Persistent Desktops

When we provisioned 2,000 persistent desktops (each with 31 GB disk capacity) we also experienced a 25 to 1 data reduction; the entire 31 TB was stored on 1.19 TB of physical media. This result accounts for the RAID-3D protection, shared data and the volume data as shown below. Note that space reporting doesn't include zeros that EZT (Flat) and ZT (thick) VMFS volumes.



Figure 6: Data reduction of 2,000 desktops: 1,000 linked clones plus 1,000 full clones

The persistent desktops had 1 GB of user data each, so for 1,000 desktops for different view Planner user customization (profiles/registry settings), the user data added up to 1.19 TB. 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. Note that the OS image doesn't add to the physical space as most blocks are deduplicated.

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.



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 the Table B below.

Component	Description
Processor	2 X Intel Xeon X5690 @ 3.47GHz (12 Cores total, 24 Logical CPUs)
Memory	192 GB @ 1333 MHz (16GB X 12)
НВА	Dual port Qlogic ISP2532-based 8Gb Fibre Channel PCIe card
NIC	Quad port Intel Corp 82576 1Gb 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 8Gb 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. 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 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 7 shows vCenter datastore details). Check Appendix B for a sample Cisco MDS zoning.

Datastore Details							
Win7YDIYolume-001 Location: /vmfs/vo Hardware Acceleratior	l Jumes/508e988b-5f6dd n: Supported	18e-b331-002590	2687b0	4.00 TB 2.33 TB	Capacity		
Refresh Storage Capa System Storage Capal User-defined Storage	abilities bility: N/A Capability: N/A			1.67 TB	Free Free		
Path Selection Round Robin (VM	Properties Volume Label:	Win7VDIVol	Extents PURE Fibre Cl	nannel Disk	.(n	4.00 TB	Storage I/O Control Disabled
Paths Total: 8 Broken: 0 Disabled: 0	Datastore Name: Formatting File System: Block Size:	Win7VDIVol VMFS 5.58 1 MB	Total Formatt	ed Capacit	Y	4.00 TB	

Figure 7: VMware VMFS datastore details

Network Configuration

Figure 8 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 8: 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 immensely in getting administrative tasks done efficiently.

Pure Storage FlashArray Best Practices for vSphere 5

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



Policy				1
Path Selection:	Round Robin (VMware)		_	Change
Storage Array 1	Type: VMW_SATP_ALUA			
Paths				
Runtime Name	Target	LUN	Status	Preferred
vmhba2:C0:T7:L	.10 52:4a:93:70:00:00:01:10 52:4a:93:70:00:00):01:10 10	🔶 Active (I/O)	
vmhba2:C0:T6:L	.10 52:4a:93:70:00:00:01:12 52:4a:93:70:00:00):01:12 10	🔶 Active (I/O)	
vmhba2:C0:T5:L	.10 52:4a:93:70:00:00:01:00 52:4a:93:70:00:00):01:00 10	🔶 Active (I/O)	
vmhba2:C0:T4:L	.10 52:4a:93:70:00:00:01:02 52:4a:93:70:00:00):01:02 10	🔶 Active (I/O)	
vmhba3:C0:T7:L	.10 52:4a:93:70:00:00:01:11 52:4a:93:70:00:00):01:11 10	🔶 Active (I/O)	
vmhba3:C0:T6:L	.10 52:4a:93:70:00:00:01:13 52:4a:93:70:00:00):01:13 10	🔶 Active (I/O)	
vmhba3:C0:T5:L	.10 52:4a:93:70:00:00:01:03 52:4a:93:70:00:00):01:03 10	🔷 Active (I/O)	
				Refresh
Name:	fc.20000024ff266871:21000024ff266871-fc.524	a937000000103:524a93	7000000103-naa.624a9	37034be0128252
Runtime Name:	vmhba3:C0:T5:L10			
Fibre Channe	ł			
Adapter:	20:00:00:24:ff:26:68:71 21:00:00:24:ff:26:68:7	1		
Targety	52·4a·93·70·00·00·01·03 52·4a·93·70·00·00·01·	03		

Figure 9: 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. Because we were using VMware View 5.0, the cluster was restricted to eight hosts, which was sufficient for deploying 1,000 virtual desktops. 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 for persistent desktops and linked-clone 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 vowing to Enterprise Plus licensing requirement. The redundant NICs were teamed in active/active mode and

VLAN configurations were done on the upstream Cisco 3750 1GE 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 10 and Figure 11 respectively.



Figure 10: 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.

Ċ2	vSwit	ch_NAT Properties			
F	Ports	Network Adapters			
	Conf	iguration	Summary	vSphere Standard Switch Properties —	
	+++	vSwitch	248 Ports	Number of Ports:	248
		NAT 10.131/16	Virtual Machine		
	ă	NAT VMkernel	vMotion and IP	- Advanced Properties	
	<u> </u>			MTU:	1500

Figure 11: 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 12 below shows how to set it using a vCenter on a host by host basis.



Figure 12: 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 glogic 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.

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 SQL server VM hosted the Microsoft SQL Server 2008 R2 database instance of the vCenter database and the VMware View events logging database. The VMware® View Planner is a product of VMware, Inc. and can be obtained via the VMware partner program website. Their newest version of View Planner 2.1 is a Centos-based appliance available in the form of a OVF to deploy in vCenter. The description of each of the infrastructure component is shown in Figure 13.



Figure 13: 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 X5690 @ 3.47GHz (12 Cores total, 24 Logical CPUs)
Memory	192 GB @ 1333 MHz (16GB X 12)
HBA	Dual port Qlogic ISP2532-based 8Gb Fibre Channel PCIe card
NIC	4 X Intel Corp 82576 1Gb 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

VMware View 5 Overview

VMware View 5 is a desktop virtualization solution that simplifies IT manageability and control while delivering one of the highest fidelity end-user experiences across devices and networks.

The VMware View solution helps IT organizations automate desktop and application management, reduce costs and increase data security through centralization of the desktop environment. This centralization results in greater end-user freedom and increased control for IT organizations. By encapsulating the operating systems, applications and user data into isolated layers, IT organizations can deliver a modern desktop. It can then deliver dynamic, elastic desktop cloud services such as applications, unified communications and 3D graphics for real-world productivity and greater business agility.

Unlike other desktop virtualization products, VMware View is built on, and tightly integrated with, vSphere, the industry-leading virtualization platform, allowing customers to extend the value of VMware infrastructure and its enterprise class features such as high availability, disaster recovery and business continuity.

View 5 includes many enhancements to the end-user experience and IT control. Some of the more notable features include:

- *PCoIP Optimization Controls*—deliver protocol efficiency and enable IT administrators to configure bandwidth settings by use case, user or network requirements and consume up to 75 percent less bandwidth
- *PCoIP Continuity Services*—deliver a seamless end-user experience regardless of network reliability by detecting interruptions and automatically reconnecting the session
- *PCoIP Extension Services*—allow Windows Management Instrumentation (WMI)–based tools to collect more than 20 session statistics for monitoring, trending and troubleshooting end-user support issues
- *View Media Services for 3D Graphics*—enable View desktops to run basic 3D applications such as Aero, Office 2010 or those requiring OpenGL or DirectX—without specialized graphics cards or client devices
- View Media Services for Integrated Unified Communications—integrate voice over IP (VoIP) and the View desktop experience for the end user through an architecture that optimizes performance for both the desktop and unified communications
- View Persona Management (View Premier editions only)—dynamically associates a user persona with stateless floating desktops. IT administrators can deploy easier-to-manage stateless floating desktops to more use cases while enabling user personalization to persist between sessions
- *View Client for Android*—enables end users with Android-based tablets to access View virtual desktops



Support for VMware vSphere 5 leverages the latest functionality of the leading cloud infrastructure platform for highly available, scalable and reliable desktop services.

For additional details and features available in VMware View 5, see the release notes.

Typical VMware View 5 deployments consist of several common components (illustrated in Figure 14 below), which represent a typical architecture. It includes VMware View components as well as other components commonly integrated with VMware View.







VMware View Configuration

VMware View 5.0.1 (build 640055) was installed on a Windows 2008 R2 VM with 4 vCPU/8GB of memory. The View Composer 2.7.0 (build 481620) was installed on the vCenter VM 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 and a 1,000 dedicated Windows 7 desktops is shown below.

General Pool Settings Provisioning	g Settings VCenter Settings Guest Customization
General	
State:	Disabled
Connection Server restrictions:	None Browse
Remote Settings	
Remote Desktop Power Policy:	Take no power action 🛛 🔻 🔇
Automatically logoff after disconnect:	Never 🛛
Allow users to reset their desktops:	No 🗸
Allow multiple sessions per user:	No 🗸
Delete or refresh desktop on logoff:	Never V
Remote Display Protocol	
Default display protocol:	PCoIP 🗸
Allow users to choose protocol:	Yes 🔻
Windows 7 3D Rendering:	Disabled 💌 Configure 3
Max number of monitors:	2 🕶 🤇
Max resolution of any one monitor:	1920×1200 ▼ ③
Adobe Flash Settings for Remote Sessi	ons
Adobe Flash quality:	Do not control 🔻 🔇
Adobe Flash throttling:	Disabled 🛛 🔻 🔇

Figure 15: Automated desktop pool settings for Windows 7

Other changes to View connection included increasing the number of concurrent operations on the Virtual Center server.

VMware View Administrator \rightarrow View Configuration \rightarrow Servers \rightarrow vCenter Servers \rightarrow Edit \rightarrow Advanced

Specify the concurrent operation limits.

Max concurrent provisioning operations:	50
Max concurrent power operations:	100

In order to do faster recompose, we changed the settings on the View Composer to do batches of 100



desktops rather than the default 12 desktops [reference 1, page 18 has details on this]. FlashArray could easily sustain the load and the operations finished in a few minutes each; see the recompose section for more details.

Desktop Operating System - Microsoft Windows 7 Configuration

The View Planner document [see reference 4] provided guidelines for configuring the base Windows 7 image. In order to get a successful View Planner run, adjustments were made to the base image before we took a snapshot and created a pool of 1,000 linked clones and a separate pool of 1,000 dedicated persistent desktops. 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 + 1 GB data disk
VMware View Agent	VMware View Agent 5.0.1 build 640055
Installed Application	MS Office 2007, Adobe Reader 9, FireFox 3.6.19 , Internet Explorer 9, and all customizations as mandated by View planner 2.1

Table D: Windows 7 virtual desktop configuration summary

Software Testing Tool – VMware View Planner 2.1

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, then accepted the resultant application latency as a metric to measure end user experience.

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 View based desktops with the

following settings:



The View Planner appliance was made accessible to Pure Storage, a VMware partner, as part of the Rapid Desktop program. Test bed was configured and Windows 7 desktop base image setup was done in strict adherence to the View Planner installation and user guide document, version 2.1.

The following parameters were tweaked in the View Planner adminops.cfg file for booting more machines at a time:

CONCURRENT_POWERONS_ONE_MINUTE=100 CONCURRENT_LOGONS_ONE_MINUTE=100 RESET_TIMER_PERIOD_IN_SECONDS=1800 POWERON_DESKTOPS=1

Testing Methodology and Success Criteria

Once the View Planner workload generator is run, it produces a View Planner score. This score indicates how many users concurrently running a standardized set of operations a particular virtualization infrastructure platform can support. The standardized view Planner workload consists of nine applications performing a combined total of 44 user operations. The tests are divided into three groups: A, B and C. The final score represents the 95th percentile score of group A operations)used to determine the Quality of Service (QoS)).



Test Results

Pure Storage achieved a score of 0.52 seconds for group A operations. This is dramatically below the 1.5 second response time for a passing score. This score was consistent for both the linked clone runs as well as the full desktop clone runs. We did multiple runs and got near identical results. Figure 16 below is the plot of the average mean response time of View Planner Group A application operations for both linked clone desktops and persistent desktops.



Figure 16: View Planner "Group A" Operations Latency

Figure 17 below shows the Pure Storage GUI dashboard, which shows that 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 100% as shown in Figure 18. We saw no ballooning and no swapping as we had 192 GB of memory on each virtual desktop server host.









Figure 18: vCenter performance data of a single host; CPU and memory utilization

SUMMARY OF RESULTS

- View Planner score of 0.52 achieved for both linked clone and persistent desktops
- 1,000 linked clone desktops created in less than two hours
- 1,000 desktops recomposed in less than 2 hours
- 1,000 desktops booted in 10 minutes and sustained sub-1ms latency and 50,000 IOPS
- Data reduction (deduplication and compression) in excess of 10-to-1 across all desktops

Benefits of Pure Storage FlashArray for View Deployment

The Pure Storage FlashArray fundamentally changes the storage experience from a virtualized data center perspective [Reference 2 talks in-depth on this topic]. Virtual desktop deployment with VMware View is a great use case for 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 half a second application response times with VMware View Planner workload.

Ease of Storage Capacity Management and Storage Provisioning

The Pure Storage GUI and CLI was designed with one goal in mind: simplify storage management. The GUI dashboard is a central location for getting capacity, IOPS/Latency/Bandwidth and system status (see Figure 17 & 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 alert you if the host FC ports are zoned correctly and has redundant connectivity to the FlashArray (see Figure 19 below). This feature is included for all platforms with no agents in the OS.



DASHBOARD S	TORAGE	ANALYSIS	SYSTEM				Search	Hosts	and Vo	olumes				
System Health	Con	nections												
Connections	Host P	Port Connectivity					ст	0			ст	1		
Configuration	HOST /	PORT		CONNECTI	ON STATUS	FC0	FC1	FC2	FC3	FC0	FC1	FC2	FC3	
o en ingui a ten	📖 📼 De	moinfra-Host		Redunda	ant Connections				ų.			1	1	
Users		21:00:00:24:ff:26:69	ee							Ψ.		Ψ.		
		21:00:00:24:ff:26:69:	ef										Ψ.	
	r≡ ES	XHost-001		Redunda	ant Connections					Ψ	Ψ			
		21:00:00:24:ff:26:68	70							Ψ.		Ψ.		
	U	21:00:00:24:ff:26:68	71				Ψ		Ψ		Ψ		Ψ.	
	📼 ES	XHost-002		Redunda	ant Connections	Ψ			ų.	Ψ	Ψ			
		21:00:00:24:ff:27:29	e2							Ψ.		Ψ.		
		21:00:00:24:ff:27:29	e3										Ψ.	
	📼 ES	XHost-003		Redunda	Int Connections	ų.	Ψ		Ψ	ų.	ų.	Ţ	1	
		21:00:00:24:ff:27:29	c0											
		21:00:00:24:ff:27:29	c1										Ψ.	
	📼 ES	XHost-004		Redunda	ant Connections							1	•	
		21:00:00:24:ff:27:2d	14											
	Show	v Hosts Only												
	Target	t Ports												
	PORT	WWN		PORT	WWN									
	CT0.F0	C0 😈 52:4a:93:7	0:00:00:01:00	CT1.FC0	52:4a:93:70:	00:00:00	1:10							
	CT0.F0	C1 52:4a:93:7	0:00:00:01:01	CT1.FC1	52:4a:93:70:	00:00:0	1:11							
	CT0.F0	C2 😈 52:4a:93:7	0:00:00:01:02	CT1.FC2	52:4a:93:70:	0:00:00	1:12							
	CT0.FC	C3 52:4a:93:7	0:00:00:01:03	CT1.FC3	52:4a:93:70:	0:00:00	1:13							

Figure 19: Connections map for hosts connected to Pure Storage FlashArray

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

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 our 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 FlashArray for virtual desktops are broadly classified into three sections and we describe them in detail in the next subsection.

- VDI operational aspects Pool maintenance (Recomposing 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

Recomposing 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 lots 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.

With the FlashArray, we were able to demonstrate a 1,000 desktop patch push out in less than two hours while sustaining 40K IOPS with half millisecond latencies. See Figure 20 below for the FlashArray GUI dashboard view during recomposing 1,000 desktops. We tuned the View composer to perform more concurrently, as mentioned in the View configuration section.

Desktop administrators can recompose their desktops anytime of the day, as the IOPS and latency are not a limiter on the FlashArray. They can even recompose a pool of desktops that is not in use while other pools are actively in use. 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 many folds and the software applications are always up-to-date.



Figure 20: Dashboard view of recomposing 1,000 desktops in less than two hours



Reduced Desktop Boot Times

When using mechanical disk storage arrays, constant pool maintenance activities can lead to unpredictable spikes in storage performance. This is a real project killer in virtual desktop deployment projects. When View admins spin up a pool 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; it creates login and logoff storms respectively.

We simulated the worst-case scenario of powering on 1,000 virtual desktops and measuring the backend IOPS. The Figure 21 below shows the Pure Storage GUI dashboard for this activity. We sustained upwards of 55K IOPS and booted 1,000 virtual desktops in less than 10 minutes while maintaining less than 1 msec latency. This is phenomenal testimony of how the FlashArray can withstand heavy load like boot storms and still deliver sub-millisecond latency.



Figure 21: Booting 1,000 VMs in less than 10 minutes with sustained IOPS upto 50K and < 1msec 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 be 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 rates are going through the roof, this becomes a huge factor in storage total cost of ownership calculations. The overall VDI operational budget needs to 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

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

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 use needs a lot more capacity and IOPS, which is readily available on the FlashArray. Throughout the various phases of this project, we deployed more than 3,000 virtual desktops of mixed types 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 (excepting the user data).

In the next section we talk more about the different configurations of 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 1,000 desktop 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 22 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 that 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.

	?	0	0	с с !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Stage	Pilot	Go Live	Expand	Scale-Up
Users	100s-1,000	Up to 1,000	Up to 2,000	5,000+
Raw Capacity	2.75 TB	2.75 TB	5.5 TB	11 TB
Usable VDI Capacit	y * 10-20 TB	10-20 TB	20-50 TB	50-100 TB





Conclusions

We set out to prove that virtual desktop deployment is an ideal use case for the Pure Storage FlashArray and we achieved unprecedented results while running an industry-standard desktop workload generator. The View Planner score of 0.52 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 >20-to-1 reduction of the VDI workload, enabling the use of either linked clone desktops or full-clone persistent desktops interchangeably, 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?



Acknowledgements

The author would like to thank VMware for providing the View Planner tool through the partner program and reviewing the work leading to this publication. Special thanks to Banit Agarwal in VMware View performance team for his help in troubleshooting View Planner configuration and View tuning. Thanks to Mac Binesh in VMware EUC reference architecture team for his support and providing the VMware View content for this document.

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.

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

References

1. VMware View 5 Performance and Best Practices: <u>http://www.vmware.com/files/pdf/view/VMware-</u> View-Performance-Study-Best-Practices-Technical-White-Paper.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-withpure-storage/

3. DSNRO, the story: <u>http://www.yellow-bricks.com/2011/06/23/disk-schednumreqoutstanding-the-story/</u>

4. VMware View Planner Installation and User Guide – Version 2.1 dated 10/24/2011



APPENDIX A

Pure Storage LUN provisioning

The following example creates a 4 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 4t
purehost create -wwwnlist 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 purehyroup setattr $_addhostlist\ HOSTLIST\ HGROUP$ is used new hosts to the host group.

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

Help Sign Out Welcome pureuser (signed in to vmw-demo)					
DASHBOARD STORAGE	ANALYSIS SYSTEM		Search Hosts and Volumes		Q
- Hosts +	Host Group PureESXCluster		80.00 TB TOTAL PROVISIONED 30.3 to 1 DATA REDUCTION		
PureESXCluster ESXHost-001 ESXHost 002	218 GB RAID-3D			7 VOLUM	62 GB E DATA 980 GB
ESXHost-002	Hosts + 📼	Connected Volumes			÷ + 🔒
r ESXHost-004	NAME	NAME	LUN PF	OVISIONED	VOLUME DATA
ESXHost-005	r ESXHost-001	SVDIVolume-001	10	4 TB	85.85 GB
D≡ ESXHost-006	r ESXHost-002	SVDIVolume-002	11	4 TB	103.14 GB
🖙 ESXHost-008	in the second s	SVDIVolume-003	12	4 TB	98.05 GB
ESXHost-009	p ESXHost-004	SVDIVolume-004	13	4 TB	98.62 GB
P ESXHost-010 S VDIVolume-001	D = ESXHost-005	SVDIVolume-005	14	4 TB	27.61 GB
VDIVolume-002	r ESXHost-006	SVDIVolume-006	15	4 TB	26.11 GB
SVDIVolume-003	r ESXHost-007	SVDIVolume-007	16	4 TB	39.38 GB
VDIVolume-004	In the second s	SVDIVolume-008	17	4 TB	59.71 GB
▼ Volumes +	in ESXHost-009	SVDIVolume-009	18	4 TB	31.95 GB
PemoInfra-Backup	p ESXHost-010	VDIVolume-010	19	4 TB	33.21 GB

The Pure Storage GUI can accomplish this using similar operation.



APPENDIX B

Cisco MDS zoning sample

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

```
# 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:







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