DEPLOYMENT GUIDE

NetApp FAS 2050HA Unified Storage

A Guide to Deploying NetApp FAS 2050HA with VMware View

Contents

1 INTRODUCTION TO VMWARE VIEW

Built on VMware's industry-leading virtualization platform, VMware® View is a Universal Client solution that lets you manage operating systems, hardware, applications and users independently of each other, wherever they may reside. VMware® View streamlines desktop and application management, reduces costs and increases data security through centralization, resulting in greater end user flexibility and IT control. VMware View enables customers to extend the value of VMware Infrastructure and virtual desktop infrastructure (VDI) environments to encompass not only desktops in the datacenter but also applications and the delivery of these environments securely to remote clients, online or off, anywhere.

VMware View transforms the way customers use and manage desktop operating systems. Desktop instances can be rapidly deployed in secure data centers to facilitate high availability and disaster recovery, protect the integrity of enterprise information, and remove data from local devices that are susceptible to theft or loss. Isolating each desktop instance in its own virtual machine eliminates typical application compatibility issues and improves and delivers a more personal computing environment.

2 ABOUT THIS GUIDE

This deployment guide provides a detailed summary and characterization for designing and configuring a NetApp FAS2050HA storage system for use with VMware View and Linked Clones. It describes a validated configuration for a 1000 user workload where 500 desktops are in persistent access mode and 500 users are in non-persistent access mode. This guide can be easily scaled up for larger deployments by simply increasing the number of servers, storage controllers, and storage needed.

The configuration presented was validated in accordance with a recommended architecture as defined in the VMware View Reference Architecture: A Guide to Large-scale VMware View Deployments. This guide is intended to offer Systems Architects and Administrators guidance with the configuration of a NetApp FAS2050HA storage system for use with such VMware View based environments. The information provided in this guide also can be helpful to anyone looking to deploy VMware View with Linked Clones using a NetApp FAS2050HA. In addition, due to all NetApp FAS storage controllers having the same feature and management interface consistency with Data ONTAP®, all of these guidelines can also be directly applied to entry-level and enterprise-class NetApp FAS storage controllers. All of the NetApp best practices applicable to this solution configuration are documented in NetApp's TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices, TR-3705: NetApp and VMware View Solution Guide and TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide and have been used in the creation and validation of this environment.

3 HARDWARE AND SOFTWARE REQUIREMENTS

3.1 TERMINOLOGY

Term	Definition
Aggregate	Pool of physical storage that contains logical volumes
CIFS	Common Internet File System
Deduplication	A technology that seeks out duplicate data, removes the duplicate data and replaces it with a reference pointer to the previously stored, identical object

NFS	Network File System Protocol
NIC	Network Interface Card
RAID	Redundant Array of Independent Disks
Snapshot	Read only copies of an entire file system in Data ONTAP
VMware View Manager	VMware View Manager manages secure access to virtual desktops, works with VMware vCenter Server to provide advanced management capabilities
VC	VMware vCenter Server
VIF	Virtual Interface
VLAN	Virtual Local Area Network
VMware View	A set of software products that provide services and management infrastructure for centralization of desktop operating environments using virtual machine technology.
Volume	Logical storage container on Data ONTAP that organizes user and system files and directories

Table 1: Glossary of terms

3.2 HARDWARE RESOURCES

The following equipment was used in this configuration:

Description	Minimum Revision
One NetApp FAS2050HA Cluster	Data ONTAP 7.3.1; NFS
Two shelves of disks	28 disks (14 per shelf); Each disk 300GB / 15K/ FC
2 Cisco 3750 stackable switches	
1 dual port Ethernet NIC per FAS2050 controller	

Ten Servers (Configured as follows)	
128 GB RAM	
4 Quad Core Xeon Processors	
2 On-board Ethernet NICs	
2 Quad port Ethernet NICs	

Table 2: Hardware Configuration

3.3 SOFTWARE RESOURCES

The following software was used in the configuration:

Description	Minimum Revision
Data ONTAP®	7.3.1
NFS License	N/A
VMware ESX Servers	3.5 Update 3
VMware vCenter Server	2.5 Update 3
Windows Servers for vCenter	Microsoft Windows Server 2003 Enterprise Edition SP 2 (32-Bit)
Desktops/Virtual Machines	Windows XP Service Pack 3 (32-Bit)
VMware Tools	3.5
Windows Server for View connection server	Microsoft Windows Server 2003 Enterprise Edition SP 2 (32-Bit)
Infrastructure servers (AD, DHCP, DNS)	Microsoft Windows Server 2003 Enterprise Edition SP 2 (32-Bit)

Table 3) Software Resources

4 PHYSICAL ARCHITECTURE

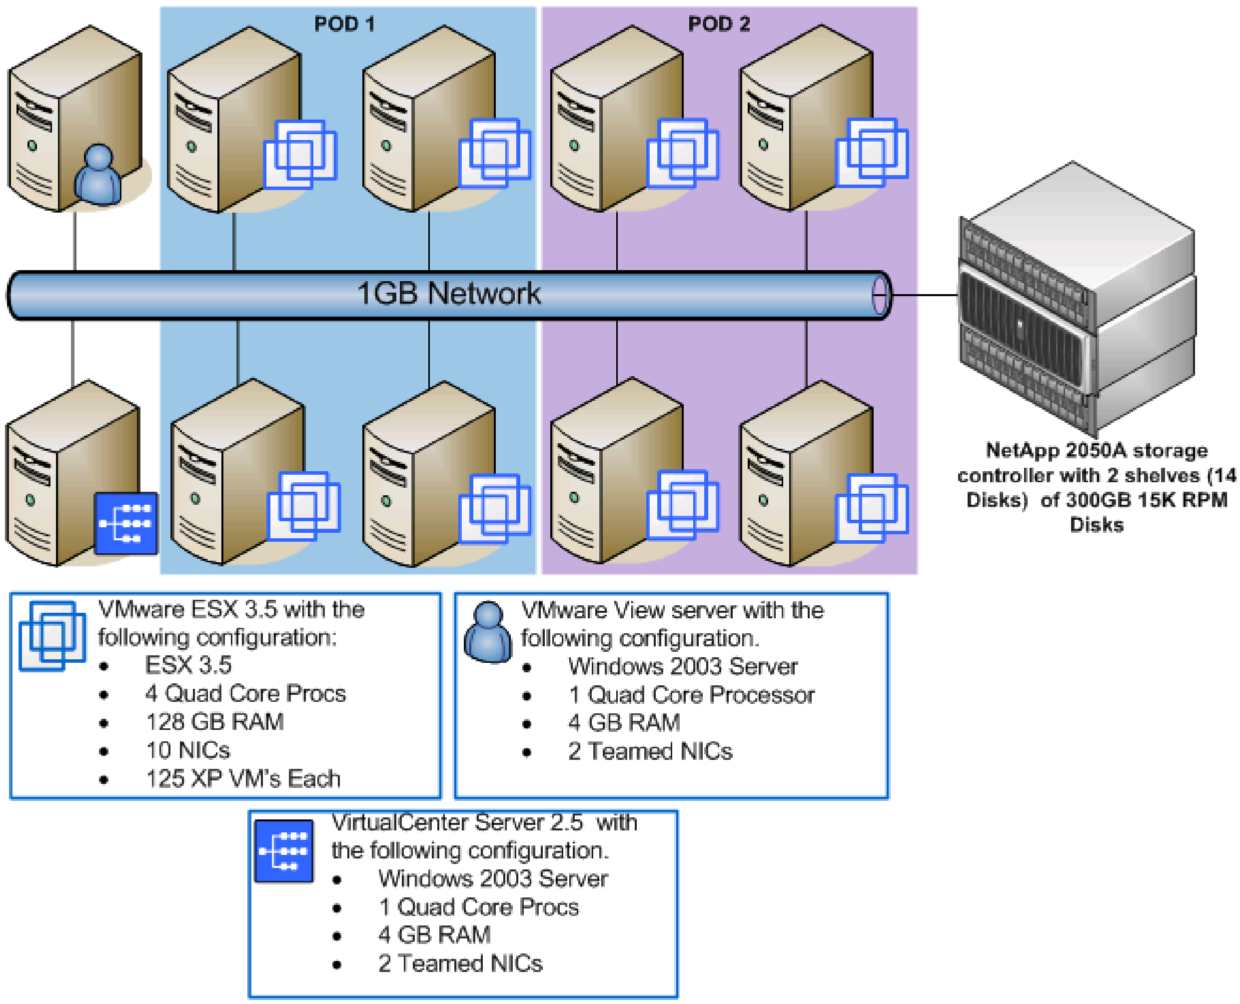

Figure 1) Physical Architecture Overview.

4.1 NETWORK ARCHITECTURE

The networks used for this test were dedicated 1Gb Ethernet. This network was split into three VLANs. One VLAN was for administrative/public traffic and CIFS access, and the other two were non-routable VLANs designed for storage and VMotion traffic. All virtual desktops were assigned an IP address using a DHCP server.

The VMware ESX Server's networking configuration consisted of ten 1Gb Ethernet Controllers (or NICs). Two were configured as NIC Teaming ports for NFS traffic and assigned to a pre-defined VMkernel port used for NFS access with a pre-defined IP address on the non-routable VLAN (i.e. 192.168.1.x). Two NICs were configured as NIC Teaming ports for VMotion traffic and assigned to a pre-defined, VMotion enabled VMkernel port with a pre-defined IP address on another non-routable VLAN used solely for VMotion traffic (i.e. 192.168.0.x). Two NICs were configured as NIC Teaming ports for the Service Console and assigned to the administrative VLAN. The last four NICs were also configured as NIC

Teaming ports and assigned to the public VLAN for access to CIFS shares and other network resource by the desktop virtual machines.

The NetApp FAS2050HA has four 1Gb NICs. The four NICs were configured as two multi-mode VIFs. One VIF was specifically for VMkernel NFS traffic and was placed on the private, non-routable VLAN. The other VIF was for CIFS (Home Directories) and management traffic. This configuration allows for an Active/Active state with both failover and a degree of redundancy in the NFS environment. The two switch ports for the NFS traffic on each storage controller were assigned to a private, non-routable VLAN previously configured and the multi-mode VIF was assigned a pre-defined IP address on this VLAN (i.e. 192.168.1.x). In addition two additional NICs on the public VLAN (with each residing on separate switches) are configured into a multi-mode VIF for CIFS traffic. This configuration can either be done in the Data ONTAP GUI or from the Data ONTAP service Console. Additionally, the multi-mode VIF configured for NFS traffic was also assigned an alias IP address to allow for mounting the NFS datastore to the ESX hosts with different IP addresses to increase throughput and redundancy of the NFS IP link.

Since the Cisco 3750s used in this configuration support cross-stack Etherchannel trunking, each storage controller requires only one physical connection for NFS traffic to each switch. The two ports were combined into one multimode LACP VIF with IP load balancing enabled.

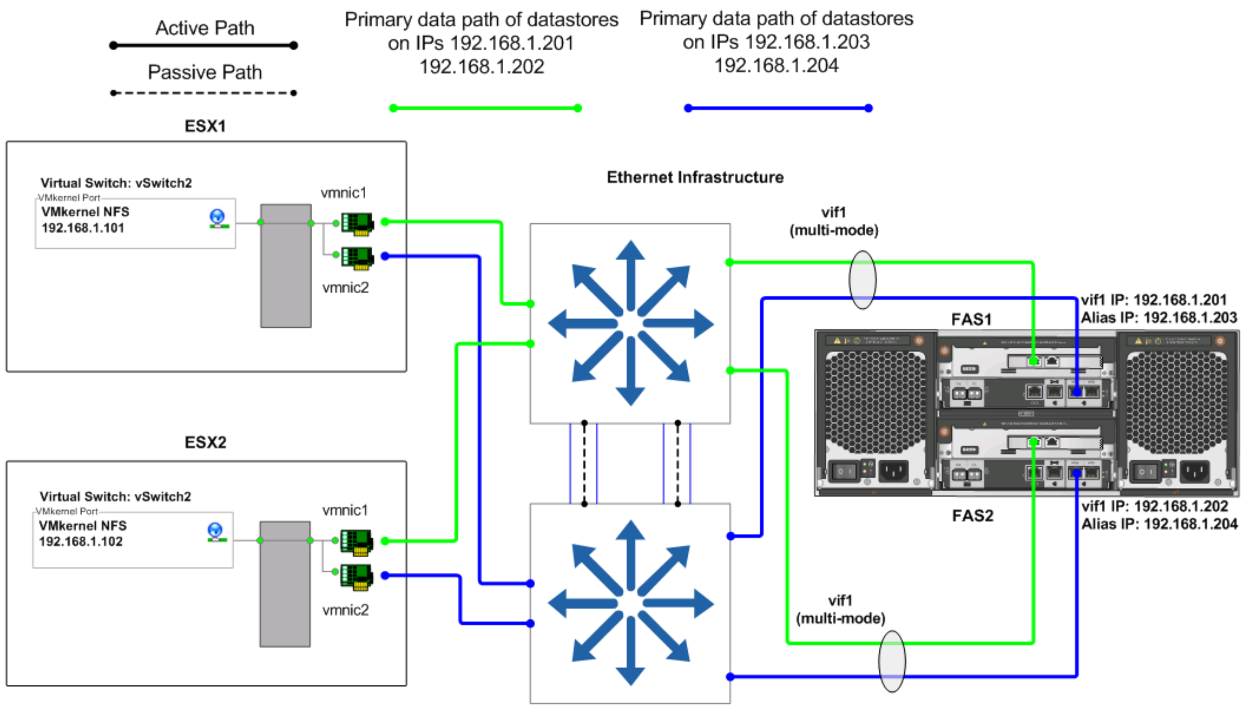

Figure 2) ESX and NetApp storage controller VIF configuration.

Shown below is a step-by-step of how to perform the network configuration for the NFS VIF in FilerView®.

NOTE: In accordance with NetApp best practices for storage controller configuration the creation of the management/CIFS VIF was necessary in order to have redundant NICs for management. The creation of the management/CIFS VIF should be done during the initial setup of the FAS2050 in order to avoid a reboot of the storage controller at a future time. In the event the management/CIFS VIF needs to be created after the initial setup of the FAS2050, please refer to the [Data ONTAP® Software Setup Guide]() for directions in rerunning the setup process for the storage controller.

1. For each NetApp controller, open NetApp FilerView® and click *Network->Add Virtual Interfaces.*
2. Name the new VIF.
3. Select the two network interfaces that will make-up the NFS VIF.
4. For *Trunk Mode* press the drop down arrow and select *Multiple*.
5. For *Load Balancing* press the drop down arrow and select *IP*.
6. Click *Apply*.

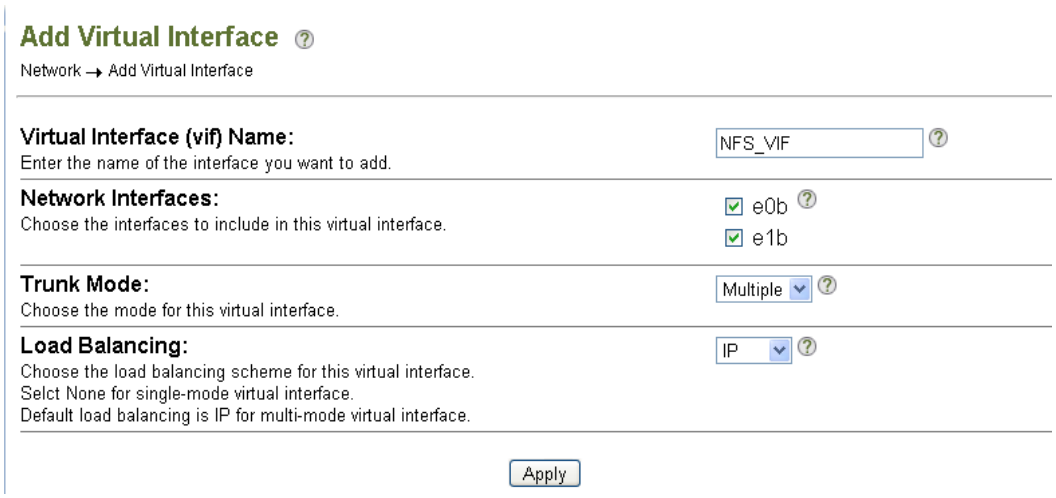

Figure 3) FAS2050 NFS VIF setup –Create Virtual Interface.

1. At the *Manage Network* Interfaces screen, click on *Modify* for the new VIF that was created in the previous steps.
2. At the *Modify Network Interfaces* screen enter a private *IP Address* that will be used for ESX host NFS connectivity and then enter the *Netmask* and click *Apply*.
3. At the *Manage Network Interfaces* screen click on *Up* for the new VIF in order to activate it.

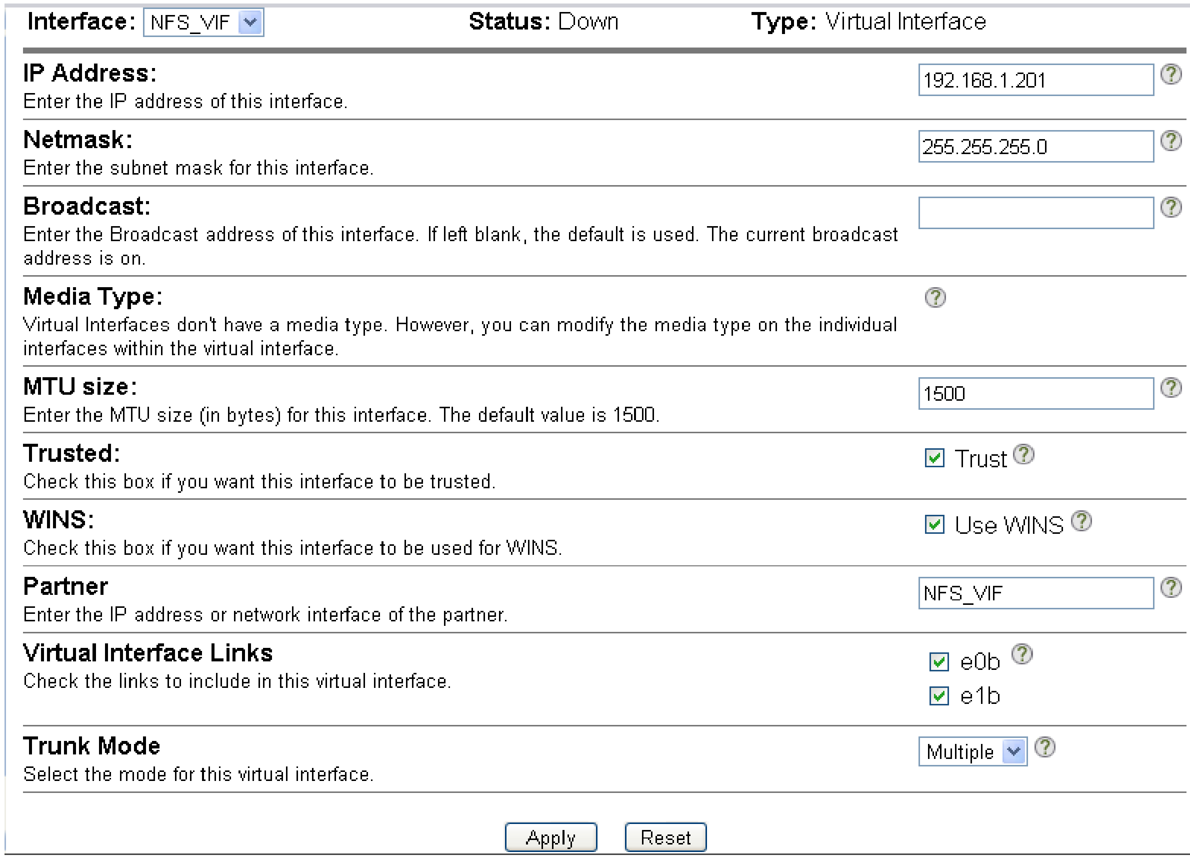

Figure 4) FAS2050 NFS VIF setup – Modify virtual interface.

The ESX hosts will need to be configured as well. For each ESX host, configure NFS connections to the storage controllers using a VMkernel port group. The network ports for this NFS VMkernel port group should be on a non-routable NFS VLAN. For each ESX host, the virtual switch for the NFS VMkernel should have two network ports that each go to different switches in the network. These two switches are stacked together with link aggregation configured for the uplink vmnics.

Note: This should be configured as a standard port channel, not using any link aggregation protocols such as LACP.

To configure an ESX server for NFS connectivity follow the below steps:

- Configuration Tab > Networking > Add networking > VMKernel > Create a New switch
- Provide a Network Label – Example: VMkernel NFS
- Set a VLAN tag (optional)
- Provide the IP address and subnet mask

NOTE: While performing the network setup the VMotion VMkernel port group and a public Virtual Machine port group should be configured as well. Also, ensure that the redundant Service Console NIC is placed in the Service Console port group for redundancy. Depicted below is an example of the networking setup for an ESX host for the environment detailed in this document.

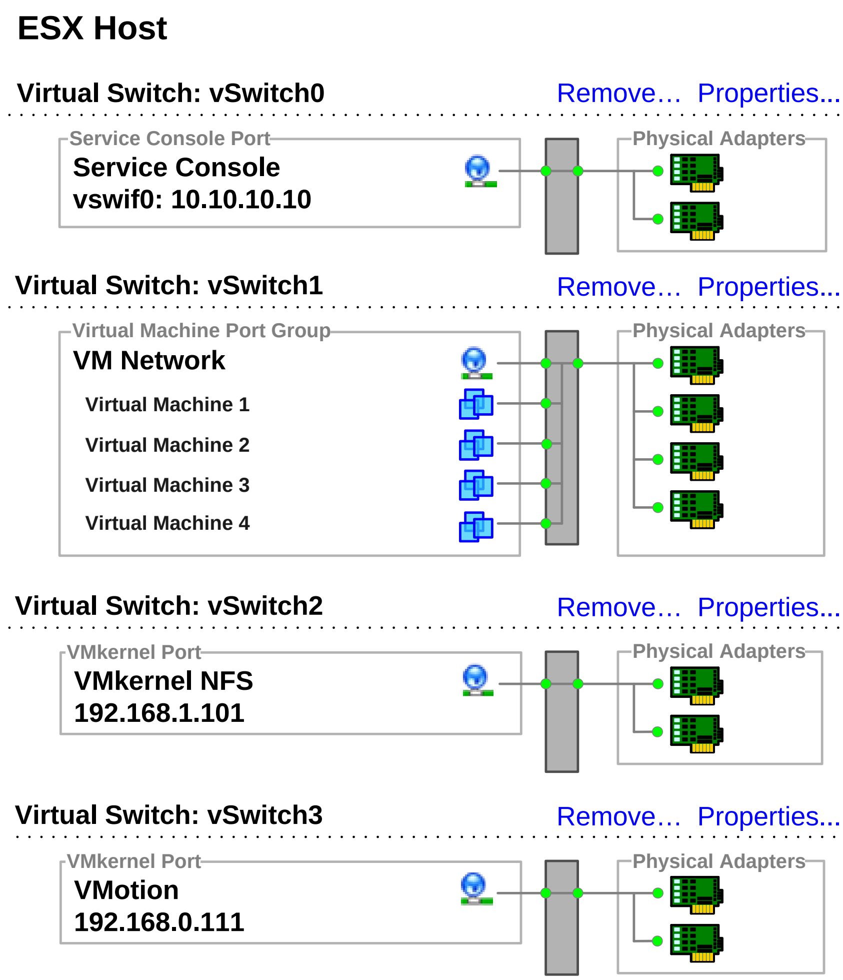

Figure 5) VMware ESX server NIC configuration

4.2 STORAGE ARCHITECTURE

The NetApp FAS2050HA was configured with one aggregate and six volumes per controller. Each storage controller had one shelf of 300GB/15K RPM disks (with 12 disks in an aggregate and two spares). **Using NetApp's RAID-DP to create one 12 disks aggregate enabled all virtual desktops hosted on the storage controller to use the pooled performance capability as required. RAID-DP also allows higher levels of resiliency as compared to RAID-5.**

NetApp Storage Controller Objects	Configuration
Volume for Parent VM	
Initial volume storage capacity	25GB
NFS capacity	25GB

Storage required after deduplication	5GB
Number of volumes used	1 (Single parent VM will be used for both PODs)
Size of VM	10GB
Volumes for Linked Clones OS Data Disks	
NFS volume storage capacity	400GB
Number of volumes used	4 (two per FAS controller)
Volume for User Data Disks for Desktops in Persistent Access Mode (500; 250 per FAS controller)	
NFS volume storage capacity	250GB (2GB per user, 50% deduplication savings)
Number of volumes used	2 (one per FAS controller)
Volume for CIFS User Data for Desktops in Non-Persistent Access Mode (500; 250 per FAS controller)	
CIFS volume storage capacity	250GB (2GB per user, 50% deduplication savings)
Number of volumes used per storage controller	2 (one per FAS controller)

Table 4: Volume configuration

4.3 NETAPP DATA ONTAP CONFIGURATION

The required NetApp FAS2050HA configurations are as follows:

NOTE: These configurations can be implemented using the Data ONTAP® FilerView or service console commands. The instructions for performing the steps below on the NetApp storage controller are for using the FilerView interface. These configurations can be implemented on any model of NetApp FAS storage controllers using the same instructions below. Also, please be sure to follow the best practices for implementing VMware on NetApp storage as recommended in TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices

This configuration requires the creation of four volumes per FAS controller. Two of these volumes will be used as NFS based OS datastores for virtual machine files. One of the volumes will be used as a datastore for hosting the user data disks for persistent desktop. The other will be a CIFS volume for hosting the user data for non-persistent desktops. Please refer to Table 4: Volume Configuration for the correct size to use for each volume.

To configure the volumes that will contain the virtual machines perform the following steps:

1. For each NetApp controller, open NetApp FilerView® and click *Volumes->Add* to activate the Volume Wizard*.*
2. Click on *Next* to go to the Volume Wizard –Volume Type Selection screen.
3. Choose “Flexible” and click *Next*.
4. Enter the new Volume name and choose the language and click *Next*.
5. Choose the Aggregate created for the VMs in the previous steps and select None for the Space Guarantee. Click on *Next*.
6. Choose “Useable Size”, enter the appropriate Volume Size (in GB), and set the Snapshot Reserve to 0%. Click *Next*.
7. Click *Commit*
8. Click *Close*
9. Repeat these steps until all volumes have been created for each storage controller.

NOTE: For scaled-up environments it is best to separate the root-aggregate from the production aggregates and to make non-root aggregates as large as possible to benefit from the I/O capacity of all the spindles in the aggregate.

After the creation of the new aggregate the snapshot reserve will need to be turned off by using the below procedure.

1. Connect to the controller’s console, using either SSH, telnet, or serial console.
2. Set the aggregate snapshot schedule by typing s*nap sched –A <aggregate-name> 0 0 0*
3. Set the aggregate snapshot reserve by typing *snap reserve –A <aggregate-name> 0*
4. Delete existing snapshots, type snap list -A <vol-name>, and then type *snap delete <vol-name> <snap-name>*
5. Log out of the NetApp console

To create the necessary volumes which will serve as an NFS datastore to the ESX servers perform the following steps:

1. Open FilerView
2. Select Volumes.
3. Select Add to open the Volume Wizard.
4. Complete the Wizard, assigning a name to the volumes you create and placing it on the aggregate created earlier.

NOTE: Data ONTAP® creates new volumes with a security style matching that of the root volume. Verify that the security style of the volume is set to UNIX®.

You will next need to ensure that three of the volumes created on the storage controller are set as NFS volumes. In order to accomplish this, please perform the following steps:

1. From the FilerView menu, select NFS.
2. Select Manage Exports to open the Manage NFS Exports screen.
3. Click on the virtual machine production volume created in the previous step NFS export.
4. Grant the export **Root Access** permissions by clicking on Next and placing a green check inside the box. Then click Next.
5. Determine that the export path is correct for the NFS export.
6. At the ***NFS Export Wizard - Read-Write Access,*** click on the Add button and enter the IP address of the NFS VMkernel for the first ESX 3.5 host server. Repeat this step for the VMkernel IP addresses for the other seven hosts until all eight IP addresses have been entered. When this is done, click Next.
7. At the ***NFS Export Wizard – Security*** screen, ensure that Unix Style is selected and click Next.
8. At the ***NFS Export Wizard – Commit*** screen, click **Commit** and at the *NFS Export Wizard – Success* screen, click **Close**.

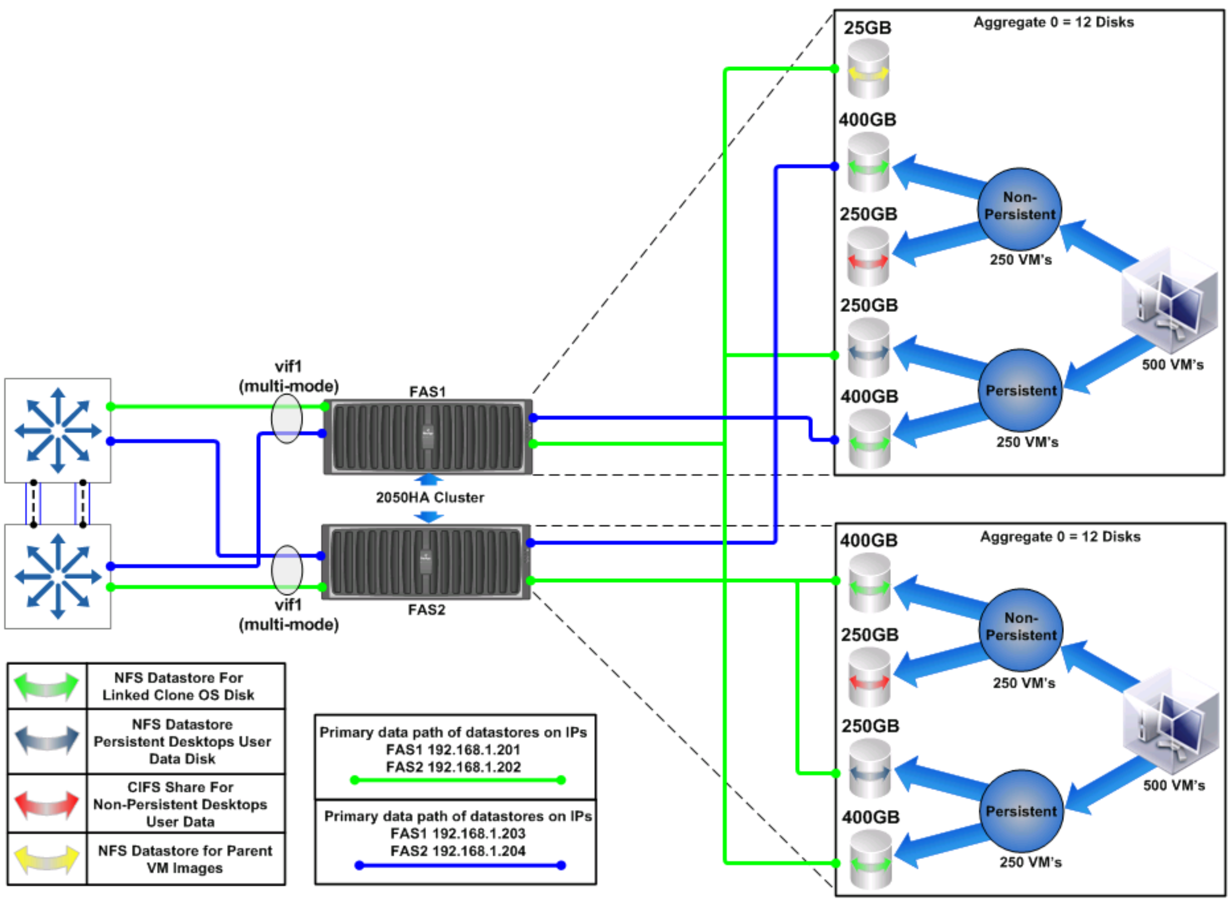

Figure 6) NetApp Datastore Layout.

4.4 INTELLIGENT READ CACHING

VDI by nature is a read-intensive technology and is very bursty in nature. This can be exemplified by boot storms, login storms, and virus scan storms which all further increase the read-intensive and bursty nature of a virtual desktop environment. Traditionally, these challenges are addressed by increasing the cache for both ESX servers and storage devices, increasing the spindle count on the storage devices, and adding more storage devices to offset the total workload.

However, NetApp intelligent caching, available natively in Data ONTAP 7.3.1, combined with VMware Linked clones addresses these challenges. These two technologies work together to help customers reduce their storage footprint for OS data and user data while increasing overall performance in the environment and decreases the overall storage solution costs. VMware Linked Clones and NetApp Intelligent Caching technologies eliminate the requirement for a large number of spindles to serve large-scale read-intensive, bursty VDI workloads.

NOTE: For large-scale environments a Performance Acceleration Module (PAM) can also be used in conjunction with the above technologies for even further performance gains and storage savings.

4.5 DEDUPLICATION SETUP

NetApp deduplication saves space on primary storage by removing redundant copies of blocks within a volume. This process is transparent to the application and can be enabled and disabled on the fly. In a VMware View environment using Linked Clones, this solution provides great value when we consider that all users in the environment will have their own user data either on the “user data disk” (for persistent desktops) and/or CIFS home directories (non-persistent desktops). In many environments user data is duplicated multiple times as various identical copies and version of documents and files are saved. For more information refer to NetApp TR-3505: NetApp Deduplication for FAS, Deployment and Implementation Guide.

Two separate User Data Disk and CIFS volumes are setup for user data for both persistent and non-persistent desktops (one volume per controller) respectively. These volumes should have deduplication enabled. To do this, please perform the following steps:

1. Connect to the controller’s system console, using either SSH, telnet, or serial console.
2. Execute the following command to enable NetApp dedupe for the gold volume:
   sis on <gold volume path>
3. Execute the following command to start processing existing data:
   sis start –s <gold volume path>
4. Execute the following command to monitor the status of the dedupe operation:
   sis status

NOTE: Savings of up to 50% have been seen in environments using deduplication for CIFS home directories. Such storage savings is also possible for the “user data disks” for persistent desktops as well.

4.6 VMWARE ESX CONFIGURATION

Each VMware ESX server should be configured to maximize NFS throughput and capability. In order to maximize NFS capabilities please perform the following steps.

1. Open VCenter Server.
2. Select an ESX host.
3. In the right pane, select the Configuration tab.
4. In the Software box, select Advanced Configuration
5. In the pop-up window, left pane, select NFS
6. Change the value of NFS.HeartbeatFrequency to 12.
7. Change the value of NFS.HeartbeatMaxFailures to 10
8. Change the value of NFS.MaxVolumes to 32
9. In the pop-up window, left pane, select Net
10. Change the value of Net.TcpIpHeapSize to 30
11. Change the value of Net.TcpIpHeapMax to 120
12. Repeat the steps for each ESX Server

Alternatively, the following commands can be run in the service console of the ESX host server in order to maximize the capabilities of NFS.

Connect to the ESX host's system console using either SSH, telnet, or serial console and log in to the console. Type each command below and hit Enter.

```
esxcfg-advcfg -s 32 /NFS/MaxVolumes
esxcfg-advcfg -s 12 /NFS/HeartbeatFrequency
esxcfg-advcfg -s 10 /NFS/HeartbeatMaxFailures
esxcfg-advcfg -s 30 /Net/TcpIpHeapSize
esxcfg-advcfg –s 120 /Net/TcpIpHeapMax
```

As stated above, the Cisco 3750 switches support multi-switch Etherchannel trunking, or virtual port channeling. Therefore, only one VMkernel port with one IP address for NFS traffic for the each of the ESX hosts is required.

On each ESX host a vSwitch was specifically created for IP storage connectivity. Two physical adapters were created configured for this vSwitch with each adapter connected to a different physical switch. The switch ports were configured into a cross-stack Etherchannel trunk. One VMkernel port was created and configured with an IP address. Finally, the NIC teaming properties of the VMkernel port were configured with Load Balancing set to "Route based on IP hash."

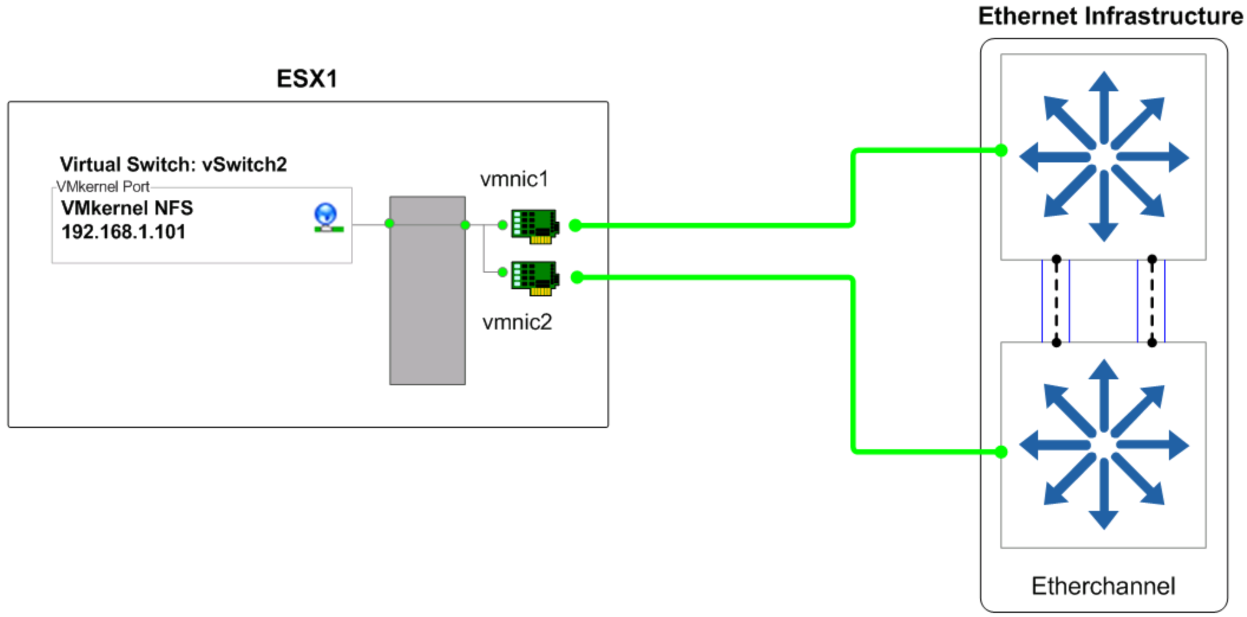

Figure 7) ESX Host NFS Network Configuration

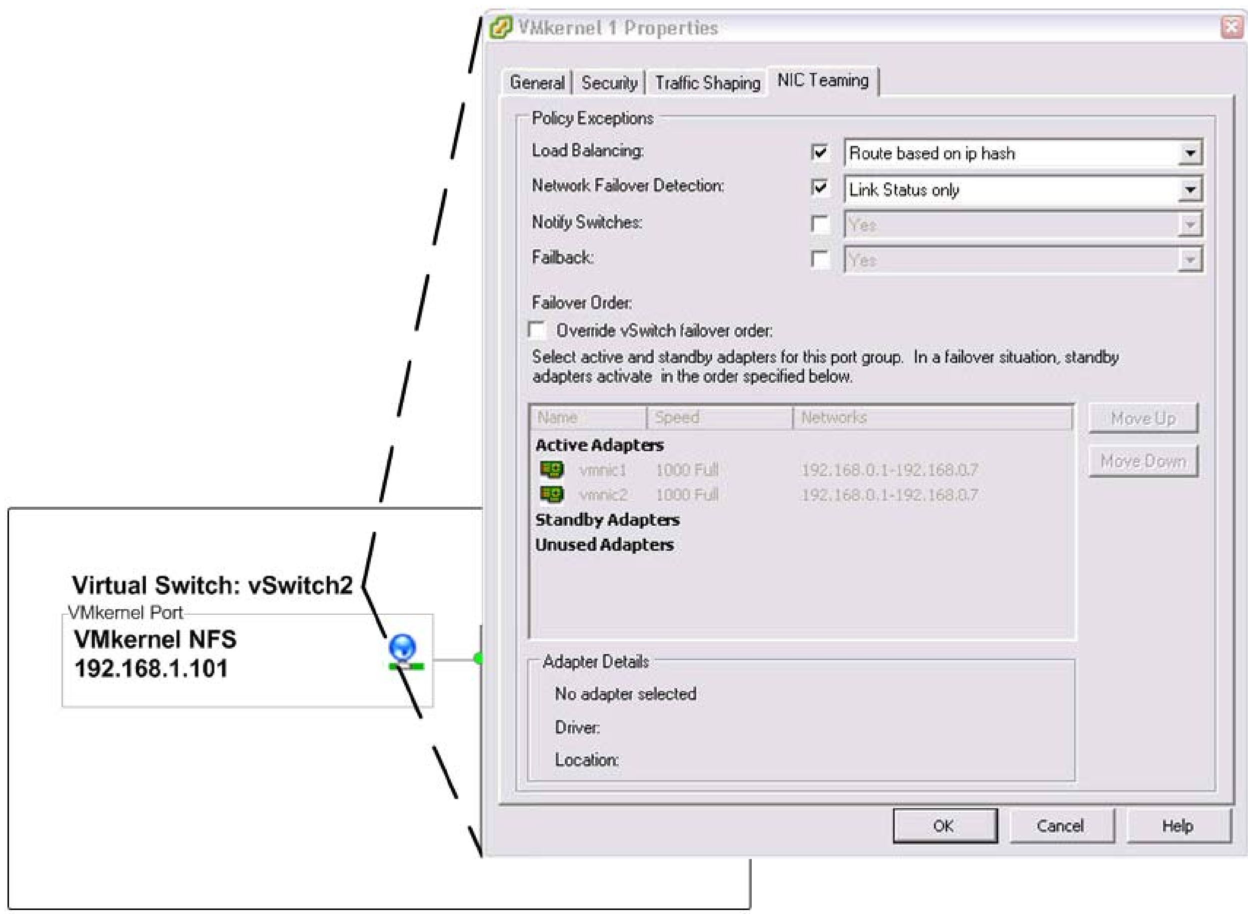

Figure 8) ESX Host NFS Load Balancing Network Configuration

The NFS datastores will need to be mounted on each of the ESX hosts. Please follow the below steps to mount NFS datastores on the ESX host server.

1. Open FilerView (http://filer/na_admin).
2. Select Volumes.
3. Select Add to open the Volume Wizard. Complete the Wizard.
4. From the FilerView menu, select NFS.
5. Select Add Export to open the NFS Export Wizard. Complete the wizard for the newly created file system, granting read/write and root access to the VMkernel address of all ESX hosts that will connect to the exported file system.
6. Open VCenter Server.
7. Select an ESX host.
8. In the right pane, select the Configuration tab.
9. In the Hardware box, select the Storage link.
10. In the upper right corner, click Add Storage to open the Add Storage Wizard.
11. Select the Network File System radio button and click Next.
12. Enter a name for the storage appliance, export, and Datastore, then click Next.
13. Click Finish.

For CIFS traffic, configure the VM Network vswitch (with 4 vmnic uplinks as shown) to span multiple physical switches, with two NICs connected to each switch. When connecting to each network switch, use ports on different network cards on the server to achieve high levels of availability.

5 VALIDATION RESULTS

The storage configuration described in this document was validated by configuring the environment described above and then performing a real world workload. A FAS2050HA (Active / Active cluster) was used to validate a 1000 seat VMware View building block architecture. The ESX infrastructure consisted of eight 16 core servers supporting 128 VMs per ESX server. All VMs were simultaneously booted and the workload was executed on all 1000 virtual desktops for the duration of a 12 hour period. This workload was based on a power user and included typical software components such as Microsoft Internet Explorer, Microsoft Word, Microsoft Excel, Microsoft Power Point, Adobe Acrobat and Firefox. The data found in this section validates that the architecture described above is capable of providing an excellent end user experience.

5.1 APPLICATION RESPONSE TIME

The graph below shows the average application execution time for all 1000 virtual desktops for the duration of the validation. All VMs were booted and the workload was simultaneously started on all 1000 guests and data was collected for a 12 hour period. The graph represents the response time of the applications and hence the end user experience that one would expect to have.

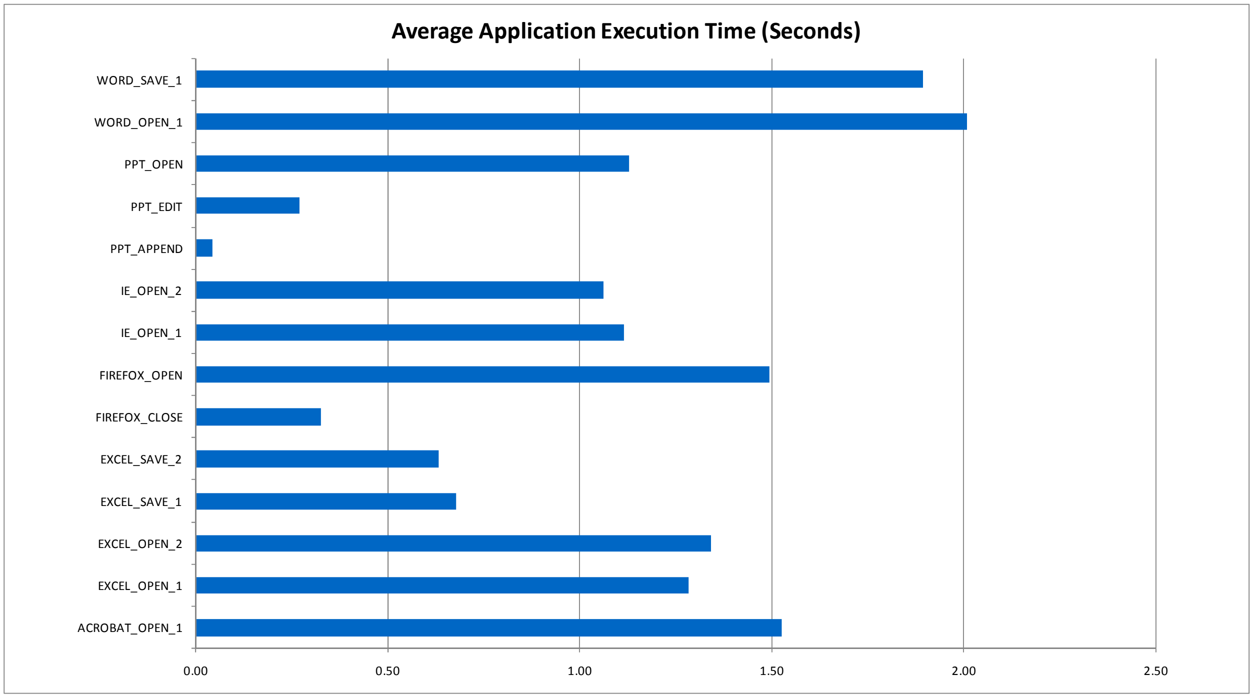

Figure 9) Average application execution time in seconds.

5.2 STORAGE SYSTEM IO SUMMARY

The graphs below show the number of NFS operations for each controller during the duration of the validation. The average number of IOPS for the cluster was 2652 IOPS or 10MB/sec. This demonstrates that each controller has sufficient capability to handle the workload for the building block.

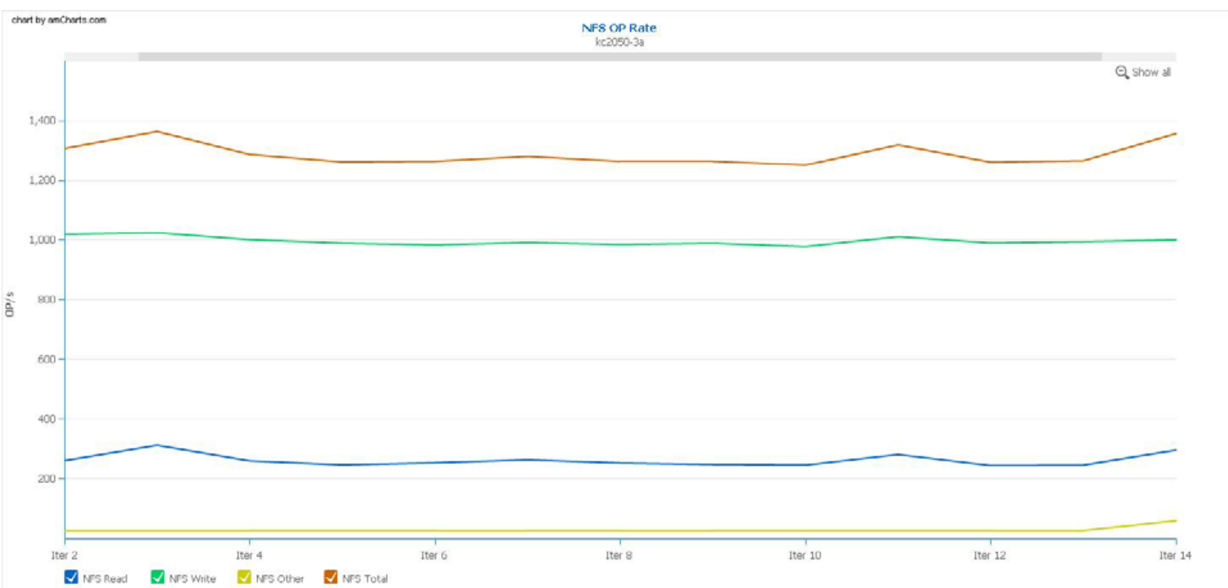

Figure 10) NetApp FAS2050 controller A average NFS OP Rate

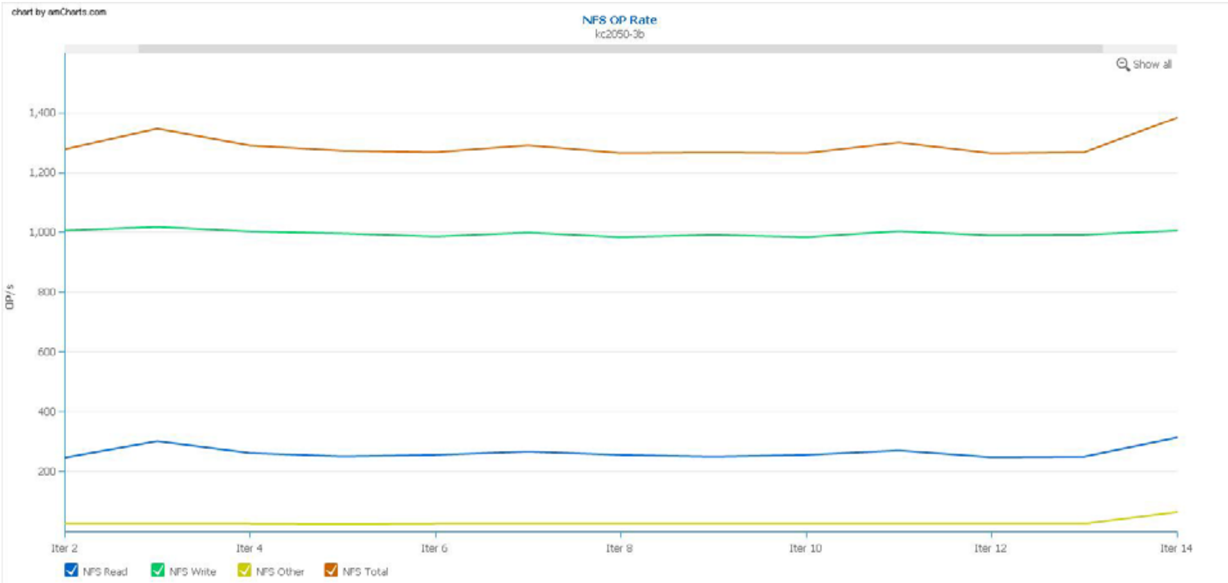

Figure 11) NetApp FAS2050 controller B average NFS OP Rate

6 SUMMARY

This reference architecture deployment guide details and validates a 1000 seat VMware View building block based architecture on NetApp. This building block based approach can be scaled linearly by adding additional building blocks. Both NetApp and VMware provide consulting services, design, architecture, deployment, and management guidelines to assist in the deployment of the solution based on your requirements. Most of this information can be found in the NetApp and VMware references below.

7 ACKNOWLEDGMENTS

The following people contributed to the solution validation and creation of this document:

Wen Yu - VMware

Fred Schimscheimer - VMware

Jack McLeod - NetApp

Abhinav Joshi - NetApp

Chris Gebhardt - NetApp

8 REFERENCES

TR-3428: NetApp and VMware Virtual Infrastructure 3 Storage Best Practices

TR-3705: NetApp and VMware View (VDI) Solution Guide

TR-3505: NetApp Deduplication for FAS Deployment and Implementation Guide

VMware View Manager Administration Guide

Introduction to View Manager

VMware Infrastructure 3 Documentation

VMware View Windows XP Deployment Guide

VMware View Composer Deployment Guide

VMware View Reference Architecture

Revision: 20090417 Item: DG-092-PRD-01-01

VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com