NETWORK APPLIANCE AND VMWARE VIRTUAL INFRASTRUCTURE 3 STORAGE BEST PRACTICES

M. Vaughn Stewart and Michael Slisinger, Network Appliance, Inc.
October 2007 | TR3428 | Version 3.1

Abstract

This document discusses the virtual storage solutions that reduce cost, increase storage utilization, increase fault tolerance, and address the challenges of backing up and restoring VMware ESX Server environments using Network Appliance™ technology.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Network Appliance technology enables companies to extend their virtual infrastructures to include advanced virtualized storage alongside their virtual servers. NetApp provides industry-leading solutions in the areas of data protection; ease of provisioning; thin storage provisioning; file-based backups; instantaneous virtual machine (VM) backup and restores; instantaneous VM cloning for testing, application development, and training purposes; and simple and flexible disaster recovery options.

This technical report reviews the best practices for implementing VMware Virtual Infrastructures on Network Appliance fabric-attached storage (FAS) systems. NetApp has been providing advanced storage features to VMware ESX solutions since the product began shipping in 2001. During that time, NetApp has developed operational guidelines for the FAS systems and ESX Server. These techniques have been documented and are referred to as best practices. This technical report describes them.

Note: These practices are only recommendations, not requirements. Not following these recommendations will not affect whether your solution is supported by Network Appliance. Not all recommendations apply to every scenario. NetApp believes that their customers will benefit from thinking through these recommendations before making any implementation decisions.

The target audience for this paper is familiar with concepts pertaining to VMware ESX Server 3.0 and Network Appliance Data ONTAP® 7.X. For additional information and an overview of the unique benefits that are available when creating a virtual infrastructure on NetApp storage, see http://www.netapp.com/library/tr/3515.pdf.

CONFIGURATION AND SETUP

VMWARE STORAGE OPTIONS

Three types of storage options are available to VMware Virtual Infrastructure 3 (VI3). The following sections summarize the unique benefits of each option.

VMFS Datastore on Fibre Channel or iSCSI

This is historically the most common method for deploying storage in VMware environments. The strengths of this solution are that it is well known, and once storage has been provisioned to the ESX Servers, the VMware administrator is free to use the storage as needed. Most operations are run exclusively through VMware VirtualCenter.

The shortcomings of this solution are that performance degrades as multiple virtual machines are deployed on the same datastore; scalability requires multiple datastores; and finding performance bottlenecks can be difficult. Figure 1 shows an example of this configuration. Note that each datastore has its own I/O queue to the NetApp FAS system. The storage network and I/O queue are shared among all virtual machines residing on the datastore. For information on accessing virtual disks stored on a Virtual Machine File System (VMFS) by using either FCP or iSCSI, see the VMware Server Configuration Guide.

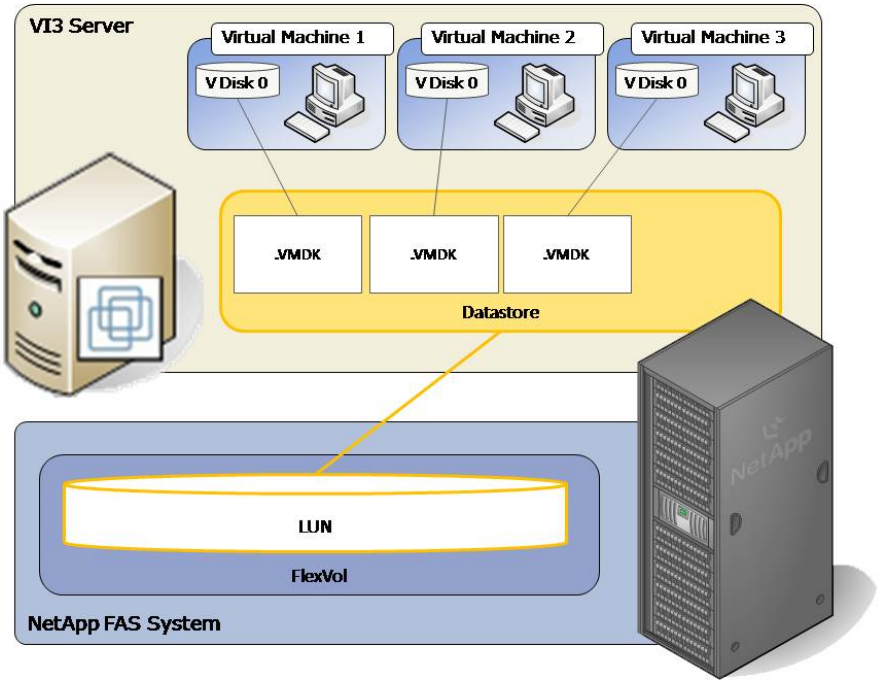

Figure 1

Raw Device Mapping over Fibre Channel or iSCSI

Raw device mapping (RDM) was introduced in VMware ESX Server 2.5. The strengths of this solution are high disk I/O performance; easy disk performance measurement; support for virtual machine host- based clustering (such as Microsoft® Cluster Server, MSCS); and easy integration with features of advanced storage systems.

The shortcomings of this solution are that VMware data centers may have to be limited in size and may require more interaction between storage and VMware administrators. Figure 2 shows an example of this configuration. Note that each virtual disk file has its own I/O directly to a LUN, which in turn is associated directly with a virtual machine. Although the storage network is shared, the storage I/O queue is not. For more information on raw device mappings over Fibre Channel and iSCSI, see the VMware Server Configuration Guide.

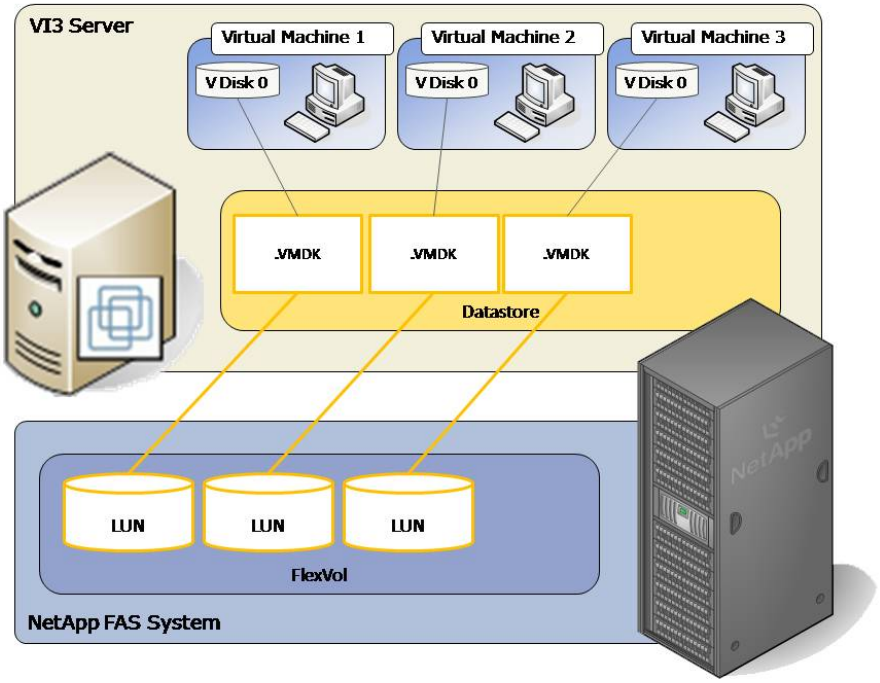

Figure 2

NFS Datastore

Support for storing virtual disks (VMDKs) on a Network File System (NFS) was introduced in VMware ESX Server 3.0. The strengths of this solution are that once storage has been provisioned to the ESX Servers, the VMware administrator is free to use the storage as needed; per port costs are lower than with a Fibre Channel solution; and VMDK files are thin provisioned by default, providing increased utilization of total storage capacity. In addition, this design is easy to integrate with features of advanced storage systems and provides high disk I/O performance. Figure 3 shows an example of this configuration. Note that each virtual disk file has its own I/O queue to the NetApp FAS system. For more information on storing VMDK files on NFS, see the VMware Server Configuration Guide.

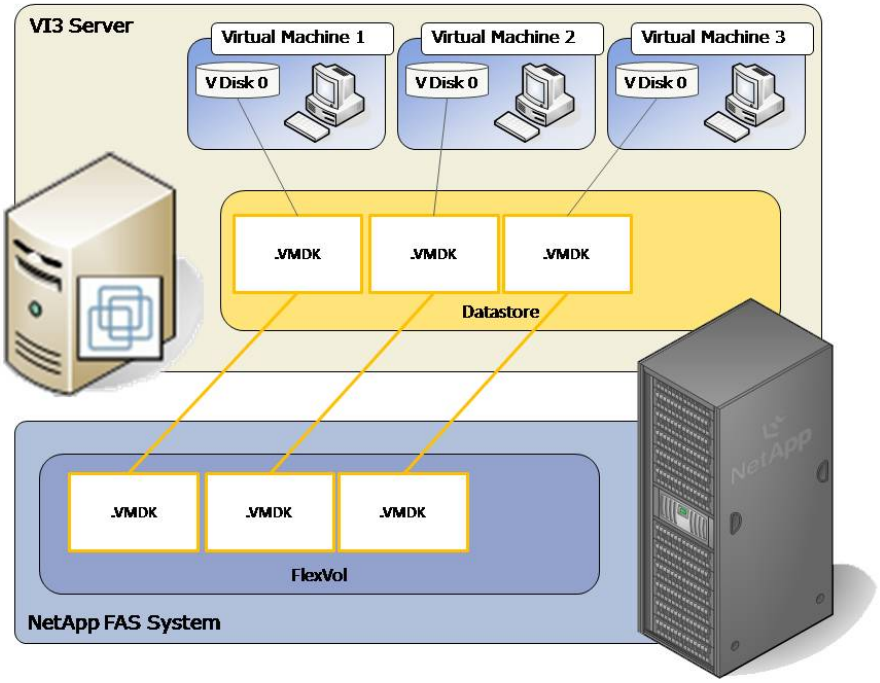

Figure 3

FAS SYSTEM STORAGE CONFIGURATION

RAID Data Protection

The amount of data addressed in today's enterprise is staggering in size compared to what it was just two or three years ago. Due to this dramatic growth, data protection has become a paramount feature of shared storage devices. Data protection in a virtual infrastructure is even more critical than in a traditional server infrastructure, because a storage failure in a VI leads to multiple virtual machines experiencing downtime and/or losing data. NetApp RAID-DP™ is an advanced RAID technology that is provided as the default RAID level on all FAS systems. RAID-DP provides protection from the simultaneous loss of two drives in a single RAID group, or more realistically the failure of a single disk drive and a media error on one of the remaining drives. This level of resiliency makes data residing on RAID-DP safer than data stored on RAID 10 or RAID 50. Because the failure of a RAID group in a virtual infrastructure would result in data loss for multiple virtual machines, NetApp recommends that RAID-DP be utilized on all RAID groups that store VMware data.

Aggregates

An aggregate is simply a collection of RAID groups across which flexible volumes are striped. NetApp recommends that whenever possible a small aggregate be created to store the files required for running the FAS system and the remaining storage be placed into a small number of large aggregates. VMware disk I/O is random by nature, and in ensure optimal performance, a large number of physical spindles should be made available to service requests. For more information, see the ESX Server 3.x Storage/SAN Compatibility Guide.

Volumes

Volumes, and specifically flexible volumes, contain LUNs or virtual disk files that are accessed by VMware ESX Servers. Virtual machines should be grouped and volumes should be created based on service levels and the accompanying Snapshot™ backup and/or SnapMirror® replication policies required to meet them. Because virtual machines can be hosted by any server in a VMware data center (which is a collection of ESX Servers) the datastore or storage is mapped to every node.

LUNs

LUNs are units of storage provisioned from a FAS system directly to the ESX Server. LUNs can be accessed by the ESX Server in two ways. The first and most common method is as storage to hold virtual disk files for multiple virtual machines. This type of usage is referred to as a Virtual Machine File System (VMFS) LUN. The second method is with raw device mapping (RDM). With RDM, the LUN is connected to the ESX Server and is passed directly to a virtual machine to use with its native file system, such as NTFS or EXT3.
As previously stated, VMFS LUNs are the traditional method for providing storage to virtual machines. VMFS LUNs provide simplicity because storage provisioned to the ESX Server can be utilized by the VMware administrator without intervention from the storage administrator. In addition, the VMware administrator can leverage the built-in storage functionality of the ESX Server, such as VMware snapshots and clones.

RDM LUNs were introduced in ESX Server 2.5 and provide several benefits not found with VMFS LUNs. RDMs allow virtual machines to be clustered with Microsoft Cluster Server, and they provide increased disk I/O performance, more specific disk performance measurement, and simpler integration with advanced storage system features such as instantaneous virtual machine recovery via NetApp SnapRestore® or FlexClone® technology. The downside of RDMs is that each VMware data center has a limit of 256 LUNs (which can lead to artificial capacity limitation of a VMware data center), and VMware administration requires an ongoing relationship with storage administration.

For more information, see the ESX Server 3.x Storage/SAN Compatibility Guide.

Storage Naming Conventions

NetApp storage systems allow human naming conventions. A well-planned virtual infrastructure implements a descriptive naming convention that aids the identification and mapping of multiple layers of storage with virtual machines. A simple and efficient naming convention also facilitates configuration of replication and disaster recovery processes.
Consider the following suggestions:

- Volume name: Matches data center name or data center name and replication policy or data type (for example, DC1_Swap, DC_4hr_repl)
- LUN name: For VMFS, NetApp suggests matching the name of the datastore
- LUN name: For RDMs, should be the hostname and volume name of the VM (for example, for Windows®, hosta_c_drive.lun, for Linux® hostb_root.lun, etc.)

Following are two examples of provisioned LUNs using descriptive naming conventions.

This is an example of a volume that stores a VMDK LUN for Data Center 2, which is replicated every 24 hours:

```
/vol/DC_2_24hr_repl/Datastore_1.lun
```

This is an example of a volume that stores RDMs for Data Center 4. In this example, the C drive LUN for a virtual machine named SQLSRVR12 has been created:

```
/vol/DC_4/SQLSRVR12_C_drive.lun
```

VIRTUAL INFRASTRUCTURE 3 CONFIGURATION BASICS

CONFIGURATION LIMITS AND RECOMMENDATIONS

When sizing storage, you should be aware of the following limits and recommendations.

NetApp Volume Options

NetApp volumes should be created as flexible volumes, with the snap reserve set to 0 and the default Snapshot schedule disabled. All NetApp Snapshot copies must be coordinated with the ESX Servers to ensure data consistency. NetApp Snapshot copies are covered in “Implementing Snapshot Copies,” later in this report. To set the volume options for Snapshot copies to the recommended setting, enter the following commands on the FAS system console.

1.	Log into the NetApp console.
2.	Set the volume Snapshot schedule: *snap sched <vol-name> 0 0 0*
3.	Set the volume Snapshot reserve: *snap reserve <vol-name> 0*

RDMs and Data Center Sizing
Currently, a VMware data center is limited to a total of 256 LUNs. This limitation typically matters only in deployments that use RDMs as the primary form of VM storage. With RDMs, you must plan for an additional pair of VFMS LUNs to store RDM pointer data and virtual machine VMX configuration files. This leaves a total of 254 LUNs available. NetApp recommends that you keep RDM pointer files and virtual machine configuration files on separate data stores.
To determine the number of ESX nodes that will be used by a single data center, you use the following formula:

> 254 / (number of RDMS per VM) / (planned number of VMs per ESX host) = number of ESX nodes in a data center

For example, if you plan to run 20 VMs per ESX Server and would like to assign 2 RDMs per VM, the formula is:

> 254/20/2 = 6.4 rounded up = 7 ESX Servers in the data center

RDM mapping files appear to be the same size as the LUN to which they point; however, each mapping file is approximately 1MB in size. NetApp recommends that you create a separate VMFS LUN to hold the mapping files and to format this LUN with the smallest available VMFS block size.

VMFS LUN Sizing
VMFS LUNs/datastores provide the simplest method of provisioning storage; however, you should plan to limit the total size of the LUN/datastore and to create multiple VMFS datastores to avoid disk I/O performance issues due to contention with file locking and SCSI reservations. Although there is no definitive recommendation, a commonly deployed size for a VMFS LUN/datastore is between 300 and 700GB. The maximum LUN size is 2TB. For more information, see the ESX Server 3.x Storage/SAN Compatibility Guide.

NFS Datastore Limits
By default, VMware ESX allows 8 NFS datastores; this limit can be increased to 32. For larger deployments, NetApp recommends that you increase this value to the maximum. To make this change, follow these steps from within the Virtual Infrastructure client. For more information refer to: NFS Mounts are Restricted to 8 by Default

1	Open VirtualCenter.
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, select NFS in the left pane.
6	Change the value of NFS.MaxVolumes to 32. See Figure 4.
7	In the pop-up window, select Net in the left pane.
8	Change the value of Net.TcpIpHeapSize to 30.
9	Repeat for each ESX server

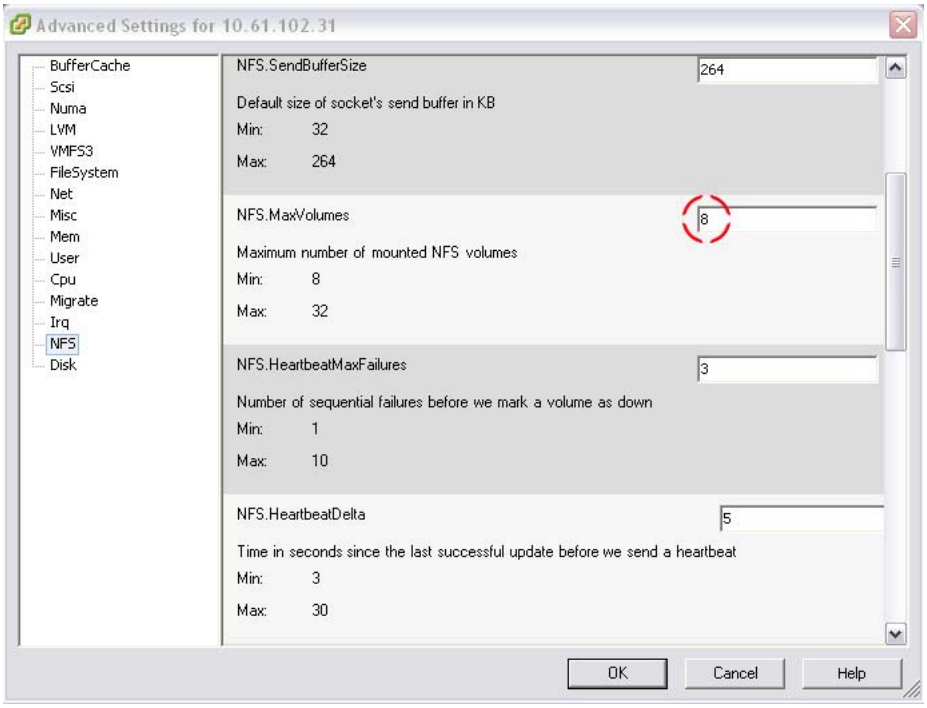

Figure 4

NFS Volume Options
When deploying VMDKs on NFS one should disable the access time updates which occur by default on NFS file systems. To make this change, follow these steps from within the FAS system console.

1	Log into the NetApp console
2	From the storage appliance console, run: *vol options <vol-name> no_atime_update on*

In addition when utilizing NetApp snapshots with NFS datastores an additional NFS setting is required to be set on each ESX server. To make this change, follow these from within the Virtual Infrastructure Client.

1	Open VirtualCenter.
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, select NFS in the left pane.
6	Change the value of NFS.Lock.Disable to 1
7	Repeat for each ESX server

Virtual Disk Starting Partition Offsets
Virtual machines store their data on virtual disks, and these disks are formatted with a file system that enables them to store that data. When formatting a virtual disk, it is important to ensure that the file systems of the VMDK, the datastore, and the storage array are in proper alignment. Misalignment of the file system can result in degraded performance. However, even if file systems are not optimally aligned, performance degradation may not

be experienced or reported, based on the I/O load relative to the ability of the storage arrays to serve the I/O and the overhead for being misaligned. Every storage vendor experiences this challenge. For details, see the VMware publication Recommendations for Aligning VMFS Partitions.

When aligning the partitions of virtual disks for use with NetApp FAS systems, the starting partition offset must be divisible by 4096. The recommended starting offset value is 32768. For Windows guest operating systems, verifying this value is easy. Run msinfo32, and you will typically find that the VM is running with a default starting offset value of 32256. See Figure 5. Msinfo32 can be run by selecting Start > All Programs > Accessories > System Tools > System Information.

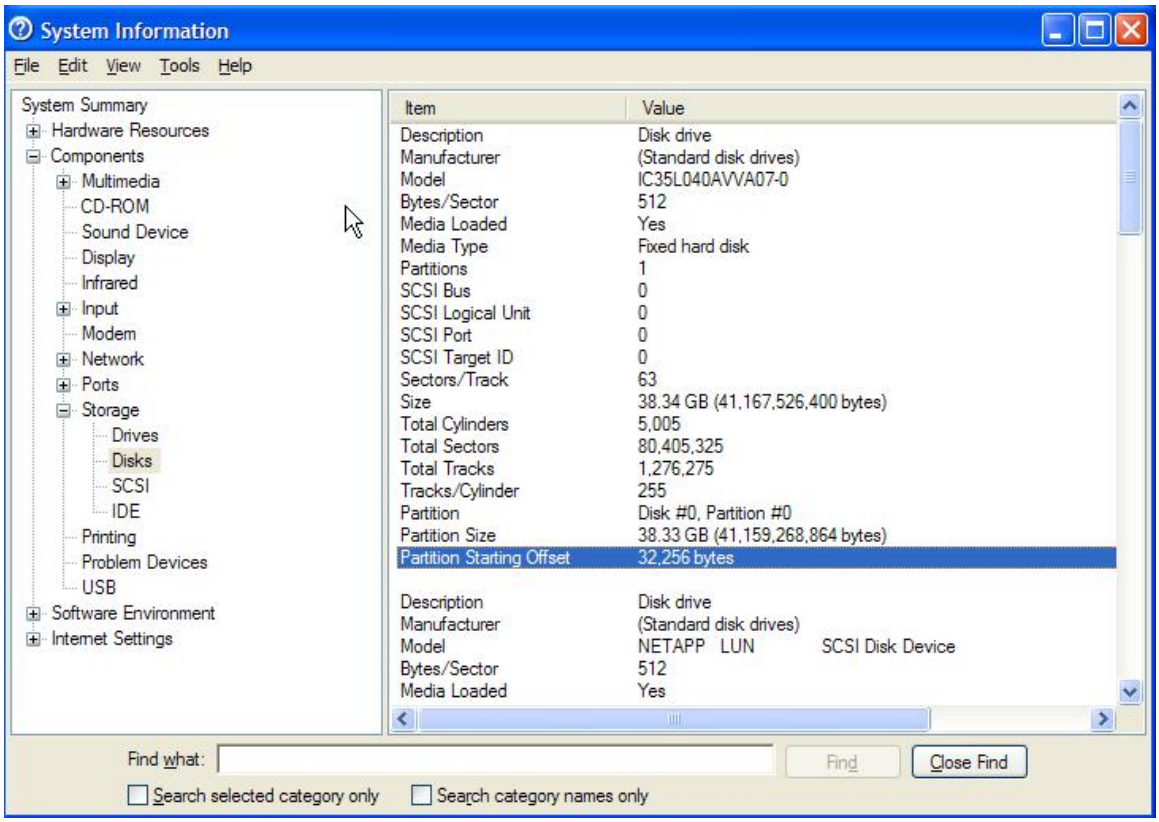

Figure 5

Correcting the starting offset is best addressed by correcting the template from which new VMs are provisioned. For currently running VMs that are misaligned, NetApp recommends correcting the offset of only VMs that are experiencing an I/O performance issue. This performance penalty should be more noticeable on systems that are completing a large number of small read and write operations. The reason for this recommendation is that in order to correct the partition offset, a new virtual disk has to be created and formatted, and the data from the VM has to be migrated from the original disk to the new one. Misaligned VMs with low I/O requirements do not benefit from these efforts.

Formatting with the Correct Starting Partition Offsets

Virtual disks can be formatted with the correct offset at the time of creation by simply booting the VM before installing an operating system and manually setting the partition offset. For Windows guest operating systems, the Windows Preinstall Environment boot CD is an excellent tool. To set up the starting offset, follow these steps and see Figure 6.

1.	Boot the VM with the WinPE CD.
2.	Select Start > Run and enter Diskpart.

3.	Enter Select Disk0.
4.	Enter Create Partition Primary Align=32.
5.	Reboot the VM with WinPE CD.
6.	Install the operating system as normal.

Figure 6

STORAGE PROVISIONING

With the introduction of VMware Virtual Infrastructure 3.0, several new storage options were introduced. This section covers provisioning storage for Fibre Channel, iSCSI, and the Network File System (NFS).

Fibre Channel and iSCSI LUN Provisioning

To provision LUNs for access via FCP or iSCSI, you begin by creating initiator groups (igroups) on the FAS system. NetApp igroups provide a form of LUN masking that controls host access to a LUN. NetApp recommends that an igroup be created for each VMware data center. In addition, NetApp suggests including in the name of the igroup the name of the data center and the protocol type (for example, DTW_DC1_FCP and DTW_DC1_iSCSI). This naming convention and method simplify the management of igroups by reducing the total number created. It also ensures that all ESX Servers in the data center see each LUN at the same ID. Each initiator group includes all of the FCP worldwide port names (WWPNs) or iSCSI qualified names (IQNs) of each of the ESX Servers in the VMware data center.

Note: If a data center will use both the Fibre Channel and iSCSI protocols, then separate igroups must be created for Fibre Channel and iSCSI.

For assistance in identifying the WWPN or IQN of the ESX Server, select storage adapters from the configuration tab for each ESX Server in VirtualCenter and refer to the SAN Identifier column. See Figure 7.

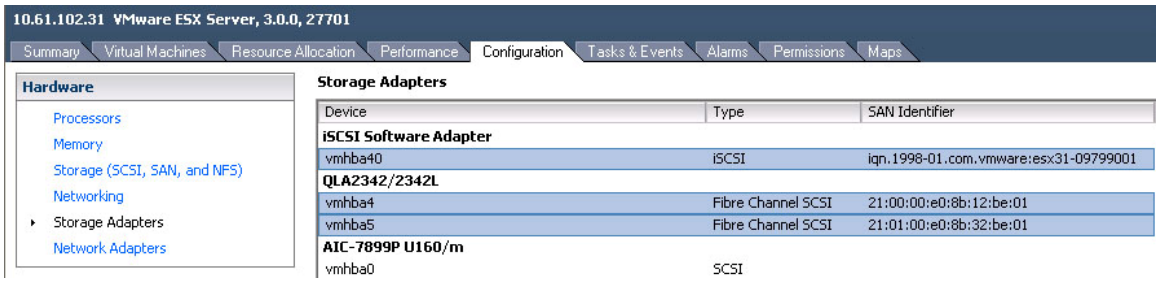

Figure 7

LUNs can be created by using the NetApp LUN wizard in the FAS system console or by using the FilerView® GUI. The following procedures demonstrate creating a LUN using the FilerView GUI.

1.	Log into FilerView.
2.	Select LUNs.
3.	Select Wizard.
4.	In the Wizard window, click Next.
5.	Enter the path. See Figure 8.
6.	Enter the LUN size.
7.	Enter the LUN type (for VMFS select VMware; for RDM select the VM type).
8.	Enter a description and click Next.

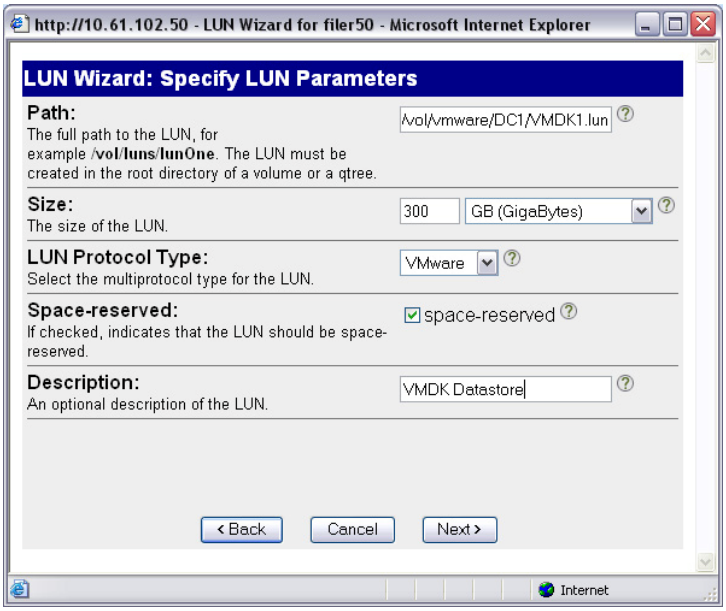

Figure 8

The next step in the LUN wizard is LUN masking. LUN masking is accomplished by assigning an igroup to a LUN. With the LUN wizard, you can either assign an existing igroup or create a new igroup. To configure LUN masking on a LUN created in the FilerView GUI, follow these steps.

1.	Select Add Group.
2.	Select the Use Existing Initiator Group radio button. Click Next and proceed to step 3a. **OR** Select the Create a New Initiator Group radio button. Click Next and proceed to step 3b.
3a.	Select the group from the list, and either assign a LUN ID or leave the field blank (the system will assign one). Click Next to complete the task.
3b.	To create a new group, supply the igroup parameters including name, connectivity type (FCP or iSCSI), and OS type (VMware), and then click Next. See Figure 9.

4.	Enter the new or select the known SAN Identifiers (WWPN or IQN) for the systems that will connect to this LUN.
5.	Click the Add Initiator button.
6.	Click Next to complete the task.

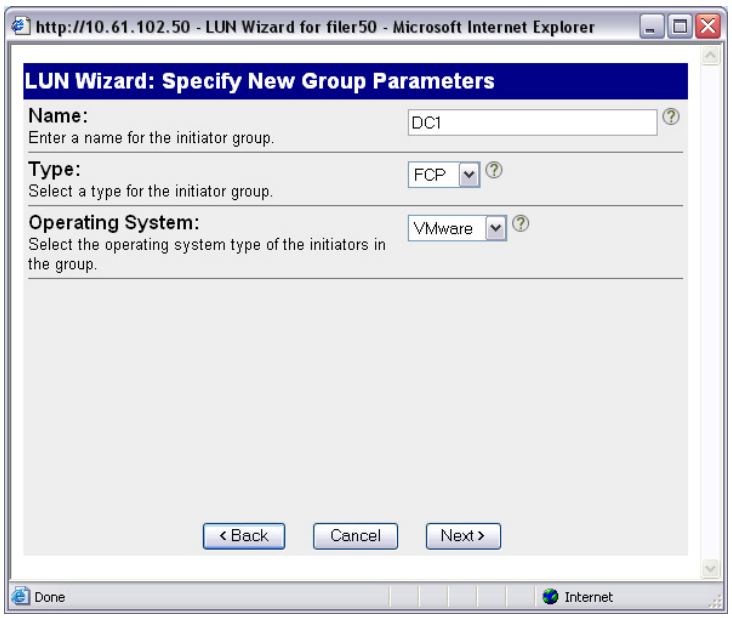

Figure 9

NETWORK FILE SYSTEM (NFS) PROVISIONING

If you would rather serve virtual disks via NFS, the process is simple. To create a file system for use as an NFS data store, follow these steps.

1	Open FilerView (http://filer/na_admin).
2	Select Volumes.
3	Select Add to open the Volume Wizard. See Figure 10. Complete the Wizard.
4	From the FilerView menu, select NFS.
5	Select Add Export to open the NFS Export Wizard. See Figure 11. 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 VirtualCenter.
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. See Figure 12.
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. See Figure 13.
13	Click Finish.

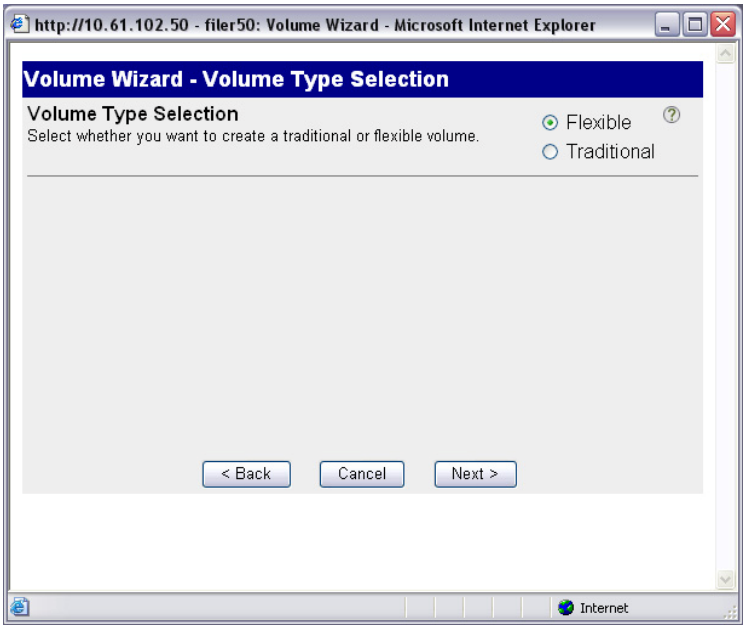

Figure 10

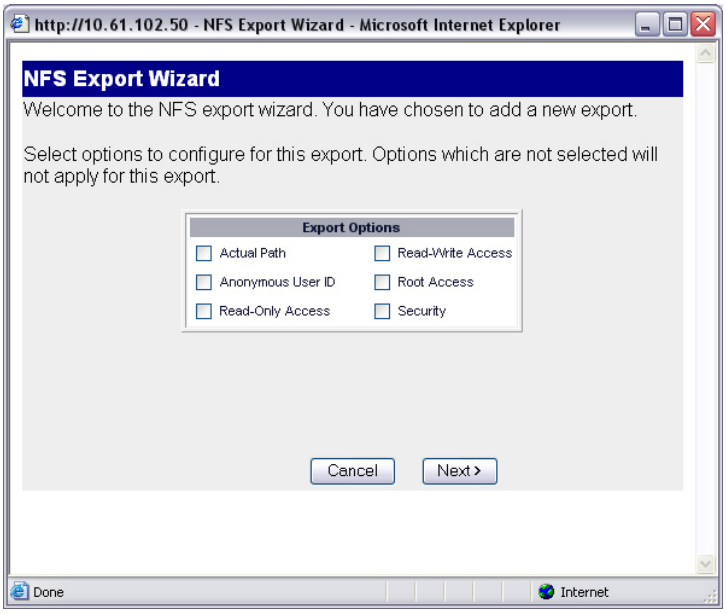

Figure 11

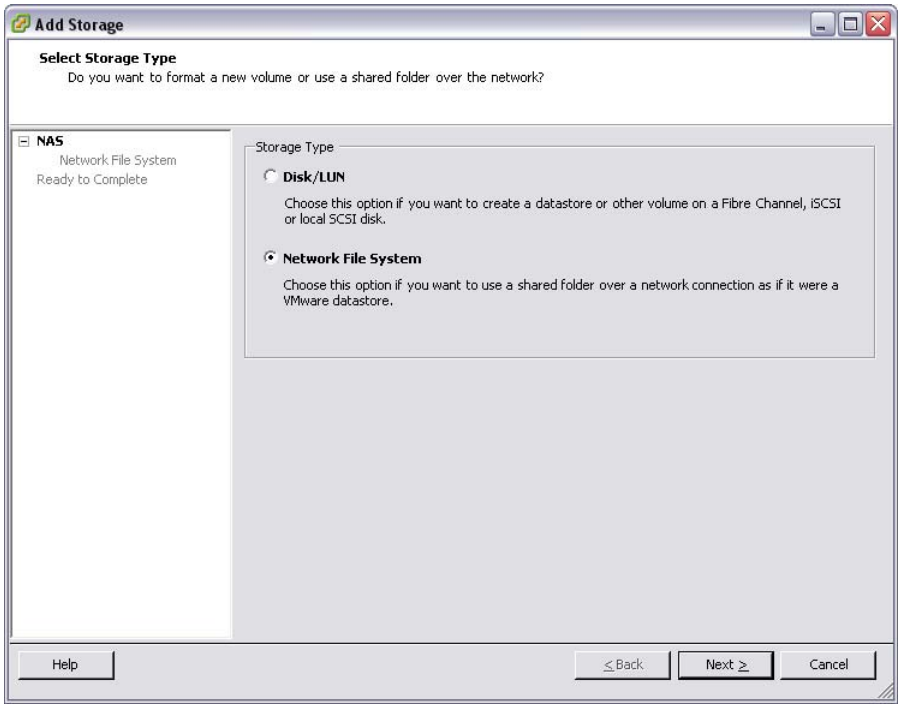

Figure 12

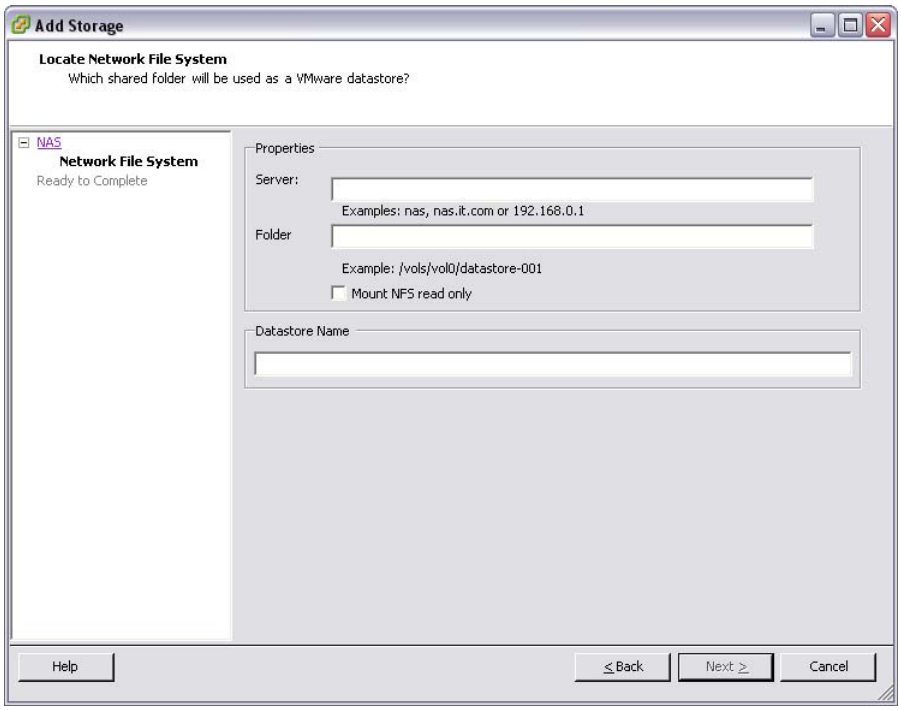

Figure 13

STORAGE CONNECTIVITY

With the introduction of VMware Virtual Infrastructure 3.0, several new storage connectivity options were introduced. This section covers the available storage options and reviews settings specific to each technology.

Fibre Channel Connectivity

To begin, you may notice that the Fibre Channel service is the only storage protocol that is running by default on the ESX Server. NetApp recommends that each ESX Server have two FC HBA ports available for storage connectivity, or at a minimum one FC HBA port and an iSCSI (software- or hardware- based) port for storage path redundancy. To connect to FC LUNs provisioned on a FAS system, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host.
3	In the right pane, select the Configuration tab.
4	In the Hardware box, select the Storage Adapters link.
5	In the upper right corner, select the Rescan link.
6	Repeat steps 1 through 5 for each ESX Server in the data center.

Selecting rescan forces the rescanning of all HBAs (FC and iSCSI) to discover changes in the storage available to the ESX Server.

Note: Some FCP HBAs require you to scan them twice to detect new LUNs (see VMware KB1798 at http://kb.vmware.com/kb/1798). After the LUNs have been identified, they can be assigned to a virtual machine as raw device mapping or provisioned to the ESX Server as a datastore.

To add a LUN as a datastore, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host.
3	In the right pane, select the Configuration tab.
4	In the Hardware box, select the Storage link and then click Add Storage to open the Add Storage Wizard. See Figure 14.
5	Select the Disk/LUN radio button and click Next.
6	Select the LUN you want to use and click Next.
7	Enter a name for the datastore and click Next.
8	Select the block size, click Next, and click Finish.

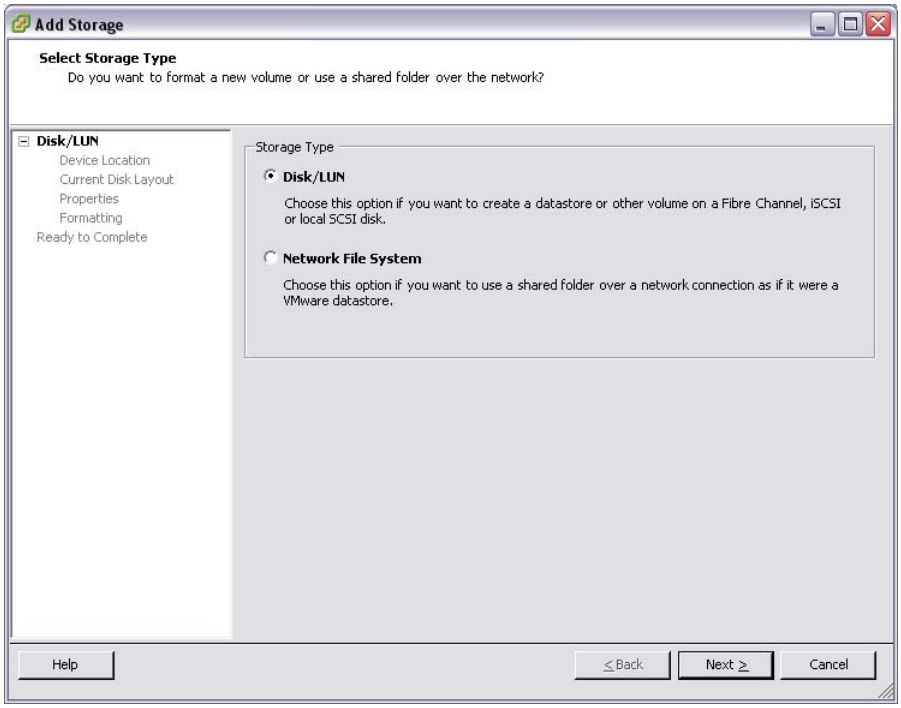

Figure 14

Note that the default block size of a Virtual Machine File System is 1MB. This block size supports storing virtual disk files up to a maximum of 256GB in size. If you plan to store virtual disks larger than 256GB in the datastore, you must increase the block size to be greater than the default. See Figure 15.

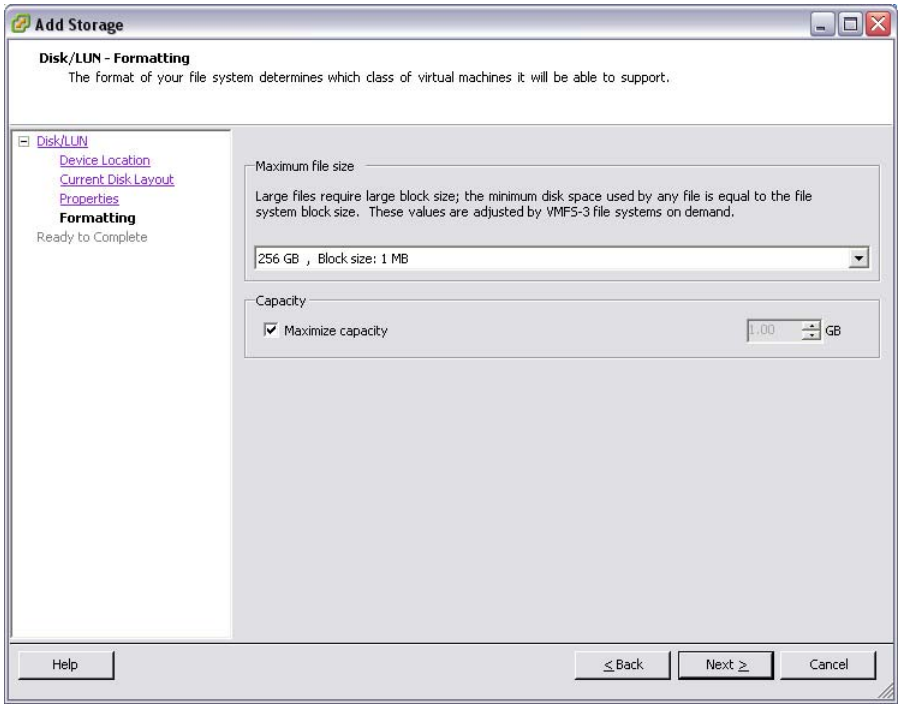

Figure 15

iSCSI/IP SAN Connectivity

As a best practice, NetApp recommends separating iSCSI traffic from other IP network traffic by implementing a separate network or VLAN than the one used for VMotion or virtual machine traffic. To enable ISCSI connectivity, the ESX Server requires a special connection type, referred to as a VMkernel port, along with a service console port. The VMkernel network requires an IP address that is currently not in use on the ESX Server. To configure the iSCSI connectivity, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host.
3	In the right pane, select the Configuration tab.
4	In the Hardware box, select Networking.
5	In the upper right corner, click Add Networking to open the Add Networking Wizard. See Figure 16.
6	Select the VMkernel radio button and click Next.
7	Either select an existing vSwitch or create a new one. **Note:** If a separate iSCSI network does not exist, create a new vSwitch.
8	Click Next.
9	Enter the IP address and subnet mask, click Next, and then click Finish to close the Add Networking Wizard. See Figure 17.
10	In the left pane of the Configuration tab select Security Profile.
11	In the right pane, select the Properties link to open the Firewall Properties window.
12	Select the Software iSCSI Client check box and then click OK to close the Firewall Properties window. See Figure 18.
13	In the right pane, select Storage Adapters in the Hardware box.

14	Highlight the iSCSI Adapter and click the Properties link in the Details box. See Figure 19.
15	Select the Dynamic Discovery tab in the iSCSI Initiator Properties box.
16	Click Add and enter the IP address of the iSCSI-enabled interface on the NetApp FAS system. See Figure 20.
17	For an additional layer of security, select the CHAP tab to configure CHAP authentication. NetApp recommends setting up and verifying iSCSI access before enabling CHAP authentication.

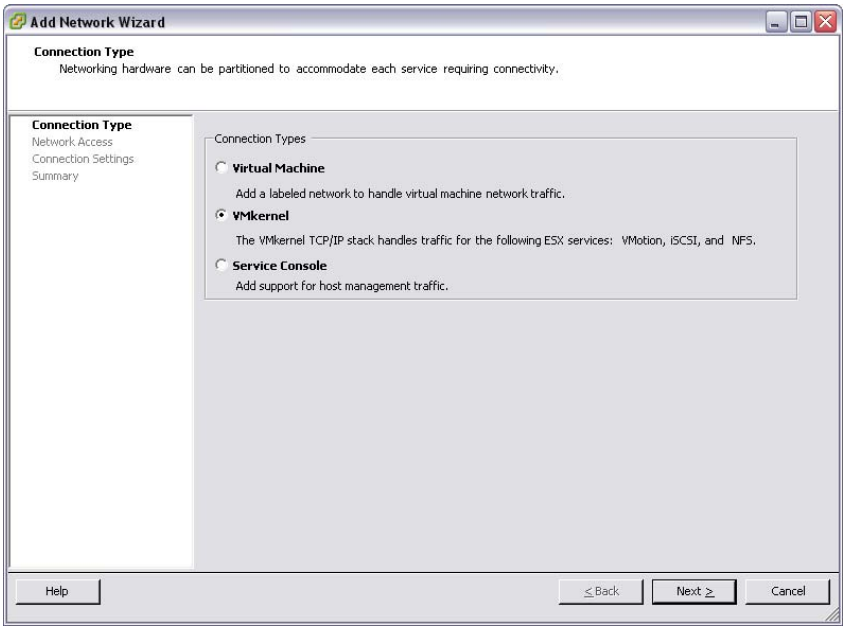

Figure 16

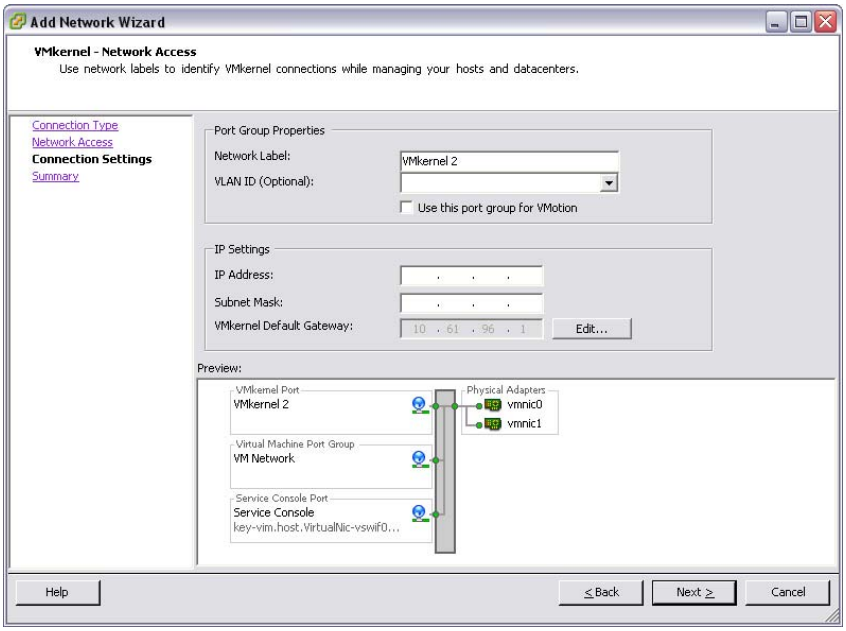

Figure 17

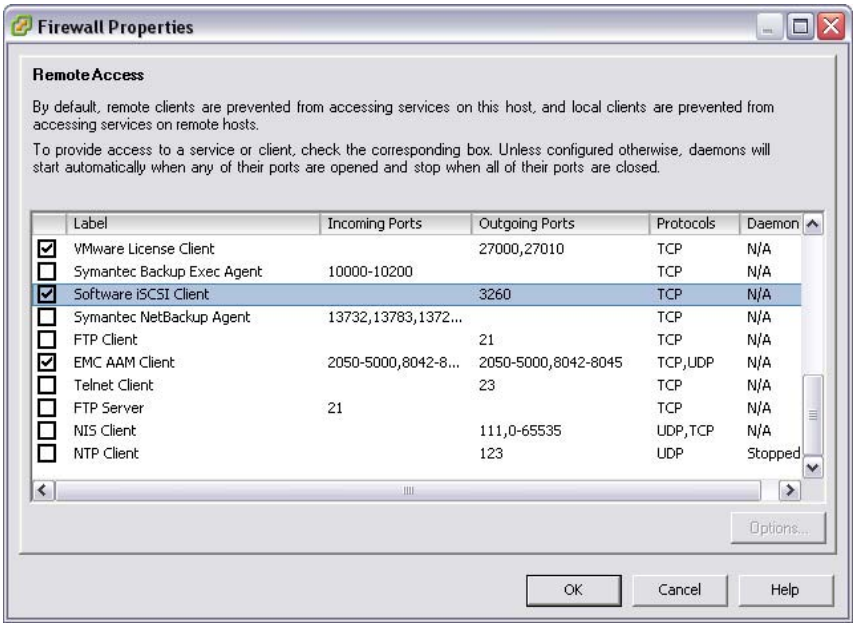

Figure 18

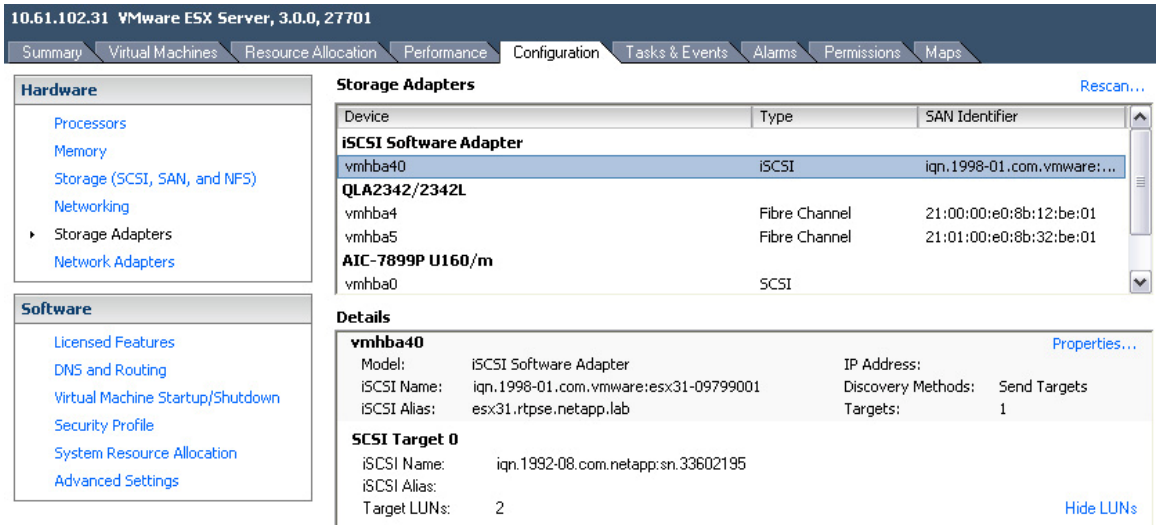

Figure19

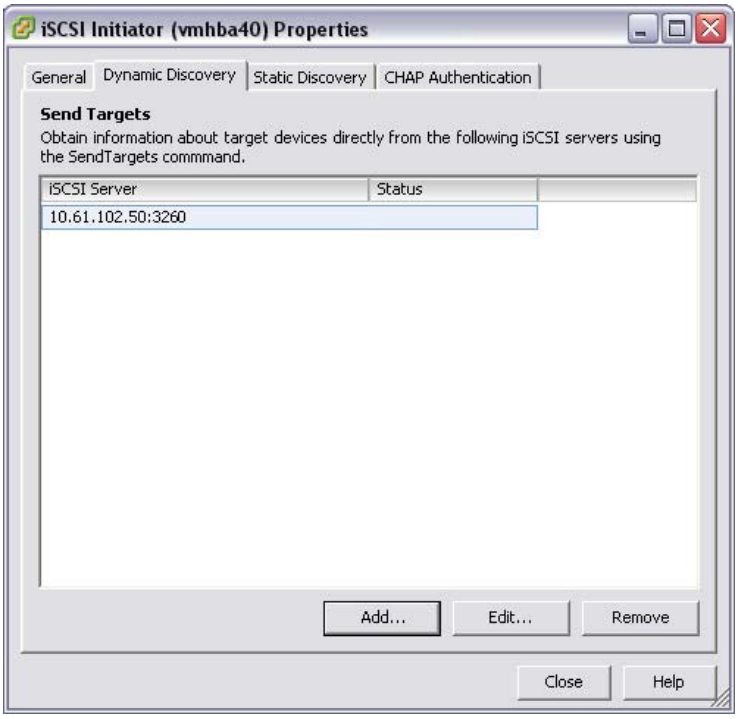

Figure 20

NetApp offers an iSCSI target host adapter for FAS systems. Using this adapter can provide additional scalability of the FAS storage controller by reducing the CPU load of iSCSI transactions. An alternative to the iSCSI target host adapter is to use TOE-enabled NICs for the iSCSI traffic. Although the iSCSI host adapters provide the greatest performance and system scalability, they do require additional NICs to be used to support all other IP operations and protocols. TOE-enabled NICs handle all IP traffic just like a traditional NIC, in addition to the iSCSI traffic.

NetApp offers iSCSI HBAs for use with iSCSI implementations. For larger deployments, scalability benefits may be realized in storage performance by implementing iSCSI HBAs. Note that this statement is not a requirement or a recommendation, but rather a consideration when designing dense storage solutions. The benefits of iSCSI HBAs are best realized on FAS systems, because the storage arrays will have a higher aggregated I/O load than that of any individual VI3 Server.

iSCSI offers several options for addressing storage. If you are not ready to use iSCSI for your primary data access, you may consider iSCSI for several other uses. iSCSI could be used to connect to data stores that store CD-ROM ISO images. Also, iSCSI can be used as a redundant or failover path for a primary Fibre Channel path. If you are using this setup, you must configure LUN multipathing. See “NetApp Fibre Channel Multipathing,” later in this document.

NFS Connectivity

When you are using NFS connectivity for storage, it is a best practice to separate the NFS traffic from other IP network traffic by implementing a separate network or VLAN than that used for VMotion or virtual machines. To enable NFS connectivity, ESX Server requires a special connection type, referred to as a VMkernel port. The VMkernel network requires an IP address that is currently not in use on the ESX Server.

When configuring virtual machines to run on NFS, it was stated at VMworld 2006 to move the VMware swap file (one is created for each VM) to a VMFS datastore. With NFS deployments an easy no cost solution to this need is to deploy an iSCSI LUN. At the time of writing this document, this recommendation could not be validated. For more information on data layout, including VMware swap file relocation, see “VMware Swap and Log File Data Layout,” later in this document.

NetApp offers TOE-enabled NICs for use with IP traffic, including NFS. For larger deployments, scalability benefits can be realized in storage performance by implementing TOE-enabled NICs. Note that this statement is not a requirement or a recommendation, but rather a consideration when designing dense storage solutions. The benefits of TOE-enabled NICs are best realized on FAS systems, because the storage arrays will have a higher aggregated I/O load than that of any individual VI3 Server.

NetApp Fibre Channel Multipathing

Network Appliance clustered FAS systems have an option known as cfmode, which controls the behavior of the system's Fibre Channel ports if a cluster failover occurs. If you are deploying a clustered solution that provides storage for a VMware data center, you must make sure that cfmode is set to either Standby or Single System Image. Standby mode supports VMware, Windows, Linux, and Solaris™ FCP hosts. Single System Image supports all FCP hosts. For a complete list of supported ESX FCP configurations, see the NetApp SAN Support Matrix. To verify the current cfmode, follow these steps.

1	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
2	Enter fcp show cfmode.
3	If cfmode needs to be changed, enter fcp set cfmode <mode type>.

Standby cfmode may require more switch ports, because the multipathing failover is handled by the FAS system and is implemented with active/inactive ports. Single System Image requires additional multipathing configuration on the VMware server. For more information on the different cfmodes available and the impact of changing a cfmode, see section 8 in the Data ONTAP Block Management Guide.

VMware Fibre Channel and iSCSI Multipathing

If you have implemented Single System Image cfmode, then you must configure ESX multipathing. When you are using multipathing, VMware requires the default path to be selected for each LUN connected on each ESX Server. To set the paths, follow these steps.

1	Open VirtualCenter.
2	Select an ESX Server.
3	In the right pane, select the Configuration tab.
4	In the Hardware, box select Storage.
5	In the Storage box, highlight the storage and select the Properties link. See Figure 21.
6	In the Properties dialog box, click the Manage Paths button.
7	Identify the path you want to set as the primary active path and click the Change button. See Figure 22.
8	In the Change Path State window, select the path as Preferred and Enabled and click OK. See Figure 23.

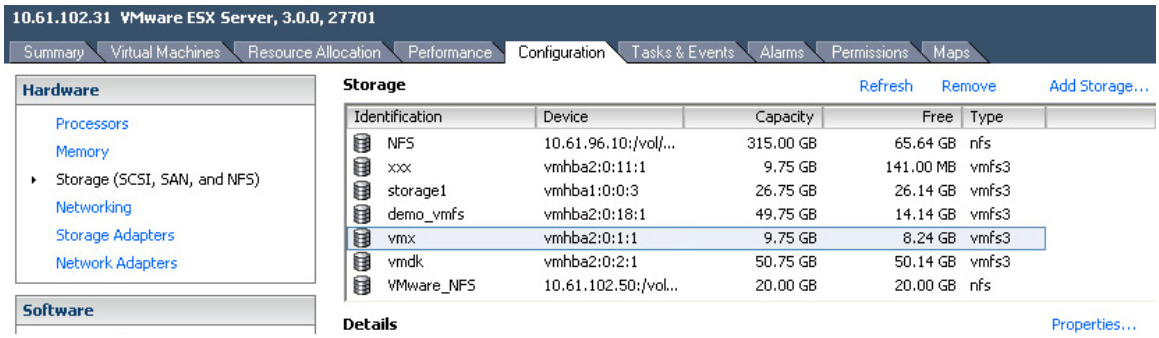

Figure 21

Figure 22

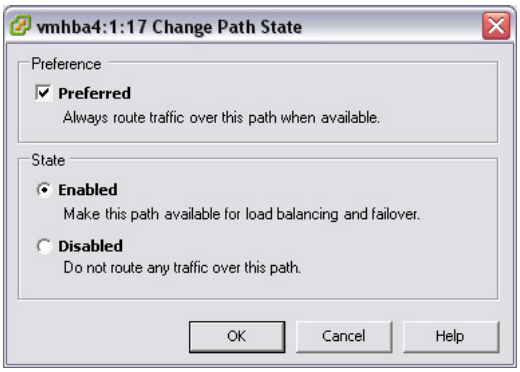

Figure 23

An alternative method for setting the preferred path for multiple LUNs is available in VirtualCenter. This method sets the preferred path for multiple LUNs. To set the paths, follow these steps.

1	Open VirtualCenter.
2	Select an ESX Server.
3	In the right pane, select the Configuration tab.
4	In the Hardware box, select Storage Adapters.
5	In the Storage Adapters pane, select a host bus adapter.
6	Highlight all of the LUNs that you want to configure.
7	Right-click the highlighted LUNs and select Manage Paths. See Figure 24.
8	In the Manage Path window, set the multipathing policy and preferred path for all of the highlighted LUNs. See Figure 25.

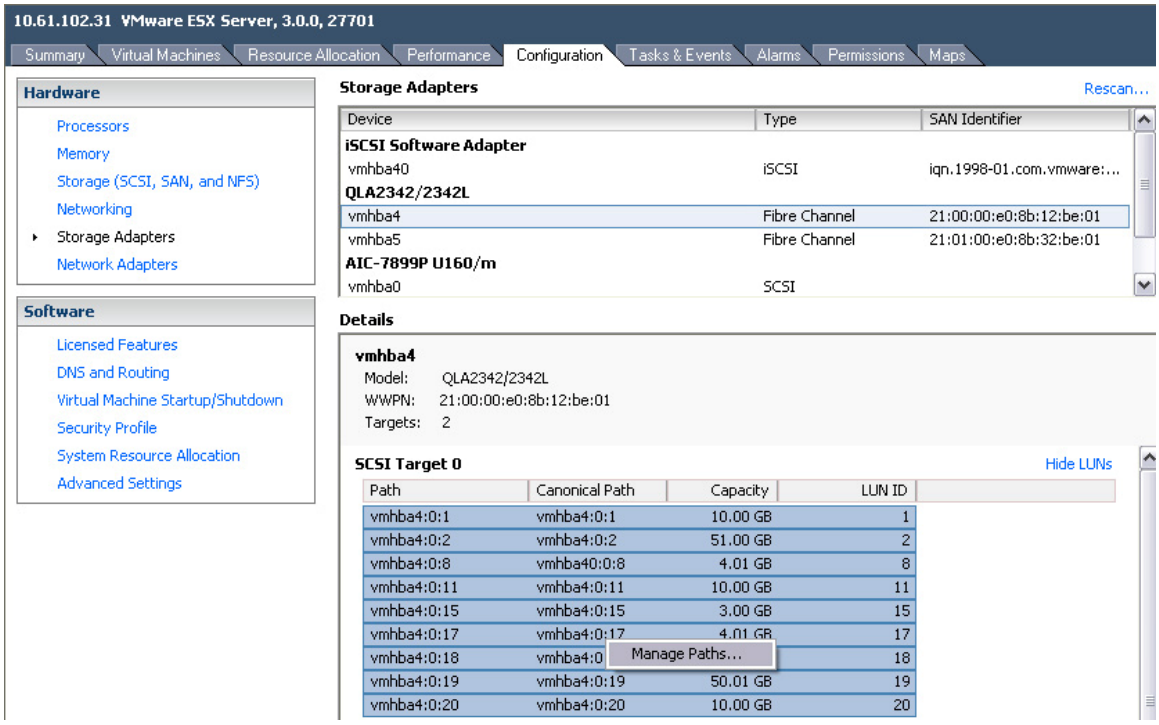

Figure 24

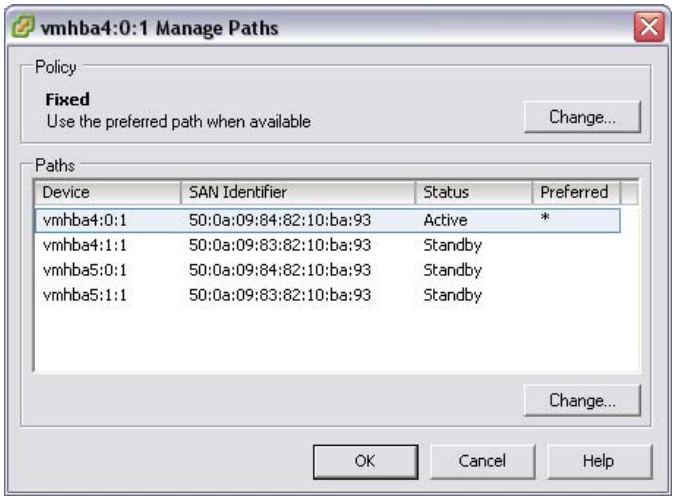

Figure 25

Managing Multipathing With NetApp ESX Host Utilities

NetApp provides a utility for simplifying the management of ESX nodes on FC SAN. This utility is a collection of scripts and executables that is referred to as the FCP ESX Host Utilities for Native OS.

One of the components of the Host Utilities is a script called config_mpath. This script reduces the administrative overhead of managing SAN LUN paths by using the procedures previously described. The config_mpath script can determine the desired primary paths to each of the SAN LUNs on the ESX Server, and then sets the preferred path for each LUN to use one of the primary paths. Multipathing configuration for large numbers of LUNs can be completed quickly and easily by simply running the config_mpath script once on each ESX Server in the data

center. If changes are made to the storage configuration, the script is simply run an additional time to update multipathing configuration based on the changes to the environment.

Other notable components of the FCP ESX Host Utilities for Native OS are the config_hba script, which sets the HBA timeout settings and other system tunables required by the NetApp storage device, and a collection of scripts used for gathering system configuration information in the event of a support issue.

More information on the FCP ESX Host Utilities for Native OS is available at:
http://now.netapp.com/NOW/knowledge/docs/san/

DATA DEDUPLICATION

One of the most popular VMware features is the ability to rapidly deploy new virtual machines from stored VM templates. A VM template includes a VM configuration file (.vmx) and one or more virtual disk files (.vmdk), which includes an operating system, common applications, and patch files or system updates. Deploying from templates saves administrative time by copying the configuration and virtual disk files and registering this second copy as an independent VM. By design, this process introduces duplicate data for each new VM deployed. Figure 26 shows an example of typical storage consumption in a VI3 deployment.

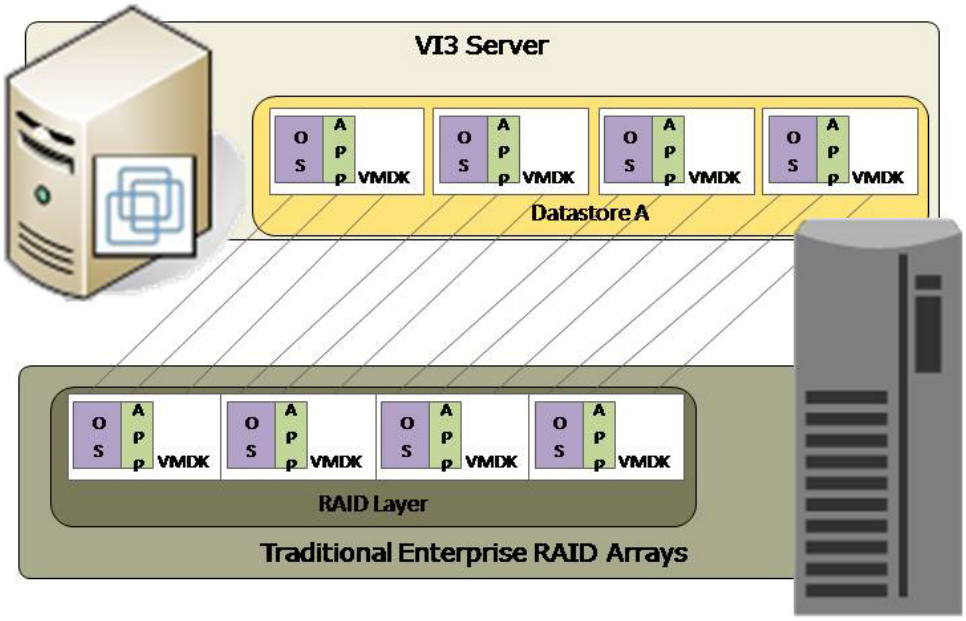

Figure 26

NetApp offers a data deduplication technology called Advanced Single Instance Storage (A-SIS). With A-SIS, VMware deployments can eliminate the duplicate data in their environment, enabling greater storage utilization. A-SIS provides a virtualization technology that enables multiple virtual machines to share the same physical blocks in a NetApp FAS system in the same manner that VMs share system memory. A-SIS can be seamlessly introduced into a virtual infrastructure without having to make any changes to VMware administration, practices, or tasks. Figure 27 shows an example of the impact of A-SIS on storage consumption in a VI3 deployment.

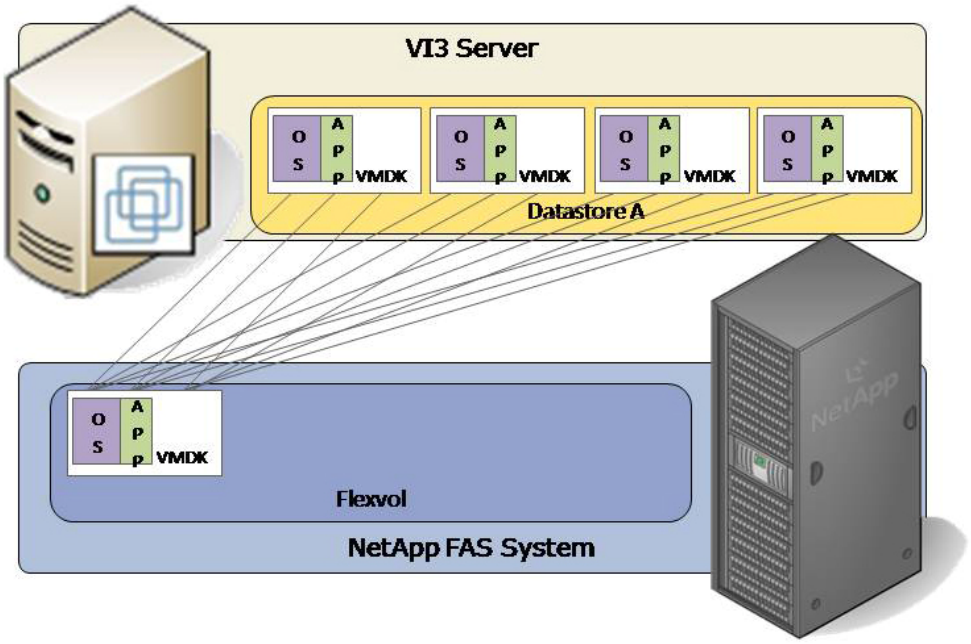

Figure 27

A-SIS is enabled on a volume, and the amount of data deduplication realized is based on the commonality of the data stored in an A-SIS-enabled volume. For the largest storage savings, NetApp recommends grouping similar operating systems and similar applications into datastores, which ultimately reside on an A-SIS-enabled volume.

A-SIS Considerations with VMFS and RDM LUNs
Enabling A-SIS when provisioning LUNs produces storage savings. However, the default behavior of a LUN is to reserve an amount of storage equal to the provisioned LUN. This design means that although the storage array reduces the amount of capacity consumed, any gains made with A-SIS are, for the most part, unrecognizable, because the space reserved for LUNs is not reduced.

To recognize the storage savings of A-SIS with LUNs, you must enable LUN thin provisioning. For details, see "Storage Thin Provisioning," later in this report. In addition, although A-SIS reduces the amount of consumed storage, this benefit is not seen directly by the VMware administrative team, because their view of the storage is at a LUN layer, and (as explained in the previous paragraph) LUNs always represent their provisioned capacity, whether they are traditional or thin provisioned.

A-SIS Considerations with NFS
Unlike with LUNs, when A-SIS is enabled with NFS, the storage savings are both immediately available and recognized by the VMware administrative team. No special considerations are required for its usage.
For A-SIS best practices, including scheduling and performance considerations, see TR 3505: NetApp A-SIS Deduplication Deployment and Implementation Guide.

STORAGE THIN PROVISIONING

VMware provides an excellent means to increase the hardware utilization of physical servers. By increasing hardware utilization the amount of hardware in a data center can be reduced thus lowering the cost of data center operations. In a typical VMware environment the process of migrating physical servers to virtual machines does not reduce the amount of data stored or the amount of storage provisioned. By default, server virtualization does not have any impact on improving storage utilization (and in many cases may have the opposite effect).

You should be very familiar with traditional storage provisioning and with the manner in which storage is preallocated and assigned to a server, or in the case of VMware, a virtual machine. It is also a common practice for server administrators to overprovision storage in order to avoid running out of storage and the associated application downtime when expanding the provisioned storage. Although no system can be run at 100% storage

utilization, there are methods of storage virtualization that allow administrators to address and oversubscribe storage in the same manner as with server resources (such as CPU, memory, networking, etc). This form of storage virtualization is referred to as *thin provisioning*.

Thin provisioning provides storage on demand; traditional provisioning preallocates storage. The value of thin-provisioned storage is that storage is treated as a shared resource pool and is consumed only as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused but provisioned areas of storage that are associated with traditional storage. The drawback to thin provisioning and oversubscribing storage is that (without the addition of physical storage) if every VM requires its maximum possible storage at the same time, there will not be enough storage to satisfy the requests.

VMware and NetApp both provide methods for increasing the storage utilization rate by using thin provisioning. VMware offers a thin format version of its virtual disk (VMDK). With this feature, storage is consumed on demand by the VM. VMDK files that reside on NFS datastores are in the thin format by default. To implement thin VMDKs on FCP or iSCSI, you issue those commands via the ESX Server console (or script). NetApp thin provisioning significantly enhances the thin provisioning built into VMware by allowing VMFS datastores and RDM LUNs to be thin provisioned. This combination of technologies allows any type of VM storage to be thin provisioned and also allows several different methods for increasing storage utilization.

To measure the impact that thin provisioning can have on increasing storage utilization (and thus reducing hardware costs), consider the scenario described in Technical Report 3515: VMware 3.0 on NetApp.

VMware Thin-Provisioning Options

By default, virtual disks preallocate the storage they require and in the background zero out all of the storage blocks. This type of VMDK format is called a *zeroed thick VMDK*. To create a thin-provisioned VMDK file, you use the `vmkfstools` command with the `-d` options switch. By using VMware thin-provisioning technology, you can reduce the amount of storage consumed on a VMFS datastore.

Note that VMDKs that are created as thin-provisioned disks can be converted to traditional zero thick format; however, you cannot convert an existing zero thick format into the thin-provisioned format, with the exception of importing ESX 2.x VMDKs into ESX 3.x.

NetApp Thin-Provisioning Options

NetApp thin provisioning extends VMware thin provisioning for VMDKs and allows LUNs that are serving VMFS data stores to be provisioned to their total capacity limit yet consume only as much storage as is required to store the VMDK files (which can be of either thick or thin format). In addition, LUNs connected as RDMs can be thin provisioned. To create a thin provisioned LUN, follow these steps.

1	Open FilerView (http://filer/na_admin).
2	Select LUNs.
3	Select Wizard.
4	In the Wizard window, click Next.
5	Enter the path.
6	Enter the LUN size.
7	Enter the LUN type (for VMFS select VMware; for RDM select the VM type).
8	Enter a description and click Next.
9	Deselect the Space-Reserved check box. See Figure 28.
10	Click Next and then click Finish.

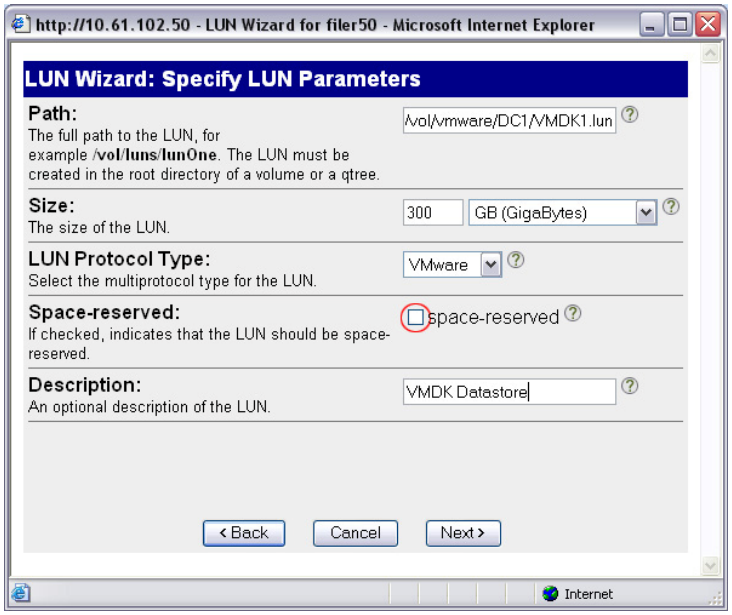

Figure 28

NetApp recommends that when you enable NetApp thin provisioning, you also configure storage management policies on the volumes that contain the thin-provisioned LUNs. The use of these policies aids in providing the thin-provisioned LUNs with storage capacity as they require it. The policies include automatic sizing of a volume, automatic Snapshot deletion, and LUN fractional reserve.

Volume Auto Size is a policy-based space management feature in Data ONTAP that allows a volume to grow in defined increments up to a predefined limit if the volume is nearly full. For VMware environments, NetApp recommends setting this value to On. Doing so requires setting the maximum volume and increment size options. To enable these options, follow these steps.

1	Log into NetApp console.
2	Set volume autosize policy: *vol autosize <vol-name> [-m <size>[k|m|g|t]] [-i <size>[k|m|g|t]] on.*

Snapshot Auto Delete is a policy-based space-management feature that automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full. For VMware environments, NetApp recommends setting this value to delete Snapshot copies at 5% of available space. In addition, you should set the volume option to have the system attempt to grow the volume before deleting Snapshot copies. To enable these options, follow these steps.

1	Log into NetApp console.
2	Set Snapshot autodelete policy: *snap autodelete <vol-name> commitment try trigger volume target_free_space 5 delete_order oldest_first.*
3	Set volume autodelete policy: *vol options <vol-name> try_first volume_grow.*

LUN Fractional Reserve is a policy that is required when you use NetApp Snapshot copies on volumes that contain VMware LUNs. This policy defines the amount of additional space reserved to guarantee LUN writes if a volume becomes 100% full. For VMware environments where Volume Auto Size and Snapshot Auto Delete are in use and you have separated the temp, swap, pagefile, and other transient data onto other LUNs and volumes, NetApp recommends setting this value to 0%. Otherwise, leave this setting at its default of 100%. To enable this option, follow these steps.

1	Log into NetApp console.

2	Set volume Snapshot fractional reserve: *vol options <vol-name> fractional_reserve 0.*

MONITORING AND MANAGEMENT

MONITORING STORAGE UTILIZATION WITH NETAPP OPERATIONS MANAGER

Network Appliance offers the Operations Manager product to monitor, manage, and generate reports on all of the NetApp FAS systems in an organization. When you are utilizing NetApp thin provisioning, NetApp recommends deploying Operations Manager and setting up e-mail and pager notifications to the appropriate administrators. With thin-provisioned storage, it is very important to monitor the free space available in storage aggregates. Proper notification of the available free space ensures that additional storage can be made available before the aggregate becomes completely full. For more information on setting up notifications in Operations Manger, see:

http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel34/html/software/admin/monitor5.htm

http://now.netapp.com/NOW/knowledge/docs/DFM_win/rel34/html/software/admin/filesys4.htm#1217177

STORAGE GROWTH MANAGEMENT

Growing VMFSs
The storage for VMFSs can be increased quite easily and without downtime because NetApp FAS systems allow dynamic resizing of LUNs, and VirtualCenter allows VMFSs to be expanded on the fly. To grow a datastore, follow these steps.

1	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
2	Select LUNs.
3	Select Manage.
4	In the left pane, select the LUN from the list.
5	Enter the new size of the LUN in the Size box and click Apply.
6	Open VirtualCenter.
7	Select an ESX host.
8	In the right pane, select the Configuration tab.
9	In the Hardware box, select the Storage Adapters link.
10	In the right pane, select the HBAs and then select the Rescan link.
11	In the Hardware box, select the Storage link.
12	In the right pane right, select the datastore that you want to grow and then select Properties.
13	Click Add Extent.
14	Select the LUN and click Next.
15	Click Next again. As long as the window shows free space available on the LUN, you can ignore the warning message. See Figure 29.
16	Make sure that the Maximize Space checkbox is selected, then click Next and Finish.

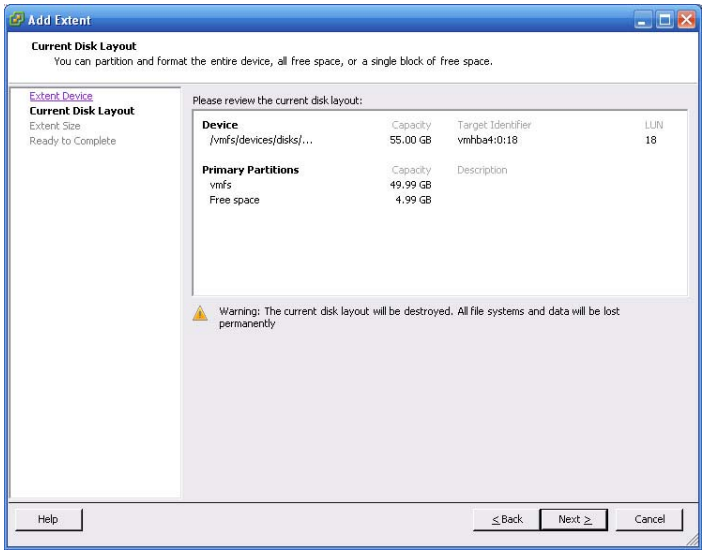

Figure 29

For more information on adding VMFS extents. see the VMware Server Configuration Guide.

Growing a Virtual Disk (VMDK)

Virtual disks can be extended; however, this process requires the virtual machine to be powered off. Growing the virtual disk is only half of the equation for increasing available storage; you will still need to grow the file system after the VM boots. Note that root volumes such as C:\ in Windows and / in Linux cannot be grown dynamically or while the system is running. For these volumes, see "Growing Bootable Volumes," later in this report. For all other volumes, you can use native operating system tools to grow the volume. To grow a virtual disk, follow these steps.

1	Connect to the ESX console (via either SSH, Telnet, or Console connection).
2	Shut down the VM: *Vmware-cmd <cfg> stop.*
3	Extend the virtual disk: *Vmkfstools -X [k|m|g|t] <path to VMDK>.*
4	Start the VM: *Vmware-cmd <cfg> start.*
5	You need to grow the file system within the virtual disk. Follow the guidelines in "Growing a VM File System," later in this report.

For more information on extending a virtual disk, see the VMware Server Configuration Guide.

Growing a Raw Device Mapping (RDM)

Growing an RDM has components of growing a VMFS and a virtual disk. This process requires the virtual machine to be powered off. To grow RDM base storage, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host and power down the VM.
3	Right-click the VM and select Edit Settings to open the Edit Settings window.
4	Highlight the hard disk to be resized and click Remove. Select the Remove from Virtual Machine radio button and select Delete Files from Disk. This action deletes the Mapping File but does *not* remove any data from the RDM LUN. See Figure 30.
5	Open FilerView (http://filer/na_admin).
6	Select LUNs.
7	Select Manage.
8	From the list in the left pane, select the LUN.

9	In the Size box, enter the new size of the LUN and click Apply.
10	Open VirtualCenter.
11	In the right pane, select the Configuration tab.
12	In the Hardware box, select the Storage Adapters link.
13	In the right pane, select the HBAs and click the Rescan link.
14	Right-click the VM and select Edit Settings to open the Edit Settings window,
15	Click Add, then select Hard Disk and click Next. See Figure 31.
16	Select the LUN and click Next. See Figure 32.
17	Specify the VMFS datastore that will store the Mapping File.
18	Start the VM. Remember that although you have grown the LUN, you still need to grow the file system within it. Follow the guidelines in “Growing a VM File System,” later in this report.

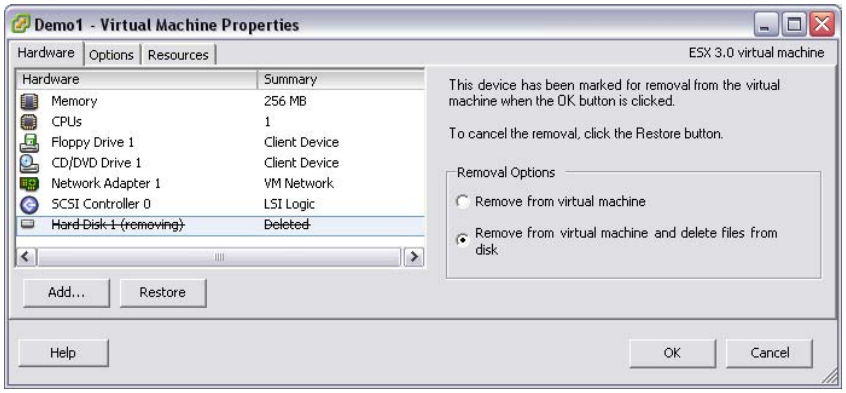

Figure 30

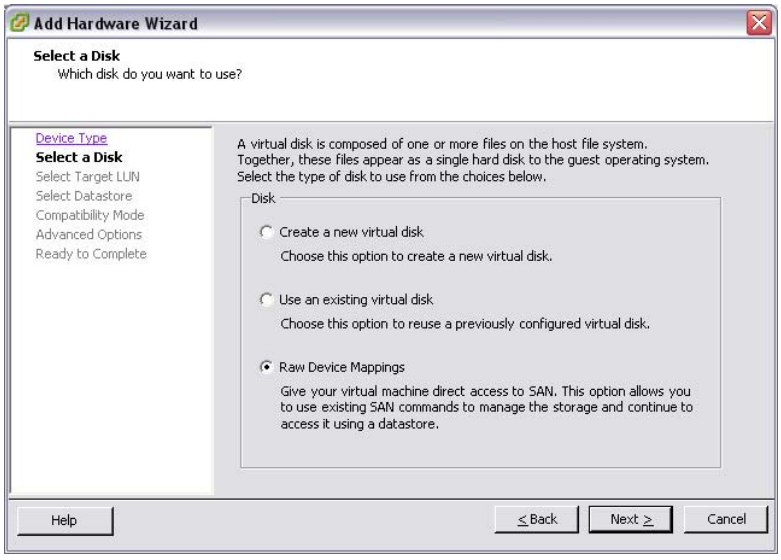

Figure 31

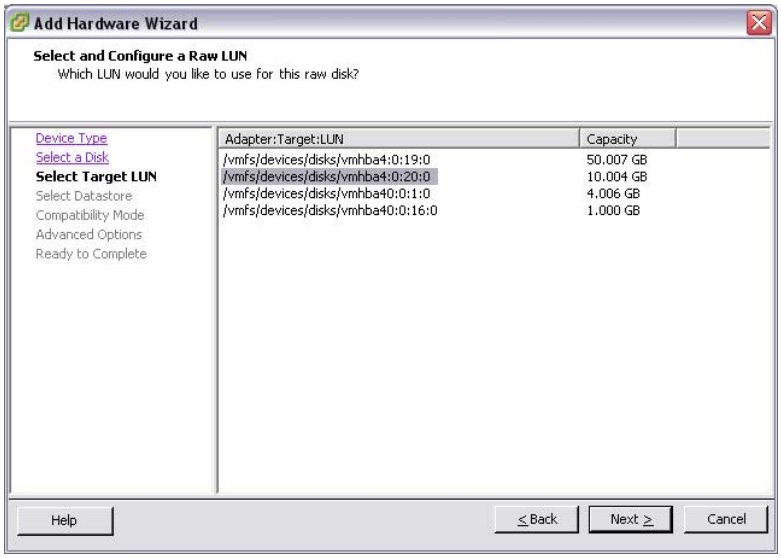

Figure 32

Growing a VM File System (NTFS or EXT3)
When a virtual disk or RDM has been increased in size, you still need to grow the file system residing on it after booting the VM. This process can be done live while the system is running, by using native or freely distributed tools.

1	Remotely connect to the VM.
2	Grow the file system.
3	For Windows VMs, you can use the diskpart utility to grow the file system. For more information, see http://support.microsoft.com/default.aspx?scid=kb;en-us;300415. **or** For Linux VMs, you can use ext2resize to grow a file system. For more information, see http://sourceforge.net/projects/ext2resize.

Growing Bootable Volumes
Root volumes such as C:\ in Windows VMs and / in Linux VMs cannot be grown on the fly or while the system is running. There is a simple way to expand these file systems that does not require the acquisition of any additional software (except for ext2resize). This process requires the VMDK or LUN that has been resized to be connected to another virtual machine of the same operating system type, by using the processes defined earlier. Once the storage is connected, the hosting VM can run the utility to extend the file system. After extending the file system, this VM is shut down and the storage is disconnected. Connect the storage to the original VM. When you boot, you can verify that the boot partition now has a new size.

BACKUP AND RECOVERY

SNAPSHOT TECHNOLOGIES

VMware Virtual Infrastructure 3.0 introduced the ability to create Snapshot copies of virtual machines. Snapshot technologies allow point-in-time copies to be made that provide the fastest means to recover a VM to a previous point in time. NetApp has been providing customers with the ability to create Snapshot copies of their data since 1992, and although the basic concept of a snapshot is similar between NetApp and VMware, you should be aware of the major differences between the two, and when you should use one rather than the other.

VMware snapshots provide simple point-in-time versions of VMs, allowing quick recovery. The benefits of VMware snapshots are that they are easy to create and use, because they can be executed and scheduled from within VirtualCenter. One concern about VMware snapshots is that they are software based, and most software-based snapshot technologies have significant scalability and performance issues. For more information on native VMware snapshots, see the VMware Basic Systems Administration Guide and the ESX Server 3.x Storage/SAN Compatibility Guide.

The patented NetApp Snapshot technology can easily be integrated into VMware environments, where it provides crash-consistent versions of virtual machines for the purpose of full VM recovery, full VM cloning, or site replication and disaster recovery. The benefits of this solution are that it is the storage industry's only snapshot technology that does not have a negative impact on system performance. VMware states that for optimum performance and scalability, hardware-based snapshot technology is preferred over software-based solutions. The shortcoming of this solution is that it is not managed within VirtualCenter, requiring external scripting and/or scheduling to manage the process. For details, see the VMware Basic Systems Administration Guide and the VMware Server Configuration Guide.

DATA LAYOUT FOR SNAPSHOT COPIES

When you are implementing either NetApp Snapshot copies or SnapMirror, NetApp recommends separating transient and temporary data off the Virtual Disks that will be copied using Snapshot or SnapMirror. Because Snapshot copies hold onto storage blocks that are no longer in use, transient and temporary data can consume a large amount of storage in a very short period of time. In addition, if you are replicating your environment for business continuance or disk-to-disk backup purposes, failure to separate the valuable data from the transient will have a large impact on the amount of data sent at each replication update.

Virtual machines should have their swap files, pagefiles, and user and system temp directories moved to separate virtual disks residing on separate datastores residing on NetApp volumes dedicated to this data type. In addition, the ESX servers create a VMware swap file for every running VM. These files should also be moved to a separate datastore residing on a separate NetApp volume. Figure 33 shows an example of this data layout.

For example, if you have a group of VMs that create a Snapshot copy three times a day and a second group that creates a Snapshot copy once a day, then you need a minimum of four NetApp volumes. For traditional virtual disks residing on VMFS, each volume contains a single LUN; for virtual disks residing on NFS, each volume has several virtual disk files; and for RDMS, each volume contains several RDM-formatted LUNs. The Snapshot backup script must be configured for each volume containing VMs and the appropriate Snapshot schedule.

Virtual Machine Data Layout

This section looks at the transient temporary data that is a component of a virtual machine. This example focuses on a Windows guest operating system, because the requirements to set up this data layout are a bit more complex than those for other operating systems; however, the same principles apply to other operating systems. To reduce the time required to create this configuration, you should make a master virtual disk of this file system and clone it with VMware virtual disk cloning when you are either creating new virtual machines or starting virtual ,machines at a remote site or location in a disaster recovery process.

Following is a registry file example of a simple registry script that sets the pagefile and temp area (for both user and system) to the D:\ partition. This script should be executed the first time a new virtual machine is created. If the D:\ partition does not exist, the system default values are used. The process of launching this script can be automated with Microsoft Setup Manager. To use the values in this example, copy the contents of this section and save it as a text file named temp.reg. The Setup Manager has a section where you can add temp.reg to the run the first time the virtual machine is powered on. For more information on automating the deployment of cloned Windows servers, see Microsoft Setup Manager.

```
Start----------

Windows Registry Editor Version 5.00


[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management]

"PagingFiles"=hex(7):64,00,3a,00,5c,00,70,00,61,00,67,00,65,00,66,00,69,00,6c,\
```

00,65,00,2e,00,73,00,79,00,73,00,20,00,32,00,30,00,34,00,38,00,20,00,32,00,

30,00,34,00,38,00,00,00,00,00

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment]

"TEMP"="D:\\"

"TMP"="D:\\"

[HKEY_CURRENT_USER\Environment]

"TEMP"="D:\\"

"TMP"="D:\\"

[HKEY_USERS\.DEFAULT\Environment]

"TEMP"="D:\\"

"TMP"="D:\\"

End ----------

VMware Swap and Log File Data Layout

The VMware ESX Server creates a swap file and logs for every running VM. The sizes of these files are dynamic; they change in size according to the difference in the amount of physical memory in the server and the amount of memory provisioned to running VMs. Because this data is transient in nature, it should be separated from the valuable VM data when implementing NetApp Snap technologies. Figure 33 shows an example of this data layout.

A prerequisite to making this change is the creation of a VMFS datastore to store the swap files. Because the VMware swap file storage requirements are dynamic, NetApp suggests creating either a large thin-provisioned LUN or a FlexVol with the auto grow feature enabled. Thin provisioned LUNs and auto grow FlexVols provide a large management benefit when storing swap files. In this design, it removes the need to micromanage the swap space or to reduce the utilization rate of the storage. Consider the alternative of storing VMware swap files on traditional storage arrays. If you undersize the swap space, the VMs fail to start; conversely, if you oversize the swap space, you have provisioned but unused storage. In order to configure these settings follow this process:

1	Open VirtualCenter.
2	Select either a Virtual Machine or a VM template
3	If the Virtual machine is running, stop it
4	Connect to the ESX console (via either SSH, Telnet, or Console connection).
5	Cd to the path of the .vmx file you wish to edit
6	Add the line workingDir = /vmfs/volumes/<volume_name_of_temp>
7	Delete the line sched.swap.dir = xxx
8	Delete the line sched.swap.derivedName = xxx
9	Restart Virtual Machine if its not a template
10	Repeat as necessary for each existing Virtual Machine

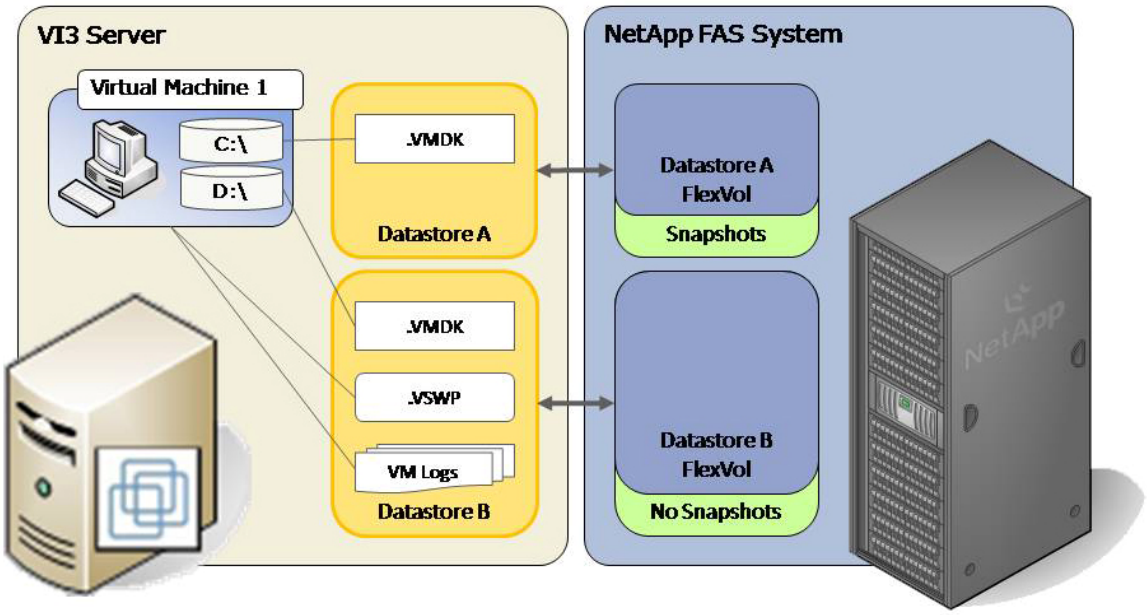

Figure 33

SNAPSHOT CONCEPTS

IMPLEMENTING SNAPSHOT COPIES

The consistency of the data contained in a Snapshot copy is paramount to a successful recovery. This section describes how to implement NetApp Snapshot copies for VM recovery, cloning, and disaster recovery replication.

ESX SNAPSHOT CONFIGURATION FOR SNAPSHOT COPIES

In a VMware Virtual Infrastructure, the storage provisioned to virtual machines is stored in either virtual disk files (residing on VMFS or NFS) or in raw device mappings (RDMs). With the introduction of VI3, administrators have the ability to mount storage-created Snapshot copies of VMFS LUNs. With this feature, customers can now connect to Snapshot copies of both VMFS and RDM LUNs from a production ESX Server. To enable this functionality, follow these steps.

1	Open VirtualCenter.
2	Select an ESX Server.
3	In the right pane, select the Configuration tab.
4	In the Software box, select Advanced Settings to open the Advanced Settings window.
5	In the left pane, select LVM.
6	In the right pane, enter the value of 1 in the LVM.EnableResignature box.
7	Repeat steps 2 through 6 for each ESX Server in the data center.

This process can also be completed from the console of the ESX Server using the following command:

```
esxcfg-advcfg -s 1 /LVM/EnableResignature
```

Note: You may need to run the following commands to see the changes made in the VirtualCenter GUI:

```
service mgmt-vmware restart
```

Note: Virtual Infrastructure 3 also allows users to connect to Snapshot copies of NFS datastores from a production ESX Server. No special configuration is required to facilitate this process.

ESX SERVER AND NETAPP FAS SSH CONFIGURATION

The most efficient way to integrate NetApp Snapshot copies is to allow the centralized management and execution of Snapshot copies. NetApp recommends configuring the FAS systems and ESX Servers to allow a single host to remotely execute commands on both systems. This management host is required to have an SSH client installed and configured.

FAS System SSH Configuration

To configure SSH access on a NetApp FAS system, follow these steps:

1	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
2	Execute the following commands: *secureadmin setup ssh* *options ssh.enable on* *options ssh2.enable on*
3	Log into the Linux or VMware system that remotely executes commands on the FAS system as root.
4	Add the Triple DES cipher to the list of available SSH ciphers; this is the only cipher recognized by the NetApp FAS system. Edit the /etc/ssh/sshd-config file and edit the Ciphers line to read as follows: *Ciphers aes128-cbc, aes256-cbc, 3des-cbc.*
5	Generate a DSA host key. On a Linux or VMware ESX Server, use the following command: *ssh-keygen –t dsa –b 1024.* When prompted for the passphrase, do not enter one; instead, press Enter. The public key is saved to /root/.ssh/id_dsa.pub.
6	Copy only the key information from the public key file to the FAS system's /etc/sshd/root/.ssh/authorized_keys, removing all information except for the key string preceded by the string ssh-dsa and a comment line. See the following example.
7	Test the connectivity from the remote host by issuing the `version` command on the FAS system. It should not prompt for a password. *ssh <netapp> version* NetApp Release 7.2: Mon Jul 31 15:51:19 PDT 2006

This is an example of the key for the remote host:

```
ssh-dsa AAAAB3NzaC1kc3MAAABhALVbwVyhtAVoaZukcjSTlRb/REO1/ywbQECtAcHijzdzhEJUz9Qh96
HVEwyZDdah+PTxfyitJCerb+1FAnO65v4WMq6jxPVYto6l5Ib5zxfq2I/hhT/6KPziS3LTZjKccwAAABUAjkL
MwkpiPmg8Unv4fjCsYYhrSL0AAABgF9NsuZxniOOHHr8tmW5RMX+M6VaH/nlJUzVXbLiI8+pyCXALQ2
9Y31uV3SzWTd1VOgjJHgv0GBw8N+rvGSB1r60VqgqgGjSB+ZXAO1EecbnjvLnUtf0TVQ75D9auagjOAA
AAYEJPx8wi9/CaS3dfKJR/tYy7Ja+MrlD/RCOgr22XQP1ydexsfYQxenxzExPa/sPfjA45YtcUom+3mieFaQ
uWHZSNFr8sVJoW3LcF5g/z9Wkf5GwvGGtD/yb6bcsjZ4tjlw==
```

ESX System SSH Configuration

To configure an ESX Server to accept remote commands by using SSH, follow these steps.

1	Log into the ESX console as root.
2	Enable the SSH services by running the following commands: *esxcfg-firewall -e sshServer* *esxcfg-firewall -e sshClient*
3	Change to the SSH server configuration directory: *cd /etc/ssh.*

4	Edit the configuration file: vi sshd_config
5	Change the following line from: *PermitRootLogin no* to: *PermitRootLogin yes*
6	Restart the SSH service by running the following command: *service sshd restart*
7	Create the SSH public key: *ssh-keygen -t dsa* This command outputs content similar to the following example. Retain the default locations, and do not use a passphrase.
8	Change to the .ssh directory: *cd /root/.ssh*
9	Run the following commands: *cat id_dsa.pub >> authorized_keys* *chmod 600 authorized_keys*
10	Repeat steps 1 through 9 for each ESX Server in the data center.

Example output:

```
Generating public/private dsa key pair.
Enter file in which to save the key (/home/root/.ssh/id_dsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/root/.ssh/id_dsa.
Your public key has been saved in /home/root/.ssh/id_dsa.pub.
The key fingerprint is:
7b:ab:75:32:9e:b6:6c:4b:29:dc:2a:2b:8c:2f:4e:37 root@hostname

Your keys are stored in /root/.ssh.
```

RECOVERING VIRTUAL MACHINES FROM A VMFS SNAPSHOT COPY

Snapshot copies of VMFS datastores provide a quick method to recover a VM from a Snapshot copy. In summary, this process powers off the VM, attaches the Snapshot copy VMFS LUN, copies the VMDK from the Snapshot copy to the production VMFS, and powers on the VM. To complete this process, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host and power down the VM.
3	Log into the ESX console as root.
4	Rename the VMDK files: *mv <current VMDK path> <renamed VMDK path>*
5	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
6	Clone the original LUN from a recent Snapshot copy, bring it online, and map it. From the storage appliance console, run: *lun clone create <original LUN path> –b <original LUN path> <Snapshot name>* *lun online <LUN path>* *lun map <LUN path> <igroup> <ID>*
7	Open VirtualCenter.
8	Select an ESX host.

9	In the right pane, select the Configuration tab.
10	In the Hardware box, select the Storage Adapters link.
11	In the upper right corner, select the Rescan link. Scan for both new storage and VMFS datastores. The Snapshot VMFS datastore appears.
12	Log into the ESX console as root.
13	Copy the virtual disks from the Snapshot datastore to the production VMFS: *cd <VMDK snapshot path>* *cp <VMDK> <production VMDK path>*
14	Open VirtualCenter.
15	Select the ESX Server and start the virtual machine.
16	Validate that the restore is to the correct version. Log into the VM and verify that the system was restored to the proper point in time.
17	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
18	Delete the Snapshot copy LUN: *lun destroy –f <LUN path>*
29	In the upper right corner, select the Rescan link. Scan for both new storage and VMFS datastores.

RECOVERING VIRTUAL MACHINES FROM AN RDM SNAPSHOT

RDMs provide the quickest possible method to recover a VM from a Snapshot copy. In summary, this process powers off the VM, restores the RDM LUN, and powers on the VM. To complete this process, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host and power down the VM.
3	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
4	Clone the original LUN from a recent Snapshot copy: *lun clone create <original LUN path> –b <original LUN path> <Snapshot name>*
5	Take the current version of the LUN in use off line: lun offline <LUN path>
6	Map and the cloned LUN and put it on line: *lun online <LUN path>* *lun map <LUN path> <igroup> <ID>*
7	Open VirtualCenter.
8	Select an ESX host and power on the VM.
9	Validate that the restore is to the correct version. Log into the VM and verify that the system was restored to the proper point in time.
10	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
11	Delete the original LUN and split the clone into a whole LUN: *lun destroy –f <original LUN path>* *lun clone split start <cloned LUN path>*
12	Rename the cloned LUN to the name of the original LUN (optional): lun mv <cloned LUN path> <original LUN path>

RECOVERING VIRTUAL MACHINES FROM AN NFS SNAPSHOT COPY

NFS provides a quick method to recover a VM from a Snapshot copy. In summary, this process powers off the VM, restores the VMDK, and powers on the VM. To complete this process, follow these steps.

1	Open VirtualCenter.
2	Select an ESX host and power down the VM.
3	Log into the ESX console as root.
4	Rename the VMDK files: *mv <current VMDK path> <renamed VMDK path>*
5	Connect to the FAS system console (via either SSH, Telnet, or Console connection).
6	Restore the VMDK file from a recent Snapshot copy: *snap restore –t file -s <snapshot-name> <original VMDK path> <original VMDK path>*
7	Open VirtualCenter.
8	Select the ESX and start the virtual machine.
9	Validate that the restore is to the correct version. Log into the VM and verify that the system was restored to the proper point in time.
10	Log into the ESX console as root.
11	Delete the renamed VMDK files: *rm <renamed VMDK path>*

SUMMARY

VMware Virtual Infrastructure gives customers several methods of providing storage to virtual machines. Each of these storage methods enables customers to have flexibility in their infrastructure design, which in turn provides cost savings, increased storage utilization, and enhanced data recovery.

This technical report is not intended to be a definitive implementation or solutions guide. Expertise may be required to solve user-specific deployments. Contact your local Network Appliance representative to make an appointment to speak with a NetApp VMware solutions expert.

Comments on this technical report are welcome. Please contact the authors here.

APPENDIX: EXAMPLE HOT BACKUP SNAPSHOT SCRIPT

This script allows effortless backup of virtual machines at the datastore level. This means that virtual machines can be grouped into datastores based on their Snapshot or SnapMirror backup policies, allowing multiple recovery point objectives to be met with very little effort. Critical application server virtual machines can have Snapshot copies automatically created based on a different schedule than second-tier applications or test and development virtual machines. The script even maintains multiple versions of Snapshots.

This script provides managed, consistent, backups of virtual machines in a VMware Virtual Infrastructure 3 environment leveraging Network Appliance Snapshot technology. It is provided as an example that can easily be modified to meet the needs of an environment. For samples of advanced scripts built from this example framework, see Loyola Marymount University and Sirius Computer Solutions.

Backing up VMs with this script completes the following process:

Quiesce all of the VMs on a given datastore
Take a crash-consistent NetApp Snapshot copy
Apply the Redo logs and restore the virtual disk files to a read-write state

```
#!/bin/sh
#
# Example code which  takes a snapshot of all VMs using the VMware
# vmware-cmd facility. It will maintain and cycle the last 3 Snapshot copies.
```

```
#
# This sample code is provided AS IS, with no support or warranties of any
# kind, including but not limited to warranties of merchantability or
# fitness of any kind, expressed or implied.
#
# 2007 Vaughn Stewart, Network Appliance
#
# ---------------------------------------------------------------------

PATH=$PATH:/bin:/usr/bin

# Step 1 Enumerate all VMs on an individual ESX Server, and put each VM in hot backup mode.
for i in `vmware-cmd -l`
do
  vmware-cmd $i createsnapshot backup NetApp quiesce
done

# Step 2 Rotate NetApp Snapshot copies and delete oldest, create new, maintaining 3.
ssh <Filer> snap delete <esx_data_vol> vmsnap.3
ssh <Filer> snap rename <esx_data_vol> vmsnap.2 vmsnap.3
ssh <Filer> snap rename <esx_data_vol> vmsnap.1 vmsnap.2
ssh <Filer> snap create <esx_data_vol> vmsnap.1

# Step 3 Bring all VMs out of hot backup mode,
for i in `vmware-cmd -l`
do
  vmware-cmd $i removesnapshots
done
```

REFERENCES

NetApp TR3612: Network Appliance and VMware Virtual Desktop Infrastructure

NetApp TR3515: Network Appliance and VMware ESX Server 3.0: Building a Virtual Infrastructure from Server to Storage

NetApp TR3482: Network Appliance and VMware ESX Server 2.5.x

NetApp TR3001: A Storage Network Appliance

NetApp TR3466: Open Systems SnapVault® (OSSV) Best Practices Guide

NetApp TR3347: FlexClone Volumes: A Thorough Introduction

NetApp TR3348: Block Management with Data ONTAP 7G: FlexVol®, FlexClone, and Space Guarantees

NetApp TR3446: SnapMirror Best Practices Guide

Total Cost Comparison: IT Decision-Maker Perspectives on EMC and Network Appliance Storage Solutions in Enterprise Database Environments

Wikipedia RAID Definitions and Explanations

VMware Introduction to Virtual Infrastructure

VMware Server Configuration Guide

VMware ESX Server 3.x Storage/SAN Compatibility Guide

VMware VMworld Conference Sessions Overview

VMware Recommendations for Aligning VMFS Partitions

VERSION TRACKING

Version 1.0	May 2006	Original Document
Version 2.0	January 2007	Major Revisions Supporting VI3
Version 2.1	May 2007	Updated VM Snapshot Script and Instructions Added Figure 28
Version 3.0	September 2007	Major Revision Update
Version 3.1	October 2007	Minor corrections, added NFS snapshot config requiremnts

www.netapp.com