Network Appliance™ and VMware ESX Server 3.0 Storage Best Practices

M. Vaughn Stewart & Michael Slisinger, Network Appliance, Inc.
May 2007 | TR-3428 | Version 2.1

Abstract

This document discusses the virtual storage solutions which 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 solutions that are industry leading 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 DR 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 will describe them.

Please note that these practices are only recommendations, not requirements. Not following these recommendations does not affect whether your solution is supported by Network Appliance. Not all recommendations will apply to every scenario. NetApp believes that their customers will benefit from thinking through these recommendations prior to making any implementation decision.

The target audience for this paper has a familiarity with concepts pertaining to VMware ESX Server 3.0 and Network Appliance Data ONTAP® 7.X. For additional information and overview of the unique benefits available when creating a Virtual Infrastructure on NetApp storage please refer to: http://www.netapp.com/library/tr/3515.pdf.

CONFIGURATION & SETUP

VMWARE STORAGE OPTIONS

There are three types of storage options available to VMware's Virtual Infrastructure 3. Each option has its own unique benefits, so we will begin by providing a brief summary of each option.

VMFS Data Store on Fibre Channel or iSCSI

This is the default 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 administrators are free to use the storage as needed. Most operations are run exclusively through VirtualCenter. The shortcomings of this solution are that performance degrades as multiple Virtual Machines are deployed on the same Data store, scalability requires multiple Data stores, and finding performance bottlenecks can be difficult. For an example of this configuration see Figure 1 below. For more information on Virtual Disks over FCP and iSCSI see page 83 of the VMware Server Configuration Guide at http://www.vmware.com/pdf/vi3_server_config.pdf.

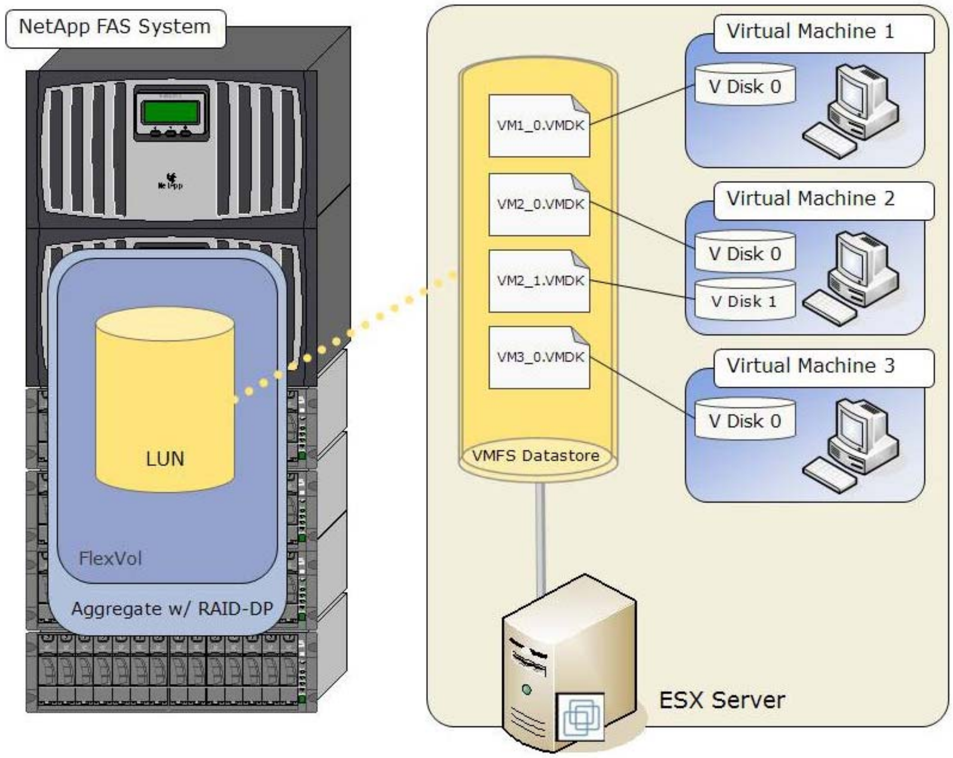

Figure 1

Raw Device Mappings over Fibre Channel or iSCSI

Raw Device Mappings (or RDMs) were introduced in VMware ESX Server 2.5. The strengths of this solution are high disk I/O performance, easy disk performance measurement, Virtual Machine host based clustering (such as 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 it may require more interaction between storage and VMware administrators. For an example of this configuration see Figure 2 below. For more information on Raw Device Mappings over Fibre Channel and iSCSI see page 142 in http://www.vmware.com/pdf/vi3_server_config.pdf.

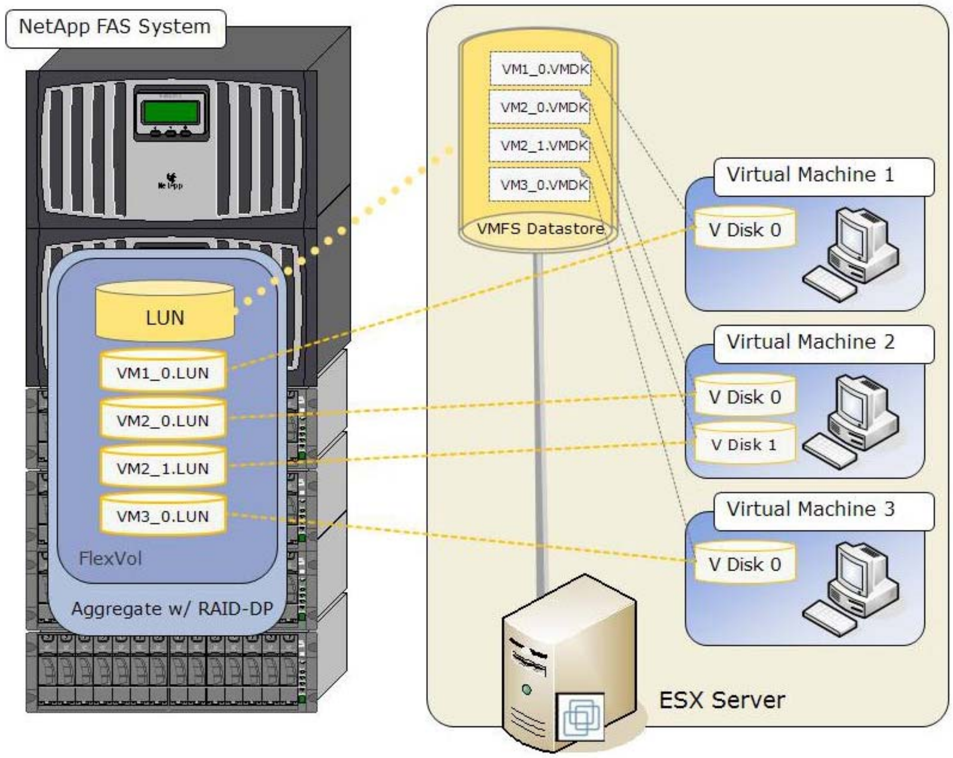

Figure 2

Virtual Disks over NFS

Support for storing Virtual Disks (or VMDKs) on a Network File System (or NFS) was introduced in VMware's ESX Server 3.0. The strength of this solution is that once storage has been provisioned to the ESX Servers, the VMware administrators are 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, thus providing an increase of utilization of total storage capacity. The shortcoming of this solution is that reaching required disk I/O levels may require either additional and/or TOE enabled Network Cards. For an example of this configuration see Figure 3 below. For more information on storing VMDK files on NFS see page 119 in http://www.vmware.com/pdf/vi3_server_config.pdf.

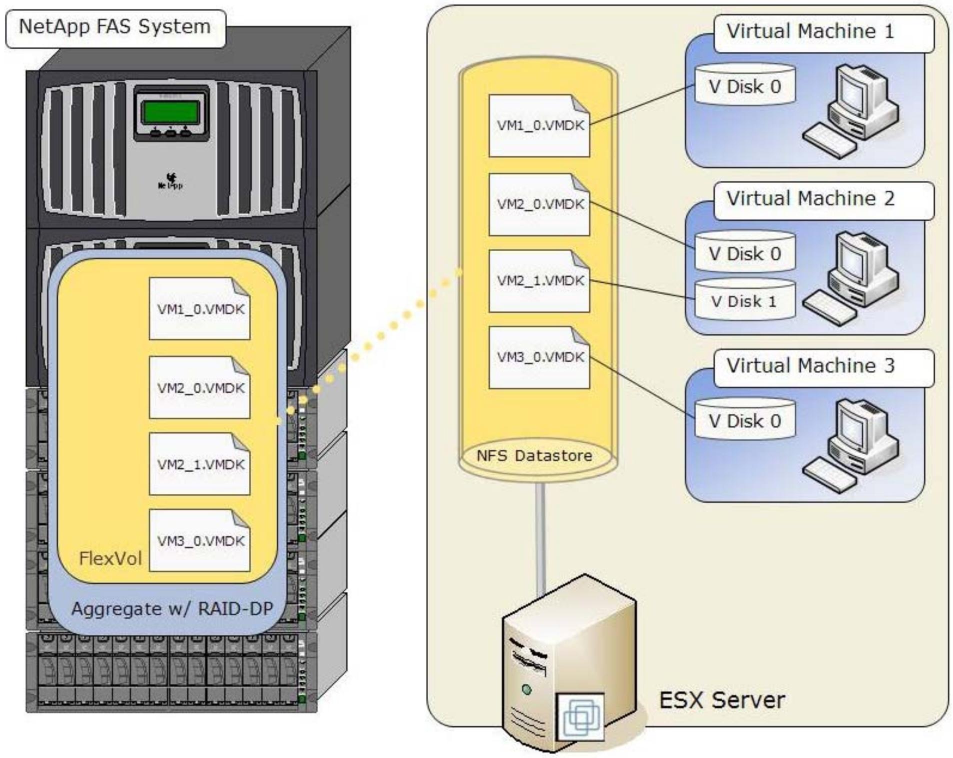

Figure 3

STORAGE CONFIGURATION

RAID Data Protection

The amount of data addressed in today's enterprise is staggering in size compared to what it was just 2 or 3 years ago. Due to this dramatic growth data protection has become paramount. Data protection in a Virtual Infrastructure is even more critical than in a traditional server infrastructure as a storage failure within a VI will lead to multiple Virtual Machines experiencing downtime and/or losing data. NetApp RAID-DP™ is an advanced RAID technology provided as the default RAID level on all Fabric Attached Storage 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 the remaining drives. The Data protection in RAID-DP is greater than that of RAID 10 or RAID 50 as RAID-DP allows for the loss of any two drives within a RAID set/group. As 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 which Flexible Volumes are striped across. It is recommended 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 very random by nature and in order to ensure optimal performance you want to ensure a large number of physical spindles are available to service requests. For more information please see pages 35 and 115 in http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

Volumes

Volumes, or more specifically Flexible Volumes, contain LUNs or Virtual Disk Files which are accessed by VMware ESX Servers. Volumes should be created based on Snapshot™ backup and/or SnapMirror® replication policies. As Virtual Machines can be hosted by any server within a VMware Data center (which is a collection of ESX Servers) storage is not directly mapped to individual ESX Servers; instead the storage is mapped to the entire Data center.

NetApp volumes allow for the creation of a particular type of directory in the root of the volume which is called a Qtree. Qtrees have several properties which distinguish them from traditional directories such as the ability to have separate replication relationships and schedules and individual security semantics. It is recommended that every LUN (FCP or iSCSI) created for an ESX Server be placed into a Qtree.

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 fashions. The first and most common method is as storage to hold Virtual Disk Files for multiple Virtual Machines. This type of usage will be referred to as a VMFS (Virtual Machine File System) LUN. The second method is as a Raw Device Mapping (RDM). With an RDM the LUN is connected to the ESX Server and 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 for simplicity as once storage is provisioned to the ESX Server it can be utilized by the VMware Administrator without intervention from the storage administrator. In addition the built in storage functionality of ESX Server can be leveraged such as VMware snapshots and clones.

RDM LUNs were introduced in ESX Server 2.5 and provide for several benefits not found with VMFS LUNs. RDMs allow for Virtual Machines to be clustered with Microsoft® Cluster Server, 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™ technologies. The downside of RDMs is that each VMware data center has a finite limit of 256 LUNs (which can lead to artificial capacity limitation of a VMware data center) and VMware administration will require an ongoing relationship with storage administration.

For more information please see page 33 in http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

Storage Naming Conventions

NetApp storage systems allow for human naming conventions, and when designing storage for virtual infrastructures implementing a descriptive naming convention is recommended. Descriptive naming conventions provide for simplicity when trying to identify/map multiple layers of storage with Virtual Machines. A simple and efficient naming convention is critical when trying to simplify replication or disaster recovery processes. Please consider the following suggestions:

> Volume Name – matches Data center Name or Data center name and Replication Policy or Data Type (i.e. DC1_Swap, DC_4hr_repl)
>
> Qtree Name – matches hostname of Virtual Machine (used for RDMs only)
>
> LUN Name – for VMFS describes storage contents (i.e. VMDK1, VMX1, Boot, Swap, etc.)
>
> LUN Name – for RDMs should be the volume name of the LUN (i.e. for Windows® c_drive, for Linux® root, etc.)

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

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

> /vol/DC_2_24hr_repl/VMDK_1.lun

This is an example of a volume which stores RDMs for Data Center 4. In this example the C Drive LUN for a Virtual machine named SQL_VM_12 has been created:

> /vol/DC_4/SQL_VM_12/C_drive.lun

CONFIGURATION BASICS

Virtual Machine Data Layout

When implementing either NetApp Snapshot copies or SnapMirror, separating transient and temporary data off of the Virtual Disks which will be copied using Snapshot or SnapMirror is highly recommended. As Snapshot copies may 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. Virtual Machines should have their swap files, page files, user and system temp directories moved to a separate Virtual Disk which resides within a NetApp volume dedicated to this data type.

To help reduce the time required to create this configuration, a master virtual disk should be made of this file system and cloned with VMware's virtual disk cloning when either creating new Virtual Machines or starting Virtual Machines at a remote site/location in a disaster recovery process.

Below is a registry file example of a simple script that will set the pagefile and temp area (for both user and system). This script is meant to direct an administrator where to locate and set the correct values and in turn create a system preparation script to be run on a Windows 2003 Server Virtual Machine master copy. Remember to reseal your master images with sysprep or a similar tool.

Start----------

```
Windows Registry Editor Version 5.00


[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Environment]
"TEMP"=
"TMP"=
 [HKEY_CURRENT_USER\Environment]
"TEMP"=
"TMP"=
[HKEY_USERS\.DEFAULT\Environment]
"TEMP"=
"TMP"=
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Session Manager\Memory Management]
"PagingFiles"=
```

End ----------

CONFIGURATION LIMITS AND RECOMMENDATIONS

When sizing storage there are some limits and recommendations that you should be aware of.

NetApp Volume Options

NetApp volumes should be created as Flexible volumes with the snap reserve set to 0 and the default Snapshot schedule should be disabled. All NetApp Snapshot copies need to be coordinated with the ESX Servers in order to ensure data consistency. NetApp Snapshot copies are covered in the Implementing Snapshot Copies section of this paper. To set the volume options for Snapshot copies to the recommended setting enter the following commands on the FAS System console:

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

RDMs and Data Center Sizing

At the time of this writing a VMware data center is limited to a total of 256 LUNs. This limitation typically only arises in deployments which will use RDMs as the primary form of Virtual Machine storage. With RDMs you will need to 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. It is recommended that RDM pointer files and Virtual Machine configuration files be kept on separate data stores.

In order to determine the number of ESX nodes that will be used by a single data center you will use the following formula:

254 / (the 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 would be: 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 which they point to; however, each mapping file is approximately 1 MB in size. It is recommended to 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/data stores provide the simplest method of provisioning storage; however, you should plan to limit the total size of the LUN/data store and to create multiple VMFS data stores in order to avoid disk I/O performance issues due to contention with file locking and SCSI reservations. While there is no definitive recommendation a common deployed size for a VMFS LUN/data store is between 300 and 700 GB. For more information please see page 35 in http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

STORAGE PROVISIONING

With the introduction of VMware Virtual Infrastructure 3.0 several new storage options were introduced. In this section we will cover 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 one needs to begin by creating initiator groups (igroups) on the FAS System. NetApp igroups provide a form of LUN masking which controls host access to a LUN. It is recommended that an igroup be created for each VMware data center. In addition the igroup should be named after or should include the name of the data center in its name. Each initiator group will include all of the FCP worldwide port names (WWPNs) or iSCSI qualified names (IQNs) installed on each of the ESX Servers within the VMware data center. For assistance in identifying the WWPN or IQN of the ESX Server select storage adapters from the configuration tab for each ESX Server within VirtualCenter and refer to the ‘San Identifier’ column. See Figure 4 below.

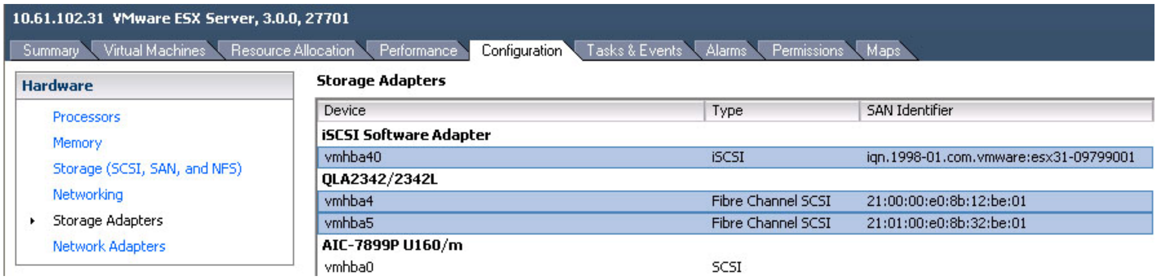

Figure 4

LUNs can be created via the NetApp LUN wizard in the FAS System console or via the FilerView® GUI. The procedures below demonstrate creating a LUN using the FilerView GUI.

1.	Log into FilerView.
2.	Select LUNs.
3.	Select Wizard.
4.	Within the Wizard window, select next.
5.	Enter the path (See Figure 5).
6.	Enter the LUN size.
7.	Enter the LUN Type (For VMFS select VMware, for RDM select the VM type).
8.	Enter the Description and select Next.

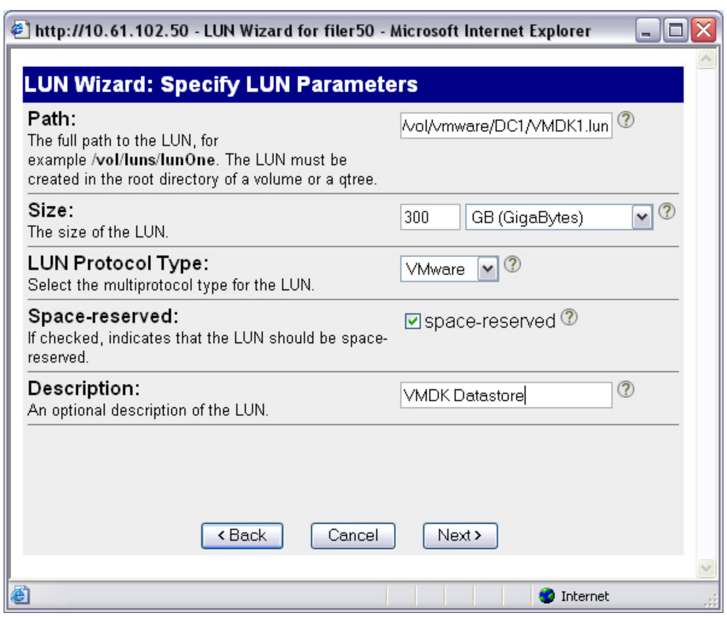

Figure 5

The next step in the LUN wizard is LUN masking. LUN masking is accomplished by assigning an igroup to a LUN. The LUN wizard allows you to either assign an existing igroup or create a new igroup. The steps below demonstrate how to configure LUN masking on a LUN created in the FilerView GUI.

1.	Select Add Group.
2.	Select the radio button for 'Use Existing Initiator Group". Select Next & Proceed to step 3a. **OR** Select the radio button for 'Create a new Initiator Group". Select Next & Proceed to step 3b.
3a.	Select the group from the list, and either assign a LUN ID or leave the field blank and the system will assign one. Select Next and you are finished.
3b.	To create a new group supply the igroup parameters including name, Connectivity type (FCP or iSCSI), and OS type (VMware) and select Next. (See Figure 6.)
4.	Enter in the new or select the known SAN Identifiers (WWPN or IQN) for the system(s) which will connect to this LUN.
5.	Select the Add Initiator button.
6.	Select Next to Complete.

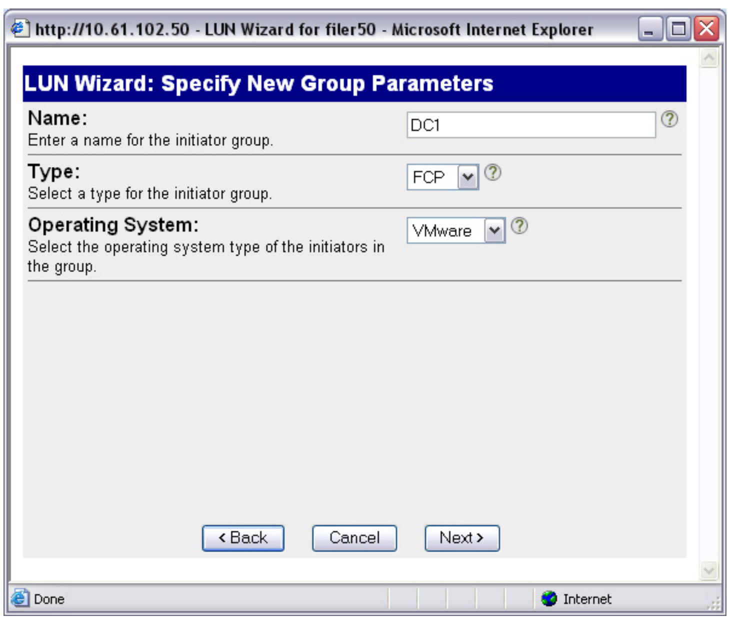

Figure 6

It is recommended that an igroup be created which contains all of the SAN Identifiers for all of the ESX hosts in a VMware data center. This method also simplifies LUN masking in that new LUNs need only be mapped to one igroup to be shared by all ESX Servers in a data center. In addition it is suggested to include in the name of the igroup the name of the data center and protocol type (i.e. DTW_DC1_FCP and DTW_DC1_iSCSI). This naming convention and method simplify the management of igroups by reducing the total number that will be created. Note: in the event that a data center will use both the Fibre Channel and iSCSI protocols, then separate igroups will need to be created for Fibre Channel and iSCSI.

NETWORK FILE SYSTEM (NFS) PROVISIONING

For those who would rather serve Virtual Disks via NFS the process is rather simple. Follow the steps below in order to create a file system for use as an NFS data store.

1	Open FilerView (http://filer/na_admin).
2	Select Volumes.
3	Select Add. This will open the Volume Wizard. See Figure 7 below. Complete the Wizard.
4	From the FilerView menu select NFS.
5	Select Add Export. This action will open the NFS Export Wizard. Complete the Wizard for the newly created file system granting read/write and root access to the ESX hosts that will connect to the exported file system.
6	After creating the file system export the file system via NFS to all of the ESX hosts contained within the VMware data center that will connect to the storage. This granted access needs to be read/write. See Figure 8 below.
7	Open VirtualCenter.
8	Select an ESX Host.
9	In the right pane select the Configuration tab.
10	Select the Storage link in the Hardware box.
11	In the upper right corner click Add Storage. The Add Storage Wizard will pop up. See Figure 9.
12	Select the Network File System radio button and click next.
13	Enter a name for the storage appliance, export, and data store, then click next. See Figure 10.
14	Click finish.

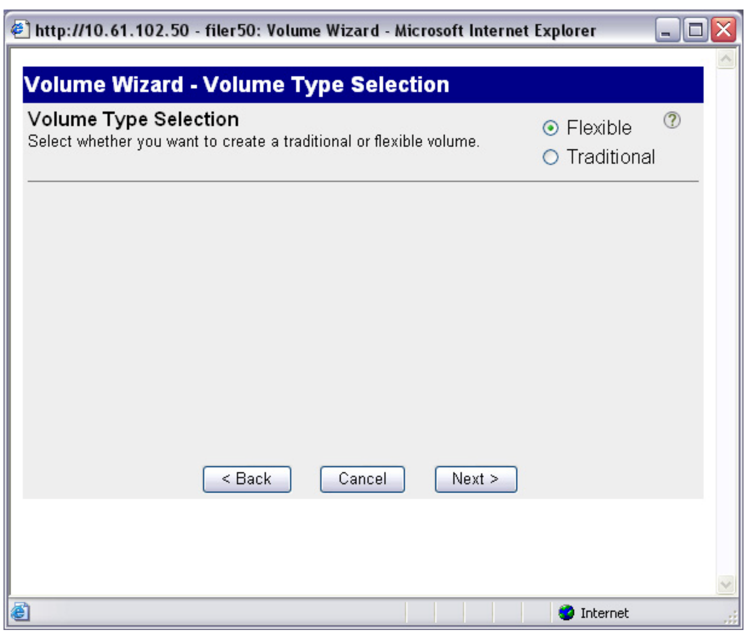

Figure 7

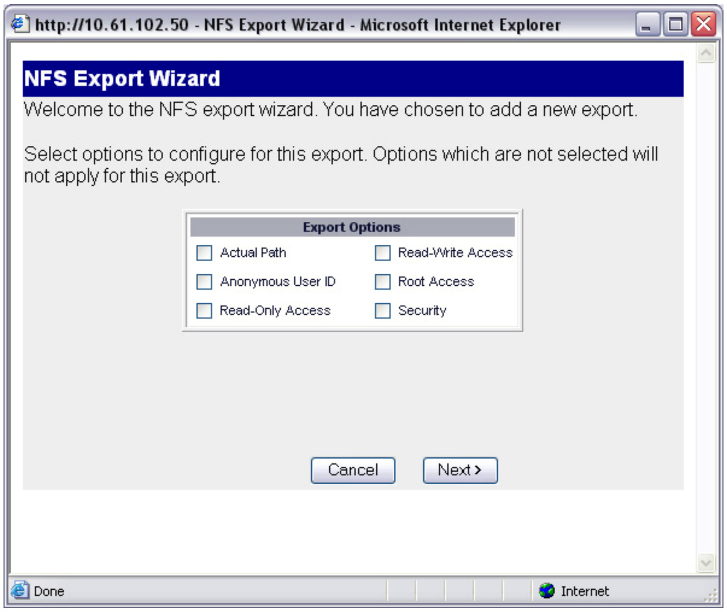

Figure 8

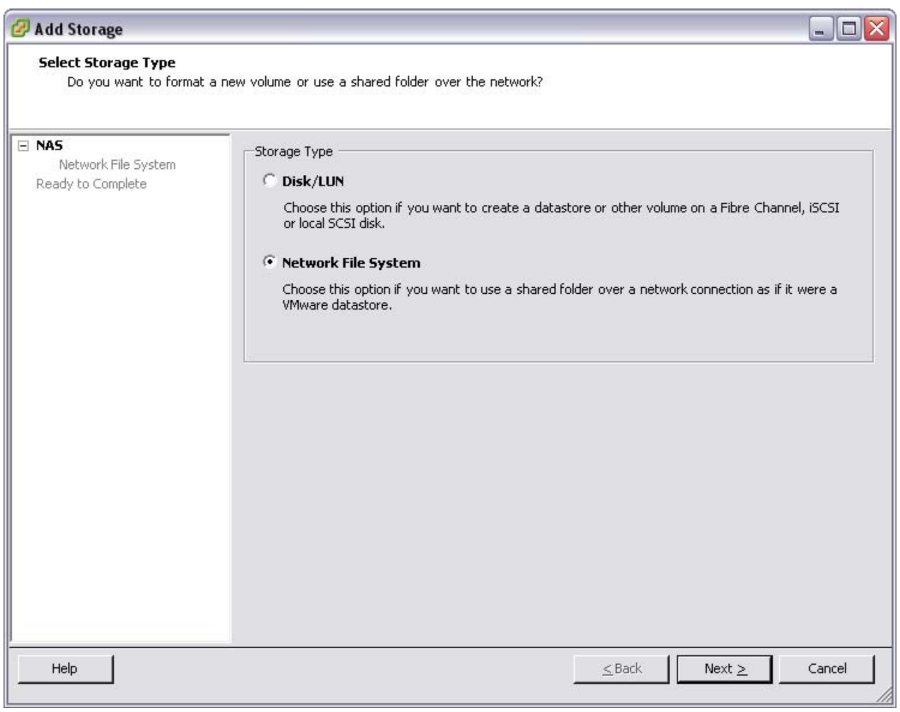

Figure 9

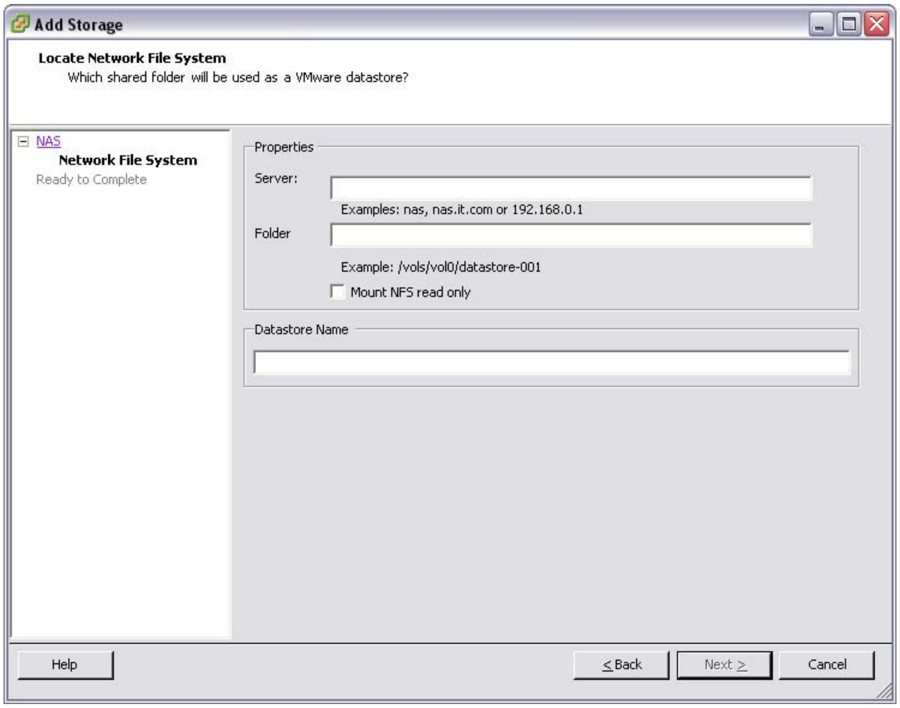

Figure 10

STORAGE CONNECTIVITY

With the introduction of VMware Virtual Infrastructure 3.0 several new storage connectivity options were introduced. In this section we will cover the available storage options and review settings specific to each technology.

FIBRE CHANNEL CONNECTIVITY

To begin you may notice that the Fibre Channel service is the only storage protocol which is running by default within ESX Server. It is recommended 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 complete the following steps.

1	Open VirtualCenter.
2	Select an ESX Host.
3	In the right pane select the Configuration tab.
4	Click the Storage Adapters link in the Hardware box.
5	In the upper right corner click the Rescan link.
6	Select the radio button for 'Create a new Initiator Group". Select Next & Proceed to step 6b.
7	Repeat steps 1-6 for each ESX Server in the data center.

Selecting rescan will force the rescanning of all HBAs (FC & iSCSI) to discover changes in the storage available to the ESX Server. Note: Some FCP HBAs require you to rescan them twice in order 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 a Raw Device Mapping or provisioned to the ESX Server as a data store. To add a LUN as a data store complete the following steps.

1	Open VirtualCenter.
2	Select an ESX Host.
3	In the right pane select the Configuration tab.
4	Select the Storage link in the Hardware box and then click Add Storage. This action will open the Add Storage Wizard. See Figure 11.
5	Select the Disk/LUN radio button and click next.
6	Choose the LUN you want to use and click next.
7	Enter a name for the data store which this LUN will be created and click Next.
8	Select the block size, click next, and finish.

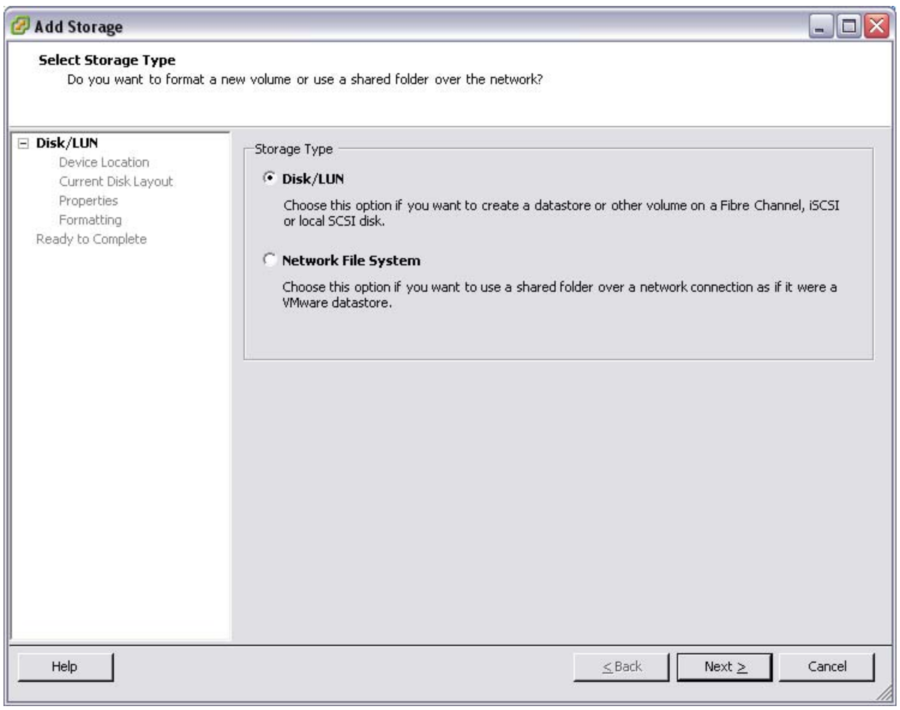

Figure 11

Note the default block size of a Virtual Machine File System is 1MB. This block size will support storing virtual disk files up to a maximum of 256GB in size. If you plan to store virtual disks larger than 256GB on the data store you will need to increase the block size to be greater than the default. See Figure 12.

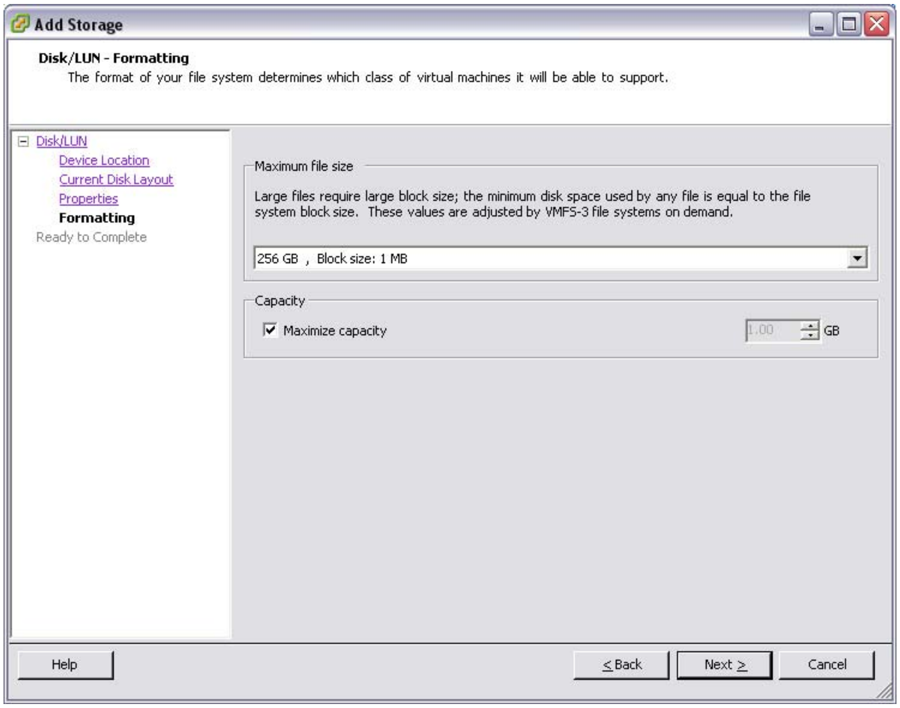

Figure 12

ISCSI / IP SAN CONNECTIVITY

It is a recommended best practice to separate iSCSI traffic from other IP network traffic. This practice is enforced as ESX Server requires a separate network connection type referred to as a VMkernel in order to enable iSCSI services. The VMkernel network will require an IP address which is currently not in use on the ESX Server. It is recommended that the VMkernel network for iSCSI traffic be a different network or VLAN than the network or VLAN used for the service console or for Virtual Machine network access. To configure the VMkernel complete the following steps.

1	Open VirtualCenter.
2	Select an ESX Host.
3	In the right pane select the Configuration tab.
4	Hardware box select Networking.
5	In the upper right cornet click Add Networking and the Add Networking Wizard will open. See Figure 13.
6	Select the VMkernel radio button and click next.
7	Select either an existing vSwitch or create a new one.
8	Click next.
9	Enter the IP address & subnet mask, click next, then finish. See Figure 14.
10	This action closes the Add Networking Wizard.
11	Select Security Profile in the Software box in the left pane of the Configuration tab.
12	In the right pane select the properties link and the Firewall Properties Window will open.

13	In the window select the Software iSCSI Client box and then click OK. See Figure 15.
14	This action closes the Firewall properties Box.
15	In the right pane select Storage Adapters in the Hardware box.
16	Highlight the iSCSI Adapter and click the properties link in the details box. See Figure 16.
17	Select the Dynamic Discovery tab on the iSCSI Initiator Properties box.
18	Select add and enter the IP address of the iSCSI enabled interfaces on the NetApp FAS system. See Figure 17.
19	For an additional layer of security CHAP authentication can be configured by selecting the tab of the same name. NetApp recommends setting up and verifying iSCSI access prior to enabling CHAP authentication.

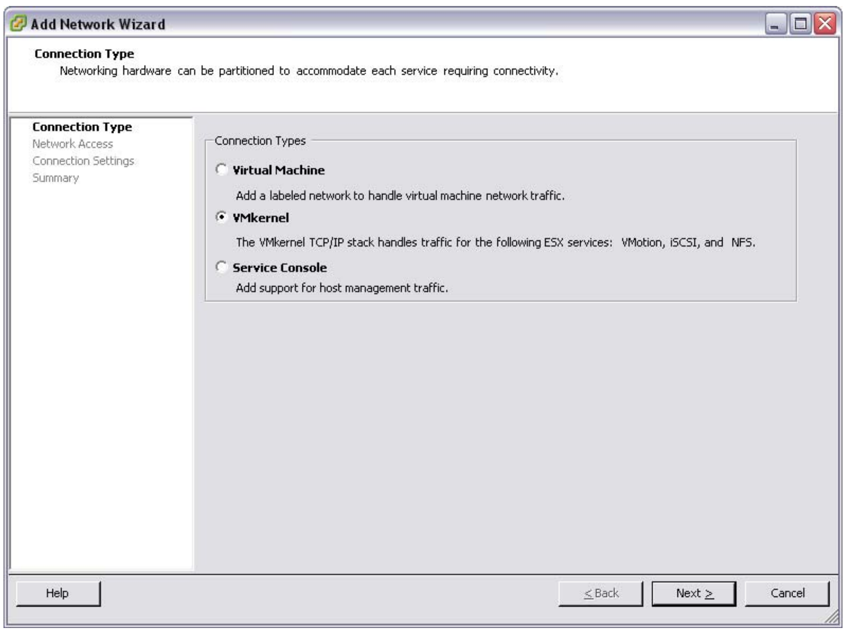

Figure 13

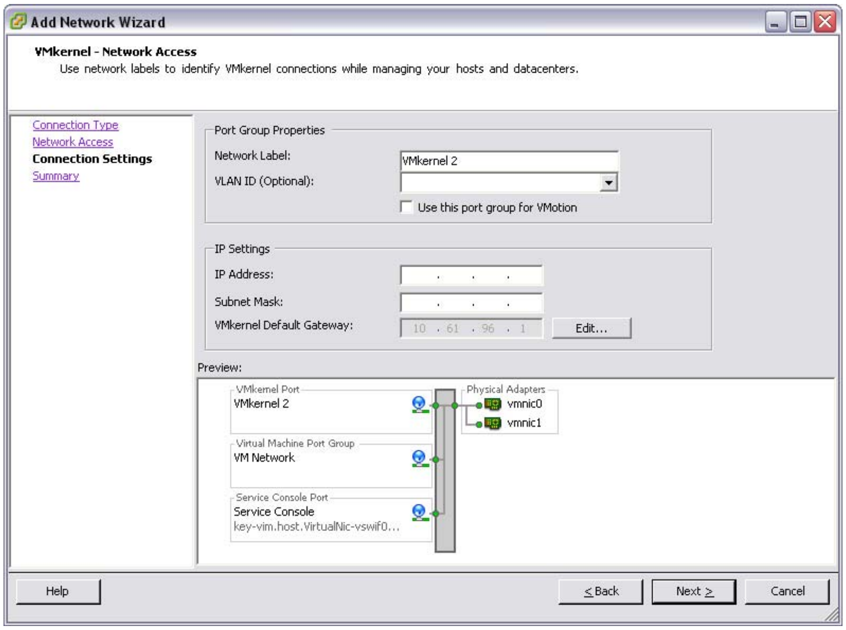

Figure 14

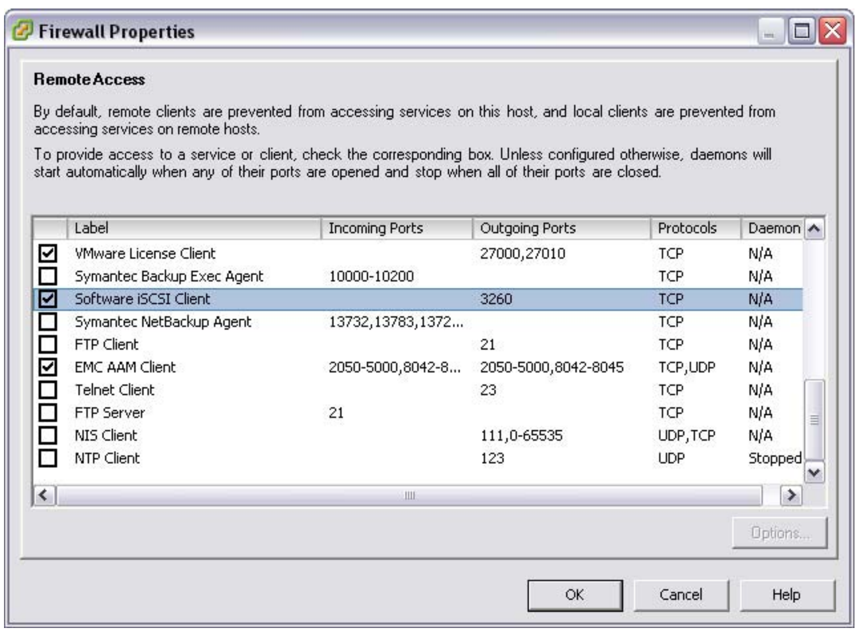

Figure 15

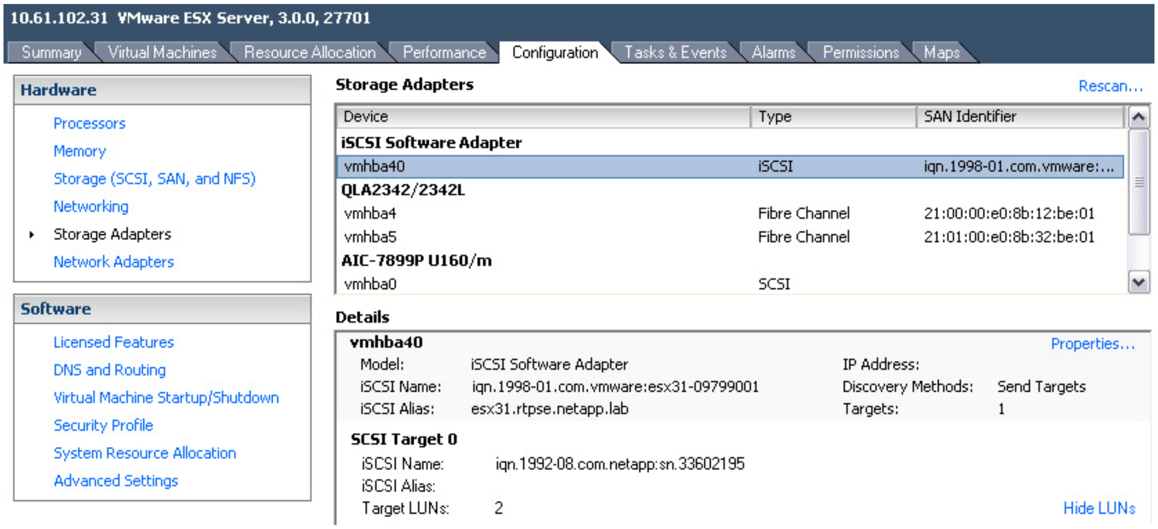

Figure 16

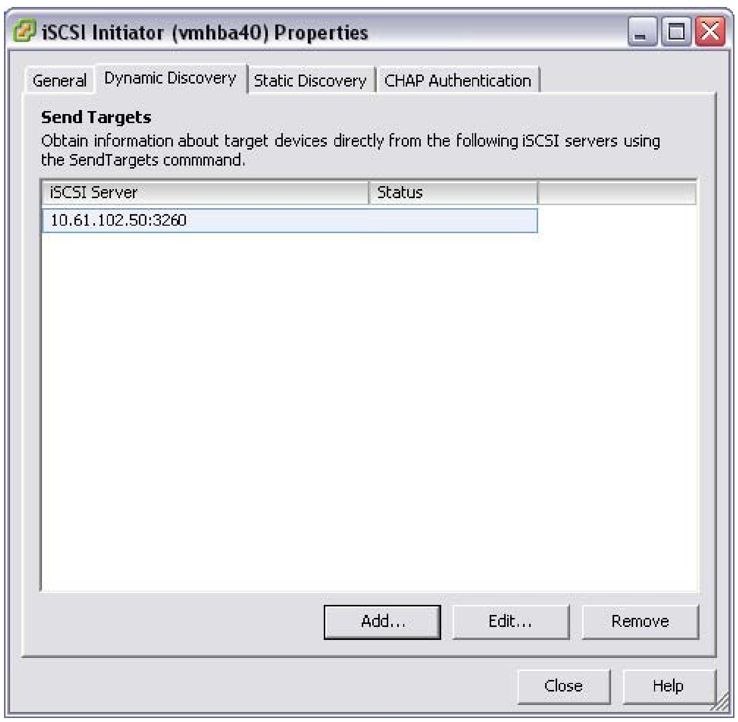

Figure 17

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. While the iSCSI host adapters provide the greatest performance and system scalability they do require additional NICs to be used in order to support all other IP operations and protocols. TOE enabled NICs handle all IP traffic as a traditional NIC in addition to the iSCSI traffic.

With the release of VMware ESX 3.0.1 support was added for iSCSI Host Bus Adapters and TOE enabled NICs. For maximum utilization of your ESX Servers it is highly recommended that you implement iSCSI HBAs or TOE enabled NICs over traditional NICs for iSCSI traffic.

iSCSI offers several options to addressing storage. If you are not ready to utilize iSCSI for your primary data access you may consider iSCSI for several other uses. iSCSI could be used to connect to data stores which store CDROM ISO images or iSCSI can be used as a redundant / failover path for a primary Fibre Channel path. If using this setup you will need to configure LUN multipathing. Please refer to the Fibre Channel Connectivity section.

NFS CONNECTIVITY

When using NFS connectivity for storage it is a recommended best practice to separate this NFS traffic from other IP network traffic. This practice is enforced as ESX Server requires a separate network connection type referred to as a VMkernel in order to enable NFS services. The VMkernel network will require an IP address which is currently not in use on the ESX Server. It is recommended that the VMkernel network for NFS traffic be a different network or VLAN than the network or VLAN used for the service console or for Virtual Machine network access.

NetApp offers TOE enabled NICs for use with IP traffic including NFS. It is recommended to use TOE enabled NICs for the greatest performance and system scalability from a NetApp FAS system. With the release of VMware ESX 3.0.1 support for TOE enabled NICs has been added. For maximum utilization of your ESX Servers it is highly recommended that you implement TOE enabled NICs over traditional NICs for NFS traffic.

NETAPP FIBRE CHANNEL MULTIPATHING

Network Appliance clustered FAS systems have an option known as CFMODE. This system controls the behavior of the system's Fibre Channel ports in the event a cluster failover occurs. If you are deploying a clustered solution providing storage for a VMware data center you will need to ensure the 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 please see the NetApp SAN Support Matrix. In order to verify the current cfmode complete the following steps.

1	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
2	Type *fcp show cfmode.*
3	If the cfmode needs to be changed, you can do so by entering *fcp set cfmode <mode type>.*

Standby cfmode may require more switch ports as 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 please 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 will need to configure ESX multipathing. When using multipathing VMware requires the default path to be selected for each LUN connected on each ESX Server. To set the paths complete the following process.

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 click the properties link. See Figure 18 below.
6	Inside of 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 19.
8	In the Change Path State window select the path as Preferred and Enabled and click OK. See Figure 20.

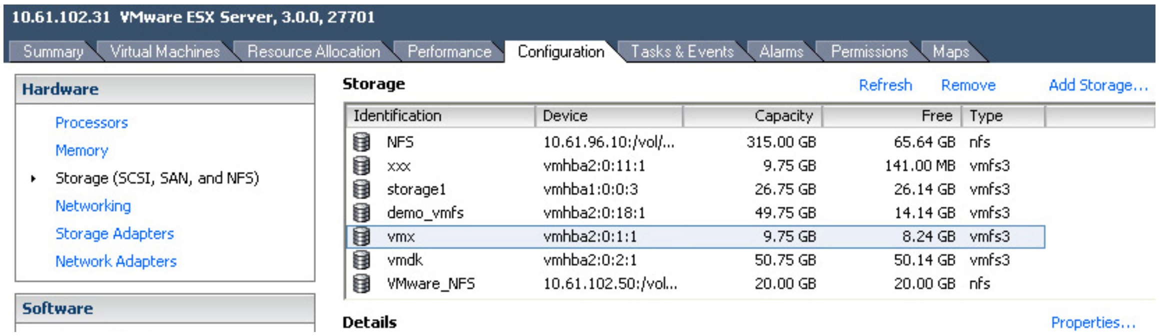

Figure 18

Figure 19

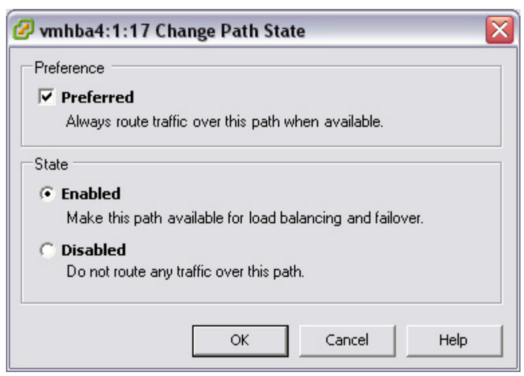

Figure 20

An alternative method for setting the preferred path for multiple LUNs is available within VirtualCenter. This method sets the preferred path for multiple LUNs. To set the paths complete the following process.

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 would like to configure.
7	Right click the highlighted LUNs, and select Manage Paths. See Figure 21.

8	Manage Path Window will open and allow you to set the multipathing policy and preferred path for all of the highlighted LUNs. See Figure 22.

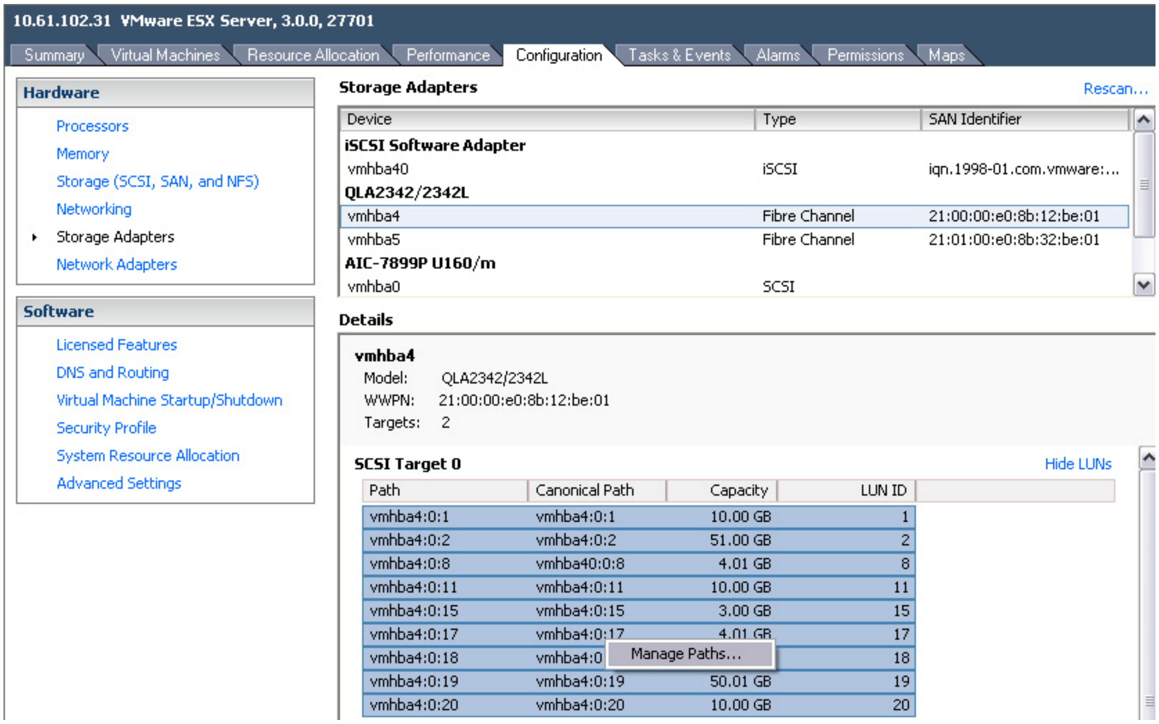

Figure 21

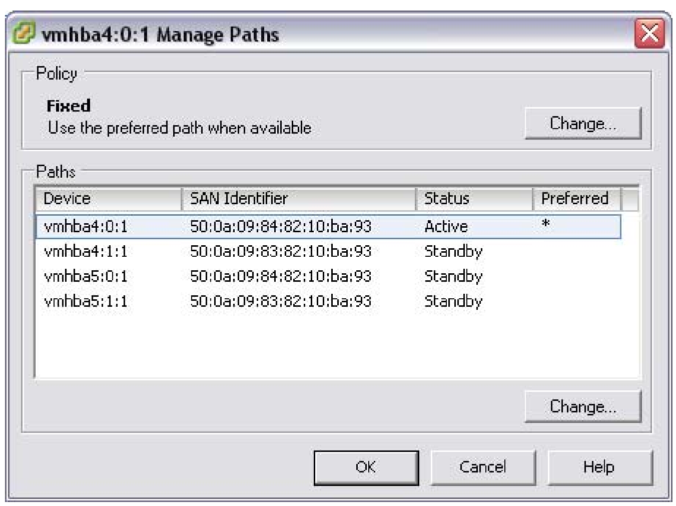

Figure 22

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 storage provisioned. By default server virtualization does not have any impact on improving storage utilization (and in many cases may have the opposite effect).

One should be very familiar with traditional storage provisioning in 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 very common practice for server administrators to overprovision storage in order to avoid running out of storage and the associated application downtime required with expanding the provisioned storage. While no system can be run at 100% storage utilization there are methods of storage virtualization which 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 while traditional provisioning preallocates storage. The value of thin provisioned storage is that storage is treated as a shared resource pool and storage is consumed as each individual VM requires it. This sharing increases the total utilization rate of storage by eliminating the unused yet provisioned areas of storage which are associated with traditional storage. The drawback to thin provisioning and oversubscribing storage is that should every VM require 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 via 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 which reside on NFS data stores are in the thin format by default. To implement thin VMDKs on FCP or iSCSI you will need to issue those commands via the ESX Server console (or script). NetApp thin provisioning significantly enhances the thin provisioning built into VMware by allowing VMFS data stores and RDM LUNs to be thin provisioned. This combination of technologies allows for any type of VM storage to be thin provisioned and allows for 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 in NetApp 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. If you would like to create a thin provisioned VMDK file you will need to use the 'vmkfstools' command with the –d options switch. By using VMware's thin provisioning technology the amount of storage consumed on a VMFS Data store can be reduced.

You should note that VMDKs which 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 for LUNs serving VMFS data stores to be provisioned to their total capacity limit yet only consume as much storage as what is required to store the VMDK files (which can be of either the thick or thin format). In addition, LUNs connected as RDMs can also be thin provisioned. In order to create a thin provisioned LUN complete the following steps.

1	Open FilerView (http://filer/na_admin).
2	Select LUNs.
3	Select Wizard.
4	Within the Wizard window, select 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 the Description and select Next.
9	Uncheck the box titled "Space-Reserved." See Figure 23.
10	Select next and Finish.

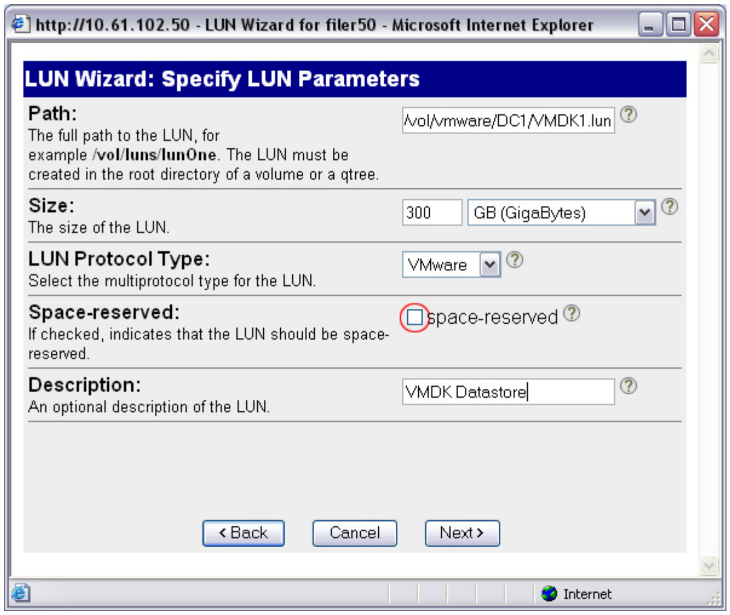

Figure 23

It is a recommended practice that when you enable NetApp thin provisioning that you also configure storage management policies on the volumes containing the thin provisioned LUNs. The use of these policies will aid 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 it is recommended to set this value to on. Doing so will require setting the maximum volume and increment size options. To enable these options complete the following process.

1	Log onto 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 which automatically deletes the oldest Snapshot copies on a volume when that volume is nearly full. For VMware environments it is recommended to set this value to delete Snapshot copies at 5% of available space. In addition you will want to set the volume option to have the system attempt to grow the volume before deleting Snapshot copies. To enable these options complete the following process.

1	Log onto 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 which is required when you use NetApp Snapshot copies on volumes containing VMware LUNs. This policy defines the amount of additional space reserved to guarantee LUN writes should a volume become 100% full. For VMware environments it is recommended to set this value to 20% if you have separated your temp, swap, pagefile, and other transient data onto other LUNs and volumes; otherwise leave this setting at its default of 100%. To enable this option complete the following process.

1	Log onto NetApp console.

2	Set volume Snapshot fractional reserve. *vol options <vol-name> fractional_reserve 20.*

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 within an organization. When utilizing NetApp thin provisioning it is recommended to deploy NetApp Operations Manager and to set 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 will ensure that additional storage can be made available before the aggregate should become completely full. For more information on setting up notifications within Operations Manger please 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 VMFS

The storage for VMFS file systems can be increased quite simply and without downtime as NetApp FAS systems allow for dynamic resizing of LUNs and VirtualCenter allows for VMFS to be expanded on the fly. The process for growing a data store is as follows:

1	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
2	Select LUNs.
3	Select Manage.
4	Select the LUN from the list in the left pane.
5	Enter the new size of the LUN in the Size box and hit Apply.
6	Open VirtualCenter.
7	Select an ESX Host.
8	In the right pane select the Configuration tab.
9	Select the Storage Adapters link in the Hardware box.
10	In the right pane select the HBA(s) and click the Rescan Link.
11	Select the Storage link in the Hardware box.
12	In the right pane right click the data store you want to grow and select properties.
13	Click Add Extent.
14	Select the LUN and click Next.
15	Select Next. Note as long as the window shows free space available on the LUN you can ignore the warning message. See Figure 24.
16	Ensure the Maximize Space checkbox is checked and select Next and Finish.

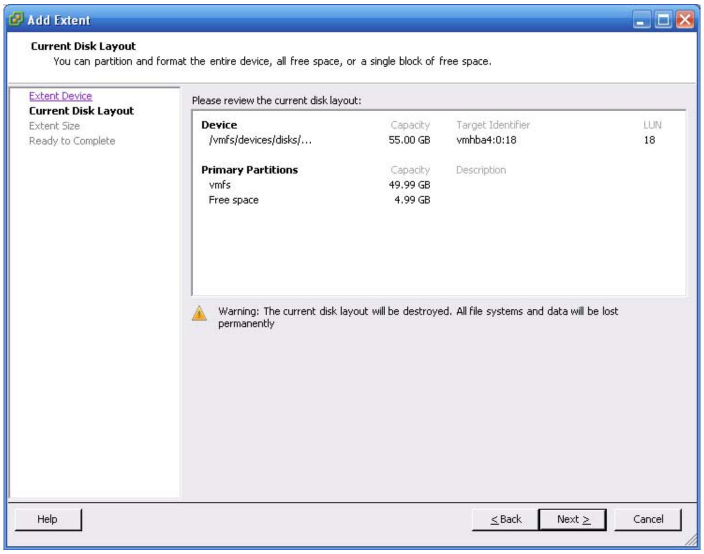

Figure 24

For more information on adding VMFS extents see page 130 of http://www.vmware.com/pdf/vi3_server_config.pdf.

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 to increasing available storage; you will still need to grow the file system after the VM boots. Please 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 the subsection below titled Growing Bootable Volumes; for all other volumes you can use native operating system tools to grow the volume. To grow a Virtual Disk complete the following 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 will need to grow the file system within the Virtual Disk. Follow the guidelines listed under Growing a VM's File System.

For more information on extending a Virtual Disk see page 285 of http://www.vmware.com/pdf/vi3_server_config.pdf.

Growing a Raw Device Mapping (RDM)

Growing a 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 complete the following process.

1	Open VirtualCenter.
2	Select an ESX Host and power down the VM.
3	Right click the VM and select Edit Settings. The Edit Settings window will open.
4	Highlight the Hard Disk to be resized, select Remove, and select the radio

	button Remove from Virtual machine and Delete Files from Disk. This action will delete the Mapping File but will NOT remove any data from the RDM LUN. See Figure 25.
5	Open FilerView (http://filer/na_admin).
6	Select LUNs.
7	Select Manage.
8	Select the LUN from the list in the left pane.
9	Enter the new size of the LUN in the Size box and hit Apply.
10	Open VirtualCenter.
11	In the right pane select the Configuration tab.
12	Select the Storage Adapters link in the Hardware box.
13	In the right pane select the HBA(s) and click the Rescan Link.
14	Right click the VM and select Edit Settings from the menu. The Edit Settings window will open.
15	Select Add.., then select Hard Disk and click Next. See Figure 26.
16	Continue on by selecting the LUN and clicking Next. See Figure 27.
17	Specify the VMFS data store which will store the Mapping File.
18	Start the VM and remember while you have grown the LUN you will still need to grow the file system within it. Follow the guidelines listed under Growing a VM's File System.

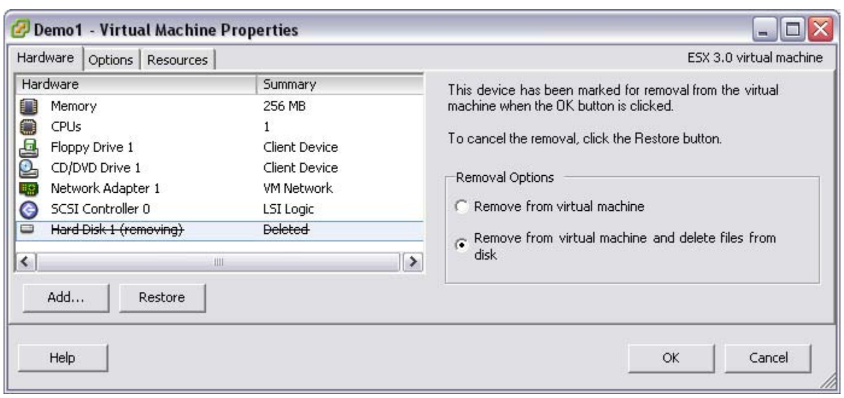

Figure 25

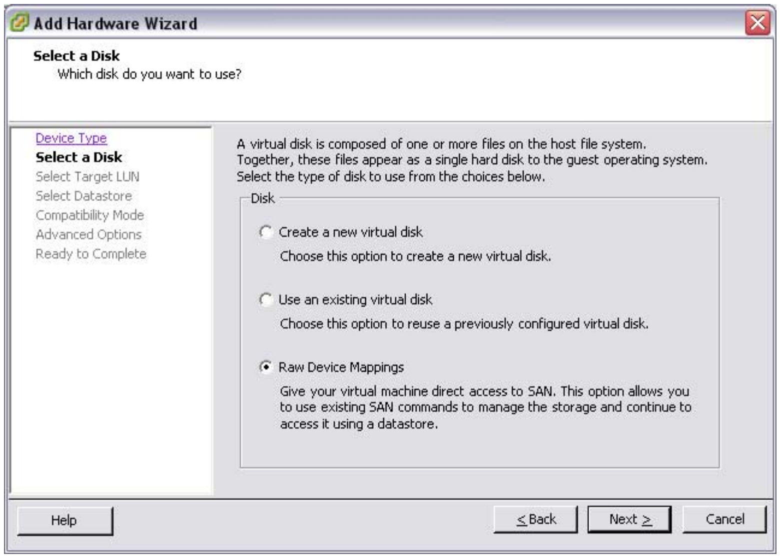

Figure 26

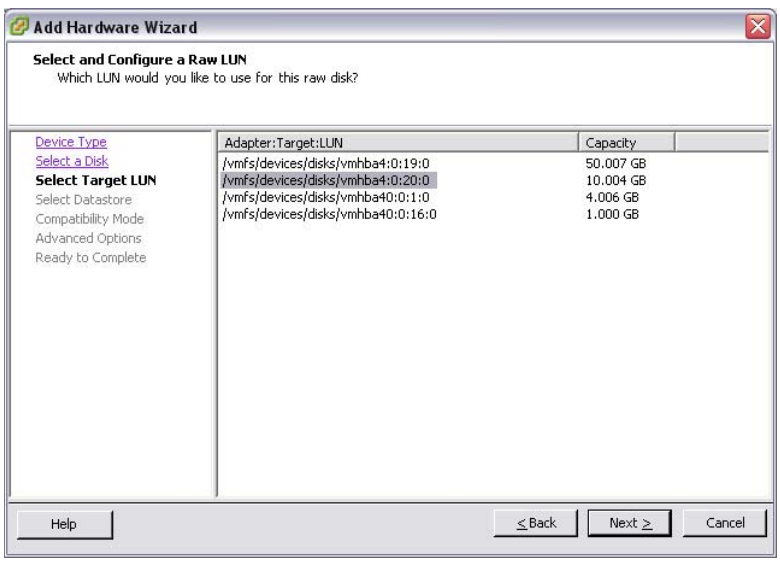

Figure 27

Growing a VM’s File System (NTFS or EXT3)

Once a virtual disk or RDM has been increased in size, you will 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 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 refer to: 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 refer to: http://sourceforge.net/projects/ext2resize.

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 rather simple way to expand these file systems which does not require the acquisition of any additional software (outside of ext2resize). This process requires the VMDK or LUN which has been resized to be connected to another Virtual Machine of the same operating system type using the processes defined above. Once the storage is connected the hosting VM can run the utility to extend the file system. After extending the file system this VM will be shut down and the storage disconnected. Connect the storage to the original VM and upon boot you will be able to 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 for point in time copies to be made which provide for 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 while the basic concepts of a snapshot is similar between NetApp and VMware one should be aware of the major differences between the two and when you should use one over the other.

VMware's snapshots provide for simple point in time versions of VMs allowing for quick recovery. The benefits of VMware snapshots are that they are easy to create and use as they can be executed and scheduled from within VirtualCenter. Some concerns with VMware's snapshots are that they are software based and most software based snapshot technologies tend to have significant scalability and performance issues. For more information on native VMware snapshots see page 167 of http://www.vmware.com/pdf/vi3_admin_guide.pdf and page 104 of http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf.

The patented NetApp Snapshot technology can easily be integrated into VMware environments and 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 which does not have a negative impact on system performance. VMware itself even states that for optimum performance & 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 more details see page 168 of http://www.vmware.com/pdf/vi3_admin_guide.pdf and page 142 of http://www.vmware.com/pdf/vi3_server_config.pdf.

DATA LAYOUT FOR SNAPSHOT COPIES

Within VMware Virtual Infrastructure 3.0 the storage provisioned to Virtual Machines is stored in either Virtual Disk Files (residing on VMFS LUNs or NFS mounts) or in Raw Device Mappings (RDM LUNs). With either storage option LUNs or Virtual Disk Files reside within NetApp volumes. NetApp Snapshot copies are created at the volume level, so it is important to group VM storage on NetApp volumes based on their Snapshot schedule.

For example, if you have a group of VMs which will create a Snapshot copy once a day and a second group which will create a Snapshot copy 4 times a day, then you will need to have a minimum of two NetApp volumes. For traditional Virtual Disks residing on VMFS each volume will contain a single LUN, for Virtual Disks residing on NFS each volume will have several Virtual Disk files, and for RDMS each volume will contain several RDM formatted LUNs. The Snapshot backup script will need to be configured for each volume containing VMs and the appropriate Snapshot schedule.

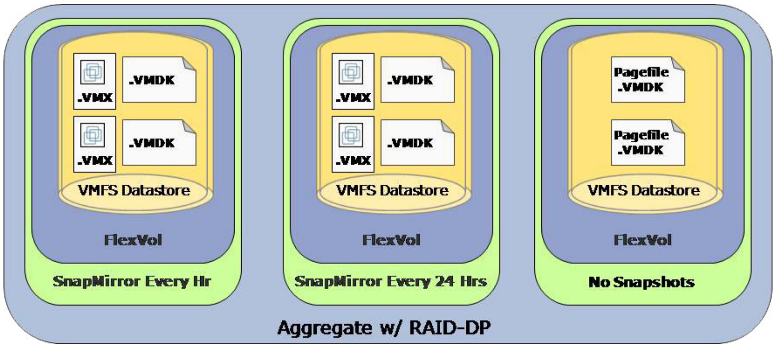

Figure 28

An example data layout configuration based on snap policies.

SNAPSHOT CONCEPTS

Within a VMware Virtual Infrastructure 3.0 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 & RDM LUNs from a production ESX Server. In order to enable this functionality execute the following commands on the console of the ESX Server:

esxcfg-advcfg -s 1 /LVM/EnableResignature

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

service mgmt-vmware restart

IMPLEMENTING SNAPSHOT COPIES

The consistency of the data contained within a Snapshot copy is paramount to a successful recovery. This section will review how to implement NetApp Snapshot copies for the purpose of VM recovery, cloning, or disaster recovery replication.

ESX SNAPSHOT CONFIGURATION FOR SNAPSHOT COPIES

With the introduction of VI3 administrators have the ability to mount storage created Snapshot copies of VMFS LUNs to production ESX Servers. In order to enable this functionality complete the following process.

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. The Advanced Settings window will pop up.
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 2-6 for each ESX Server in the Data Center.

ESX SERVER & NETAPP FAS SSH CONFIGURATION

The most efficient way to integrate NetApp Snapshot copies is to allow for centralized management and execution of Snapshot copies. Configuring the FAS systems and ESX Servers to allow for a single host to remotely execute commands on both systems is recommended. 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 complete the following process.

1	Connect to the FAS systems 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 onto the Linux or VMware system which will remotely execute commands on the FAS system as root.
4	Add the Triple DES cipher to the list of available SSH ciphers, as this is the only cipher recognized by the NetApp FAS system. Edit the file /etc/ssh/ssh_config 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; simply press [Enter]. The public key will be saved to /root/.ssh/id_dsa.pub.
6	Copy only the key information from the public key file to FAS System's /etc/sshd/root/.ssh/authorized_keys, removing all information but the key string preceded by the string "ssh-dsa" and a comment line. See the example below.
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 via SSH complete the following process.

1	Log onto 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 ssh_config

5	Change the follow 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 example listed below. 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 1-10 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

Snapshots of VMFS data stores allow for a quick method to recover a VM from a Snapshot copy. This section details the procedures required to complete this process. 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 the steps listed below.

1	Open VirtualCenter.
2	Select an ESX Host and power down the VM.
3	Log onto the ESX console as root.
4	Rename the VMDK files: *mv <current VMDK path> <renamed VMDK path>*
5	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
6	Clone the original LUN from a recent Snapshot, 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	Click the Storage Adapters link in the Hardware box.
11	In the upper right corner click the Rescan link. Scan for both new storage and VMFS Data stores.
12	The Snapshot VMFS data store will appear.
13	Log onto the ESX console as root.
14	Copy the Virtual Disks from the Snapshot data store to the Production VMFS: *cd <VMDK snapshot path>* *cp <VMDK> <production VMDK path>*
15	Open VirtualCenter.
16	Select the ESX and start the virtual machine.
17	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.
18	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
19	Delete the Snapshot copy LUN: *lun destroy –f <LUN path>*
20	In the upper right corner click the Rescan link. Scan for both new storage and VMFS data stores.

RECOVERING VIRTUAL MACHINES FROM AN RDM SNAPSHOT

RDMs allow for the quickest possible method to recover a VM from a Snapshot copy. This section details the procedures required to complete this process. In summary this process powers off the VM, restores the RDM LUN, and powers on the VM. To complete this process follow the steps listed below:

1	Open VirtualCenter.
2	Select an ESX Host and power down the VM.
3	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
4	Clone the original LUN from a recent Snapshot: *lun clone create <original LUN path> –b <original LUN path> <Snapshot name>*
5	Offline the current version of the LUN in use: lun offline <LUN path>
6	Map and online the cloned LUN: *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 systems 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 allows for a quick method to recover a VM from a Snapshot copy. This section details the procedures required to complete this process. In summary, this process powers off the VM, restores the VMDK, and powers on the VM. To complete this process follow the steps listed below.

1	Open VirtualCenter.
2	Select an ESX Host and power down the VM.
3	Log onto the ESX console as root.
4	Rename the VMDK files: *mv <current VMDK path> <renamed VMDK path>*
5	Connect to the FAS systems console (via either SSH, Telnet, or Console connection).
6	Restore the VMDK file from a recent Snapshot: *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 onto the ESX console as root.
11	Delete the renamed VMDK files: *rm <renamed VMDK path>*

SUMMARY

As you can see VMware's Virtual Infrastructure provides customers with 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 can provide cost savings, increased storage utilization, and enhanced data recovery.

This paper is not intended to be a definitive implementation or solutions guide. Many factors are not addressed in this document. Expertise may be required to solve user specific deployments. Please contact your local Network Appliance representative and schedule to speak with one of our VMware solutions experts. Comments on this technical report are welcome. Please contact the authors here.

EXAMPLE HOT BACKUP SNAPSHOT SCRIPT

This script provides managed, consistent, backups of virtual machines in a VMware Virtual Infrastructure 3 environment leveraging Network Appliance Snapshot technology. The script leverages the backup management functionality provided by VMware's Consolidated Backup software and NetApp's patented no high performance, space savings snapshot technology. The 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.

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 thus allowing multiple Recovery Point Objectives (RPO's) to be met with very little effort. Critical application server Virtual Machines can have Snapshots automatically created based upon a different schedule than second tier applications or test and development Virtual Machines.

The script works by leveraging the VM management functionality provided by VCB to identify the running virtual machines which globally reside on a Datastore. Once the virtual machines have been identified, the script places the running VMs into a hot backup mode which is a recoverable, crash consistent state. At this point the VM data is captured in a NetApp snapshot. The script even maintains multiple versions of snapshots.

Backing up VM's with VMBKUP.vbs

Backing up Virtual Machines (VM) with the VMBKUP.vbs script consists of 3 stages:

- Quiesce all of the VM's on a given Datastore
- Take a crash consistent NetApp snapshot
- Appling the REDO logs and restoring the Virtual Disk files to a read write state

NTAPVMBKUP Script Usage

The NTAPVMBKUP.vbs script allows for the control over which DataStores and NetApp volumes are backed up. The usage for the script follows:

Cscript //nologo VMBKUP.vbs [File IP Address] [VMFS DataStore Name] [NetApp Volume] [LogFile]

Script Requirements

VMware Consolidated Backup Framework installed in default location.

A username and password with read access to the Virtual Center Database, and execute permission for the VCB Framework utilities.

SSH access from the VCB server to the NetApp system(s)

Running the Script

This script can be scheduled to run via the Windows Task Scheduler

```
'=============================================================
'
' NAME:      VMBKUP.vbs
'
' AUTHOR:    Evan Battle
'                              Senior Technical Consultant
'                              Sirius Computer Solutions
'                              evan.battle@siriuscom.com
'
' DATE  :    4/25/2007
'
' VERSION:   1.05 - OpenSSH Support
'
' TO DO:     -Add support for NetApp Qtree SnapMirror Update
'
' COMMENT:            This script leverages Virtual Center to provide consistent, storage array based snapshots,
'                     for specific VMFS volumes, reather than Snaping all VI managed VM's.  This script
'                     gives you the ability to control snapshots one Volume at a time.
'
'                     This version of the script does NOT support Qtree SnapMirror.  Volume level SnapMirror
'                     can be used.
'
' USAGE:     cscript //nologo VMBKUP.vbs [Filer IP] [VMFS DataStore Name] [NetApp Volume] [Log File]
'
'=============================================================
'On Error Resume Next
'=============================================================
'            Global Variables
'=============================================================
Dim arrDataStoreInfo,arrVMInfo,arrHostInfo,arrSnapInfo
Dim arrVM_DSID,arrIDMapGlobal
Dim strDataStoreArg,strConnection,strServerInstance,strDatabase,strStorageIDGlobal,strFilerIP,strVolume
Dim objConn,objRS
'=============================================================
'            Custom Variables
'=============================================================
'\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/
strServerInstance = "##########################\msde_vc" '<<<< Set the Server\SQL Instance Here
strDatabase = "VCDB" '<<<< Set the VC Database here
strVCUserID = "########################" '<<<< User ID w/ access to Virtual Center
strVCPassword = "######################" '<<<< Password for the above User ID
intNumSnaps = 5 '<<<< Number of Array Based Snapshots to retain
'/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\/\
'strDataStoreArg = "storage1"
'strVolume = "ntap_volume_vmfs"
strStorageSubsystem = "NetApp"
strStorageAccess = "ssh" 'can be either rsh Or ssh << if ssh is to be used ensure Plink
                                                             '(http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html) Is saved to
root of the filesystem C:\
strvcbFrameworkPath = "c:\Program Files\VMware\VMware Consolidated Backup Framework"
'=============================================================
'            Global Constants
'=============================================================
Const adOpenStatic      =        3
Const adLockOptimistic = 3
Const adUseClient       =        3
Const ForReading        =        "1"
```

```
Const ForWriting		=		"2"
Const ForAppending		=		"8"
'==============================================================
'		This section identifies the arguments that are passed from
'		the command line
'==============================================================
If WScript.Arguments.Count > 0 Then
				strFilerIP = WScript.Arguments(0)
		Else
				Help
End If
If WScript.Arguments.Count > 1 Then
				strDataStoreArg = WScript.Arguments(1)
		Else
				Help
End If
If WScript.Arguments.Count > 2 Then
				strVolume = WScript.Arguments(2)
		Else
				Help
End If
If WScript.Arguments.Count > 3 Then
				strLogFile = WScript.Arguments(3)
		Else
				Help
End If
'==============================================================
'		Object Declaration
'==============================================================
wscript.echo strfilerIP
Set objConn = CreateObject("ADODB.Connection")
Set objRS = CreateObject("ADODB.Recordset")
Set objShell = CreateObject("WScript.Shell")
Set objFSO = CreateObject("Scripting.FileSystemObject")
CMD = objShell.ExpandEnvironmentStrings("%Comspec%")
strConnection = "driver=SQL Server;server="&strServerInstance&";database="&strDatabase
'==============================================================
'		Main Program Body
'==============================================================
WriteDestFile strLogFile,ForAppending,Now&" ########################################################"
WriteDestFile strLogFile,ForAppending,Now&" Begin Session"
WriteDestFile strLogFile,ForAppending,Now&" ########################################################"
Call GetVCDBID
Call GetHostInfo
Call GetVMInfo
Call GetDataStoreInfo
strStorageIDGlobal = VolumeID(strDataStoreArg)
WScript.Echo "The ID For "&strDataStoreArg&" is: "&strStorageIDGlobal
WriteDestFile strLogFile,ForAppending,Now&" The ID For "&strDataStoreArg&" is: "&strStorageIDGlobal
arrIDMapGlobal = VMTargetID(strStorageIDGlobal)
For a=0 To UBound(arrIDMapGlobal)
		For b=0 To UBound(arrVMInfo,2)
				If arrVMInfo(0,b)=arrIDMapGlobal(a) Then
						strDNS_NAME = arrVMInfo(1,b)
						strIP_ADDRESS = arrVMInfo(2,b)
						strGUEST_OS = arrVMInfo(3,b)
				End If
		Next
		If arrVMInfo(4,strIndex)="running" Then
				 Call QuiesceDisk (arrIDMapGlobal(a))
		End If
Next

If strStorageSubsystem="NetApp" Then Call NTAPSnapVolume(intNumSnaps)

Call GetSnapInfo
Call MatchSnap
WScript.Quit(0)

Sub NTAPSnapVolume(intNumSnaps1)
'==============================================================
'		This operation takes the NetApp snapshots And cycles the
'		snapshot names.  The newest snapshot ends in _snap.0
'
'		SSH Support Will Be added in a future version
'
'==============================================================
		x=intNumSnaps1-1
		WScript.Echo Now&" snap delete "&strVolume&"_snap."&x
		WriteDestFile strLogFile,ForAppending,Now&" snap delete "&strVolume&"_snap."&x
		If strStorageAccess="rsh" Then
				objShell.Run CMD & " /c rsh "&strFilerIP&" "&Chr(34)&"snap delete "&strVolume&" "&strVolume&"_snap."&x&Chr(34),0,True
		Elseif strStorageAccess="ssh" Then
				objShell.Run CMD & " /c ssh -l root "&strFilerIP&" "&Chr(34)&"snap delete "&strVolume&" "&strVolume&"_snap."&x&Chr(34),0,True
		End If
		While x > 0
				x=x-1
				WScript.Echo Now&" snap rename "&strVolume&"_snap."&x&" To "&strVolume&"_snap."&x+1
```

```
				WriteDestFile strLogFile,ForAppending,Now&" snap rename "&strVolume&"_snap."&x&" To "&strVolume&"_snap."&x+1
				If strStorageAccess="rsh" Then
					objShell.Run CMD & " /c rsh "&strFilerIP&" "&Chr(34)&"snap rename "&strVolume&" "&strVolume&"_snap."&x&"
"&strVolume&"_snap."&x+1&Chr(34),0,True
				Elseif strStorageAccess="ssh" Then
					objShell.Run CMD & " /c ssh -l root "&strFilerIP&" "&Chr(34)&"snap rename "&strVolume&" "&strVolume&"_snap."&x&"
"&strVolume&"_snap."&x+1&Chr(34),0,True
				End If
		Wend
		WScript.Echo Now&" snap create "&strVolume&"_snap."&x
		WriteDestFile strLogFile,ForAppending,Now&" snap create "&strVolume&"_snap."&x
		If strStorageAccess="rsh" Then
				objShell.Run CMD & " /c rsh "&strFilerIP&" "&Chr(34)&"snap create "&strVolume&" "&strVolume&"_snap."&x&Chr(34),0,True
		Elseif strStorageAccess="ssh" Then
				objShell.Run CMD & " /c ssh -l root "&strFilerIP&" "&Chr(34)&"snap create "&strVolume&" "&strVolume&"_snap."&x&Chr(34),0,True
		End If
End Sub

Sub MatchSnap
'==============================================================
'		Match the VM ID to the Script Generated Snapshot ID
'==============================================================
		For x = 0 To UBound(arrIDMapGlobal)
				For y = 0 To UBound (arrSnapInfo,2)
					If arrIDMapGlobal(x)=arrSnapInfo(1,y) And arrSnapInfo(2,y)="Quiesce" Then Call
UnQuiesceDisk(arrIDMapGlobal(x),arrSnapInfo(0,y))
				Next
		Next
End Sub

Function VolumeID(strStorageName)
'==============================================================
'		Extract the Datastore ID
'==============================================================
		For x = 0 To UBound(arrDataStoreInfo,2)
				If arrDataStoreInfo(1,x)=strStorageName Then VolumeID=arrDataStoreInfo(0,x)
		Next
End Function

Function VMTargetID(strStorageID)
'==============================================================
'		Gathers the VM ID's that reside on the Target Datastore
'==============================================================
		Dim cnt1
		Dim arrIDMap
		cnt1=0
		ReDim arrIDMap(cnt1)
		For x = 0 To UBound(arrVM_DSID,2)
				ReDim Preserve arrIDMap(cnt1)
				If arrVM_DSID(1,x)=strStorageID Then arrIDMap(cnt1)=arrVM_DSID(0,x)
				cnt1=cnt1+1
		Next
		VMTargetID=arrIDMap
End Function

Sub QuiesceDisk (intVMID)
'==============================================================
'		Quiesce the VM using the VCB Framework
'==============================================================
		If intVMID = "" Then Exit Sub
		WScript.Echo Now&" Taking Snapshot of Host: "& strDNS_NAME & " - IP Address: "& strIP_ADDRESS &" - Guest OS: "&strGUEST_OS
		WriteDestFile strLogFile,ForAppending,Now&" Taking Snapshot of Host: "& strDNS_NAME & " - IP Address: "& strIP_ADDRESS &" - Guest OS:
"&strGUEST_OS
		WScript.Echo Now&" Taking Snapshot of VM ID: "&intVMID
		WriteDestFile strLogFile,ForAppending,Now&" Taking Snapshot of VM ID: "&intVMID
		objShell.Run CMD & " /c "&Chr(34)&strvcbFrameworkPath&"\vcbsnapshot"&Chr(34)&" -h localhost -u "&strVCUserID&" -p "&strVCPassword&" -L 6 -
c moref:vm-"&intVMID&" Quiesce",0,True
End Sub

Sub UnQuiesceDisk (intVMID,intSnapID)
'==============================================================
' Un-Quiesce the VM using the VCB Framework
'==============================================================
		WScript.Echo Now&" Releasing Snapshot - VM ID: "&intVMID&", SnapshotID: "&intSnapID
		WriteDestFile strLogFile,ForAppending,Now&" Releasing Snapshot - VM ID: "&intVMID&", SnapshotID: "&intSnapID
		objShell.Run CMD & " /c "&Chr(34)&strvcbFrameworkPath&"\vcbsnapshot"&Chr(34)&" -h localhost -u "&strVCUserID&" -p "&strVCPassword&" -L 6 -
d moref:vm-"&intVMID&" ssid:snapshot-"&intSnapID,0,True
End Sub

Sub GetSnapInfo()
'==============================================================
'		Gather information on the VM Snapshots
'==============================================================
		Dim strSnapQuery
		Dim cnt1
		cnt1=0
		strSnapQuery = "Select * From VPX_SNAPSHOT"
		objConn.ConnectionString = strConnection
```

```
		objConn.Open
		If Err.Number = 0 Then
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
		Else
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error: "&Err.Number
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error:
"&Err.Number
		End If
		objRS.Open strSnapQuery, objConn, adOpenStatic, adLockOptimistic
		ReDim arrSnapInfo(2,cnt1)
		While Not objRS.EOF
			ReDim Preserve arrSnapInfo(2,cnt1)
			arrSnapInfo(0,cnt1) = objRS.Fields("ID").value
			arrSnapInfo(1,cnt1) = objRS.Fields("VM_ID").value
			arrSnapInfo(2,cnt1) = objRS.Fields("SNAPSHOT_NAME").value
			cnt1=cnt1+1
			objRS.MoveNext
		Wend
		objConn.Close
End Sub

Sub GetDataStoreInfo()
'=============================================================
'		Gather information on the Datastores
'=============================================================
		Dim strDSQuery
		Dim cnt1
		cnt1=0
		strDSQuery = "Select * From VPX_DATASTORE"
		objConn.ConnectionString = strConnection
		objConn.Open
		If Err.Number = 0 Then
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
		Else
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error: "&Err.Number
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error:
"&Err.Number
		End If
		objRS.Open strDSQuery, objConn, adOpenStatic, adLockOptimistic
		ReDim arrDataStoreInfo(5,cnt1)
		While Not objRS.EOF
			ReDim Preserve arrDataStoreInfo(5,cnt1)
			arrDataStoreInfo(0,cnt1) = objRS.Fields("ID").value
			arrDataStoreInfo(1,cnt1) = objRS.Fields("NAME").value
			arrDataStoreInfo(2,cnt1) = objRS.Fields("CAPACITY").value
			arrDataStoreInfo(3,cnt1) = objRS.Fields("FREE_SPACE").value
			arrDataStoreInfo(4,cnt1) = objRS.Fields("TYPE").value
			cnt1=cnt1+1
			objRS.MoveNext
		Wend
		objConn.Close
End Sub

Sub GetVMInfo()
'=============================================================
'		Gather information on the Virtual Machines
'=============================================================
		Dim strVMQuery
		Dim cnt1
		cnt1=0
		strVMQuery = "Select * From VPX_VM"
		objConn.ConnectionString = strConnection
		objConn.Open
		If Err.Number = 0 Then
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
		Else
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error: "&Err.Number
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed With Error:
"&Err.Number
		End If
		objRS.Open strVMQuery, objConn, adOpenStatic, adLockOptimistic
		ReDim arrVMInfo(9,cnt1)
		While Not objRS.EOF
			ReDim Preserve arrVMInfo(9,cnt1)
			arrVMInfo(0,cnt1) = objRS.Fields("ID").value
			arrVMInfo(1,cnt1) = objRS.Fields("DNS_NAME").value
			arrVMInfo(2,cnt1) = objRS.Fields("IP_ADDRESS").value
			arrVMInfo(3,cnt1) = objRS.Fields("GUEST_OS").value
			arrVMInfo(4,cnt1) = objRS.Fields("GUEST_STATE").value
			arrVMInfo(5,cnt1) = objRS.Fields("MEM_SIZE_MB").value
			arrVMInfo(6,cnt1) = objRS.Fields("NUM_VCPU").value
			arrVMInfo(7,cnt1) = objRS.Fields("NUM_NIC").value
			arrVMInfo(8,cnt1) = objRS.Fields("BOOT_TIME").value
			cnt1=cnt1+1
			objRS.MoveNext
```

```
		Wend
		objConn.Close
End Sub

Sub GetHostInfo()
'==============================================================
'		Gather the Host information
'==============================================================
		Dim strHostQuery
		Dim cnt1
		cnt1=0
		strHostQuery = "Select * From VPX_HOST"
		objConn.ConnectionString = strConnection
		objConn.Open
		If Err.Number = 0 Then
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
		Else
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed with Error: "&Err.Number
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed with Error: 
"&Err.Number
		End If
		objRS.Open strHostQuery, objConn, adOpenStatic, adLockOptimistic
		ReDim arrHostInfo(13,cnt1)
		While Not objRS.EOF
			ReDim Preserve arrHostInfo(13,cnt1)
			arrHostInfo(0,cnt1) = objRS.Fields("DNS_NAME").value
			arrHostInfo(1,cnt1) = objRS.Fields("IP_ADDRESS").value
			arrHostInfo(2,cnt1) = objRS.Fields("VMOTION_ENABLED").value
			arrHostInfo(3,cnt1) = objRS.Fields("PRODUCT_VERSION").value
			arrHostInfo(4,cnt1) = objRS.Fields("PRODUCT_BUILD").value
			arrHostInfo(5,cnt1) = objRS.Fields("HOST_VENDOR").value
			arrHostInfo(6,cnt1) = objRS.Fields("HOST_MODEL").value
			arrHostInfo(7,cnt1) = objRS.Fields("CPU_MODEL").value
			arrHostInfo(8,cnt1) = objRS.Fields("CPU_COUNT").value
			arrHostInfo(9,cnt1) = objRS.Fields("CPU_CORE_COUNT").value
			arrHostInfo(10,cnt1) = objRS.Fields("CPU_HZ").value
			arrHostInfo(11,cnt1) = objRS.Fields("MEM_SIZE").value
			arrHostInfo(12,cnt1) = objRS.Fields("NIC_COUNT").value
			arrHostInfo(13,cnt1) = objRS.Fields("HBA_COUNT").value
			cnt1=cnt1+1
			objRS.MoveNext
		Wend
		objConn.Close
End Sub

Sub GetVCDBID()
'==============================================================
'		Gather the VM to DS mappings
'==============================================================
		Dim strIDQuery
		Dim cnt1
		cnt1=0
		strIDQuery = "Select * FROM VPXV_VM_DATASTORE"
		objConn.ConnectionString = strConnection
		objConn.Open
		If Err.Number = 0 Then
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Successful"
		Else
			WScript.Echo Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed with Error: "&Err.Number
			WriteDestFile strLogFile,ForAppending,Now&" - Connection To "&strServerInstance&" - "&strDatabase&" Failed with Error: 
"&Err.Number
		End If
		objRS.Open strIDQuery, objConn, adOpenStatic, adLockOptimistic
		ReDim arrVM_DSID(1,cnt1)
		While Not objRS.EOF
			ReDim Preserve arrVM_DSID(1,cnt1)
			arrVM_DSID(0,cnt1) = objRS.Fields("VM_ID").value
			arrVM_DSID(1,cnt1) = objRS.Fields("DS_ID").value
			cnt1=cnt1+1
			objRS.MoveNext
		Wend
		objConn.Close
End Sub

Sub WriteDestFile(DestFile,Mode,DestLine)
'==============================================================
'		This subroutine creats a log file based on the parameters
'		passed to it
'==============================================================
 Dim OutputFile
 Set OutputFile = objFSO.OpenTextFile(DestFile, Mode, True)
 OutputFile.WriteLine(DestLine)
 OutputFile.Close
End Sub

Function Help ()
```

```
'============================================================
'               Echo's the Help String
'============================================================
                Dim strUsage
                strUsage = vbcrlf&"******************************************************************************"&_
                                vbcrlf&"******************************************************************************"&vbcrlf&vbcrlf&_
                                    "Usage:"&vbcrlf&_
                                    "cscript //nologo VMBKUP.vbs "&Chr(34)&"Filer IP Address"&Chr(34)&" "&Chr(34)&"VMFS Datastore Name"&Chr(34)&" "&Chr(34)&"NetApp Volume"&Chr(34)&" logfile.Log"&Chr(34)&vbcrlf&vbcrlf&_
                                    "*This script will create NetApp snapshots of VMFS on a per volume basis"&vbcrlf&_
                                vbcrlf&"Example:"&vbcrlf&_
                                                "cscript //nologo VMBKUP.vbs 192.168.1.1 VMFS_LUN esx_vol c:\outputlog.log"&vbcrlf&_
                                vbcrlf&"******************************************************************************"&_
                                vbcrlf&"******************************************************************************"&vbcrlf
                WScript.Echo(strUsage)
                WScript.Quit(0)
End Function
```

REFERENCES

TR3428: Network Appliance & VMware ESX Server
Instantaneous Backup & Recovery Including Single File Restoration with NetApp Snapshot Technology
http://www.netapp.com/library/tr/3428.pdf

TR3393: Using Network Appliance Snapshot Technology with VMware ESX Server

http://www.netapp.com/library/tr/3393.pdf

TR3401: Using a Network Appliance SAN with VMware to Facilitate Storage and Server Consolidation

http://www.netapp.com/library/tr/3401.pdf

Wikipedia RAID Definitions and Explanations

http://en.wikipedia.org/wiki/Redundant_array_of_independent_disks

TR3001: A Storage Network Appliance

http://www.netapp.com/library/tr/3001.pdf

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

http://www.netapp.com/library/ar/ar1038.pdf

VMware Introduction to Virtual Infrastructure

http://www.vmware.com/pdf/vi3_intro_vi.pdf

VMware Server Configuration Guide

http://www.vmware.com/pdf/vi3_server_config.pdf

VMware SAN Configuration Guide

http://www.vmware.com/pdf/vi3_esx_san_cfg.pdf

VMware VMworld Conference Sessions Overview

http://www.vmware.com/vmtn/vmworld

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

http://www.netapp.com/library/tr/3466.pdf

TR3347: FlexClone Volumes: A Thorough Introduction

http://www.netapp.com/library/tr/3347.pdf

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

http://www.netapp.com/library/tr/3348.pdf

SnapMirror Software Overview

http://www.netapp.com/products/software/snapmirror.html

TR3446: SnapMirror Best Practices Guide

http://www.netapp.com/library/tr/3446.pdf

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 & instructions Added Figure 28

NetApp®

www.netapp.com