NETAPP TECHNICAL REPORT

DISK-TO-DISK BACKUP IN A NETAPP/VMWARE ENVIRONMENT

Jeremy Merrill, NetApp
Eric Hardcastle, VMware, Inc.
TR-3610

ABSTRACT

This document describes how NetApp products increase the storage availability of a VMware ESX environment by providing disk-to-disk backup with SnapVault®. It outlines a plan for the setup, configuration, and functional testing of a disk-to-disk backup environment. Specific equipment, software, and functional tests are included along with results.

TABLE OF CONTENTS

1 EXECUTIVE SUMMARY

As compute servers become more and more powerful, customers are turning to software virtualization technologies that enable them to take advantage of the power provided by these servers. While these virtualization techniques increase the utilization of the server, backup considerations need to be kept in mind for any virtualization architecture (or design). NetApp and VMware have combined resources to keep this piece of data center design relatively simple, yet robust and scalable at the same time. By using NetApp SnapVault technology with VMware Virtual Infrastructure 3 software, customers can rely on secure, fast, and simple backups of their existing infrastructure.

This technical report discusses how to implement NetApp SnapVault in a VMware virtual infrastructure environment.

1.1 DOCUMENT PURPOSE

The intent of this document is to provide an example of how NetApp products provide a disk-to-disk solution for VMware ESX in a NetApp environment.

The purpose of this reference configuration is to show

- How NetApp products and Host Virtualization can work together to provide a robust disk-to-disk backup solution
- How NetApp Snapshot™ and SnapVault technology provide an efficient backup and restore solution

This document does not include performance-related information nor is it intended to be any kind of formal certification.

1.2 ASSUMPTIONS

Throughout this document, the authors assume that we have two physical systems named “ESX-Primary” and “ESX-Secondary.” ESX-Primary is representative of the storage system (or site) in which ESX is running its normal operations. ESX-Secondary is the backup target system, which could reside within the same data center or a DR site. Also, since SnapVault is qtree based, all VMware datastores (LUN and NFS) are created in a qtree.

The authors also assume that the reader has a good understanding of NetApp Snapshot and SnapVault technologies. This document will demonstrate how to enable SnapVault for replication of crash consistent backups. If consistent backups are required, please refer to TR-3428 for an example script to create consistent Snapshot copies.

2 PRODUCT OVERVIEW

2.1 VMWARE

VMware products provide enterprise-class virtual machines that increase server and other resource utilization, improve performance, increase security, and minimize system downtime, reducing the cost and complexity of delivering enterprise services. By leveraging existing technology, VMware enables the roll-out of new applications with less risk and lower platform costs.

VIRTUAL INFRASTRUCTURE 3

VMware Infrastructure 3 is a feature-rich suite that delivers the production-proven efficiency, availability, and dynamic management needed to create a responsive data center. The suite includes:

- VMware ESX Server. Platform for virtualizing servers, storage, and networking.
- VMware VMFS. High-performance cluster file system for storage virtualization.

- VMware Virtual SMP. Multi-processor support for virtual machines.
- VMware VirtualCenter. Centralized management, automation, and optimization for IT infrastructures.
- VMware High Availability (HA). Cost-effective high availability for virtual machines.
- VMware DRS. Dynamic balancing and allocation of resources for virtual machines.
- VMware VMotion. Live migration of virtual machines without service interruption.
- VMware Consolidated Backup. Centralized backup software for virtual machines.

2.2 NETAPP

SnapVault is based on the replication and vaulting of Snapshot backups. In addition to being fast, SnapVault is very storage and network efficient, allowing a greater number of backups to be affordably stored on disk. SnapVault technology can also efficiently move backups over existing networks to centralize backups or move them off-site for disaster protection.

Traditional backups store data on tape in proprietary formats accessible only through the backup application. SnapVault creates online backup archives in an accessible file system format. Users and/or administrators can securely and rapidly search and recover data directly from backups. SnapVault also provides near-instant disaster recovery by enabling users and applications to failover directly to the backup copy.

FlexClone® and **FlexVol**® technologies enable entirely new opportunities and ways of working for organizations grappling with the challenges of increased overhead, management costs, and data risk. NetApp FlexVol technology delivers true storage virtualization solutions that can lower overhead and capital expenses, reduce disruption and risk, and provide the flexibility to adapt quickly and easily to the dynamic needs of the enterprise. FlexVol technology pools storage resources automatically and enables you to create multiple flexible volumes on a large pool of disks.

NetApp FlexClone technology enables true cloning/instant replication of data volumes and data sets without requiring additional storage space at the time of creation. Each cloned volume is a transparent, virtual copy that you can use for essential enterprise operations, such as testing and bug fixing, platform and upgrade checks, multiple simulations against large data sets, remote office testing and staging, and market-specific product variations.

Open Systems SnapVault provides the advantages of NetApp SnapVault advanced backup and recovery technology to heterogeneous storage environments. Open Systems SnapVault enables automated backups, rapid Snapshot copies, low bandwidth utilization, and reduced storage requirements. Today, Open Systems SnapVault can be installed within a virtual machine to back up the file systems inside the guest OS. While this will provide fast and efficient backups of files, it isn't able to efficiently recover an entire virtual machine. For more information on Open Systems SnapVault, please see the Open Systems SnapVault Best Practices Guide.

This test will also verify the functional interaction of the following NetApp products both during normal operation and under the defined failure scenarios:

- SnapVault
- FlexClone

3 TIERS OF PROTECTION

The key to this architecture is that the entire primary storage system can be backed up to the secondary storage system via SnapVault, then cloned via FlexClone for testing or development against data created by the primary site. In addition, with Data ONTAP® 7.2.1 or later, the primary and secondary systems can be an active-active system with one controller acting as the SnapVault primary system and the other controller the SnapVault secondary system. This configuration will enable SnapVault to move the primary data from Fibre Channel drives to lower cost, more dense ATA drives.

The overall solution uses NetApp SnapVault to replicate the data between NetApp storage systems. It also includes two ESX servers at the primary site; at the secondary site is a single ESX server used for dev/test

and recovery purposes. The servers are running ESX 3.0 with six virtual machines (VMs) on each. Windows® 2003 will be the guest operating systems and five virtual machines will be accessing their storage via the iSCSI protocol. There will also be one virtual machine accessing its data via NFS to help demonstrate the flexibility of NetApp and VMware. Figure 1 depicts the general layout of components used in this sample configuration.

For an additional level of protection, a third NetApp storage system could be added to the solution. In the event the primary and secondary systems are located in the same data center, you could utilize NetApp's Volume SnapMirror (VSM) technology to replicate the data to the tertiary site for added disaster recovery protection. In addition, in the event the secondary system became unavailable, SnapVault transfers could continue to occur to the tertiary system.

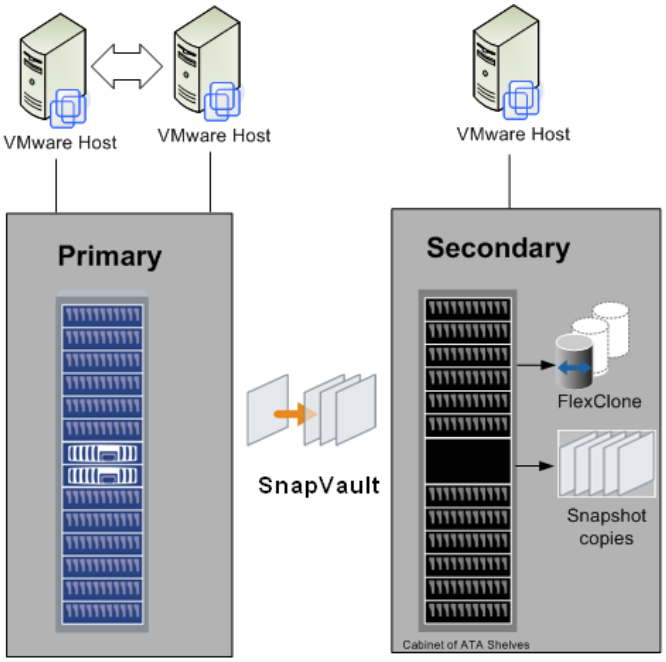

Figure 1) Topology.

4 OPERATIONAL SCENARIOS

The following subsections include test scenarios that will be executed upon successful build of the solution previously described in this document. The purpose of these scenarios is to examine and document, from a backup perspective, the reaction of a VMware/NetApp environment to various scenarios.

4.1 NORMAL BACKUP OPERATIONS

In normal backup operations, SnapVault will transfer the data at specified times. The native way of scheduling a SnapVault transfer is to create a Snapshot creation schedule on the SnapVault primary system and a SnapVault transfer (update) schedule on the secondary system. If the VM is required to be in a consistent state, then a script will be required to create a consistent Snapshot copy. If this is the case, then only a SnapVault retention schedule will be created on both the primary and secondary systems.

Task	First, perform the baseline transfer, then configure SnapVault Snapshot create and transfer schedules.
	
	1. To perform the baseline, use the `snapvault start` command. Using this command, you must specify the source and destination qtree (please refer to the SnapVault man pages for all options). This command is executed from the SnapVault secondary system.
	
	`ESX-SECONDARY> snapvault start -S 10.61.132.10:/vol/VMNFS/sv_source /vol/VMNFS_SV/sv_dest`
	`Snapvault configuration for the qtree has been set.`
	`Transfer started.`
	`Monitor progress with 'snapvault status' or the snapmirror log.`
	
	A relationship must be created for each qtree on the primary system.

	2. Once you have created the relationship, you must specify the backup schedule using the `snapvault snap sched` command. On the primary system, a SnapVault create schedule is created and on the destination system a SnapVault transfer schedule is created. To verify the schedule, use the `snapvault snap sched` command.
	
	On the primary system the schedule is set for hourly Snapshot copies at 12 AM, 6 AM, 12 PM, and 6 PM, but only retain the four most recent Snapshot copies.
	`ESX-PRIMARY> snapvault snap sched  VMNFS sv_hourly 4@0,6,12,18`
	`ESX-PRIMARY> snapvault snap sched`
	`create VMNFS sv_hourly 4@0,6,12,18`
	`create VM_VOL sv_hourly 4@0,6,12,18`
	
	On the secondary system the transfer schedule is set to run at the same intervals as the primary system, but we will retain 16 hourly Snapshot copies instead of four.
	`ESX-SECONDARY> snapvault snap sched -x VMNFS_SV sv_hourly 16@0,6,12,18`
	`ESX-SECONDARY> snapvault snap sched`
	`create VMNFS_SV  0@-`
	`xfer   VMNFS_SV sv_hourly 16@0,6,12,18`
	`create VM_VOL_SV  0@-`
	`xfer   VM_VOL_SV sv_hourly 16@0,6,12,18`
	
	NOTE: On the secondary system, “-x” was specified, which tells SnapVault to get the data from the primary system.

Expected Results	Baseline transfer will occur, transferring a full backup of the entire datastore. Once that is complete a base Snapshot copy will be created that will be used for future updates. After the baseline is completed, SnapVault will perform BLI updates only.
Actual Results	The baseline transfer occurred and the schedule was created. From this point on, we will only transfer the 4K block updates.

4.2 RECOVERY OF AN ENTIRE LUN DATASTORE

In the event an entire datastore on a primary system becomes corrupt, it’s important to be able to restore that datastore in an acceptable amount of time. The following shows the procedure for restoring an entire datastore in a LUN environment. For this test, the volume on the primary system was destroyed, erasing all the data. Prior to running through the procedure below, we had to enable auto volume resignaturing. You can do this from the “Advanced Setting” option under the “Configuration” tab for the ESX Server.

Task	Use the `snapvault restore` command to restore an entire datastore that resides on a LUN. 1. Verify that the virtual machines are no longer accessible.

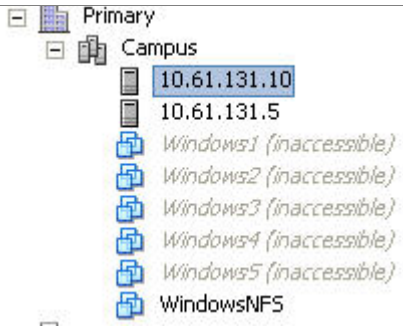

2. To restore the LUN, you must restore it to a new qtree on the primary storage system. Use the `snapvault restore` command from the primary system. Once the restore is executed, you must wait until the restore is complete before receiving the command prompt on the primary system.

```
ESX-PRIMARY> snapvsnapvault restore -S 10.61.131.101:/vol/VM_VOL_SV/dest /vol/VM_VOL/sv_restore
snapvsnapvault not found.  Type '?' for a list of commands
ESX-PRIMARY> snapvault restore -S 10.61.131.101:/vol/VM_VOL_SV/dest /vol/VM_VOL/sv_restore
Restore from 10.61.131.101:/vol/VM_VOL_SV/dest to /vol/VM_VOL/sv_restore started.
Monitor progress with the 'snapvault status' command.
Abort the restore with ^C.
```

NOTE: The qtree is created as part of the restore process so it cannot exist prior to the restore operation. If the same qtree name is required, it must be deleted prior to the restore.

3. Once the restore is completed the LUN needs to be mapped

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/VM_TMP/vmtmplun	VM Temp data	10 GB	online	vmware-prod : 1
/vol/VM_VOL/sv_restore/vm_lun	Main Datastore for VMware	100 GB	online	No Maps

To map the LUN, click on “No Maps.”

Initiator Group	LUN ID	Unmap
vmware-prod	0	☐

Apply

4. Once the LUN is mapped, you will need to rescan the adapter within the VI Client.

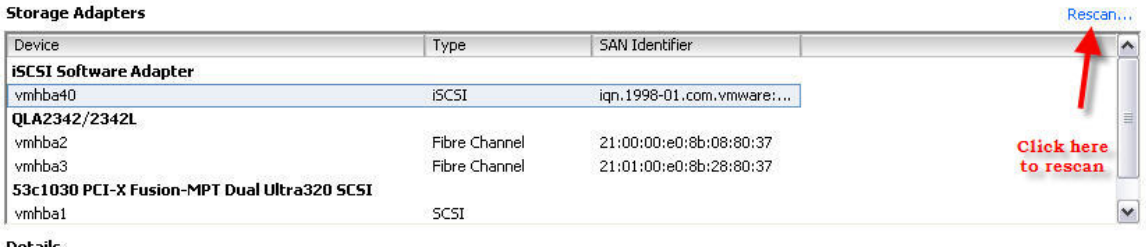

Once the rescan is complete, the LUN will appear in the VI Client.

Path	Canonical Path	Capacity	LUN ID
vmhba40:1:0	vmhba40:1:0	100.00 GB	0
vmhba40:1:1	vmhba40:1:1	10.00 GB	1

5. The datastore will now appear, but will have a “snap” prefix associated with it.

6. Now you can add the virtual machines back to the inventory by selecting the ESX Server, clicking the Configuration tab, and selecting “Storage.” Once you see the datastores, double click the datastore that was just restored to browse the contents. Browse the datastore into the virtual machine name (in this case “Windows1”) and locate the “.vmx” file. Right click on the file and select “Add to Inventory.”

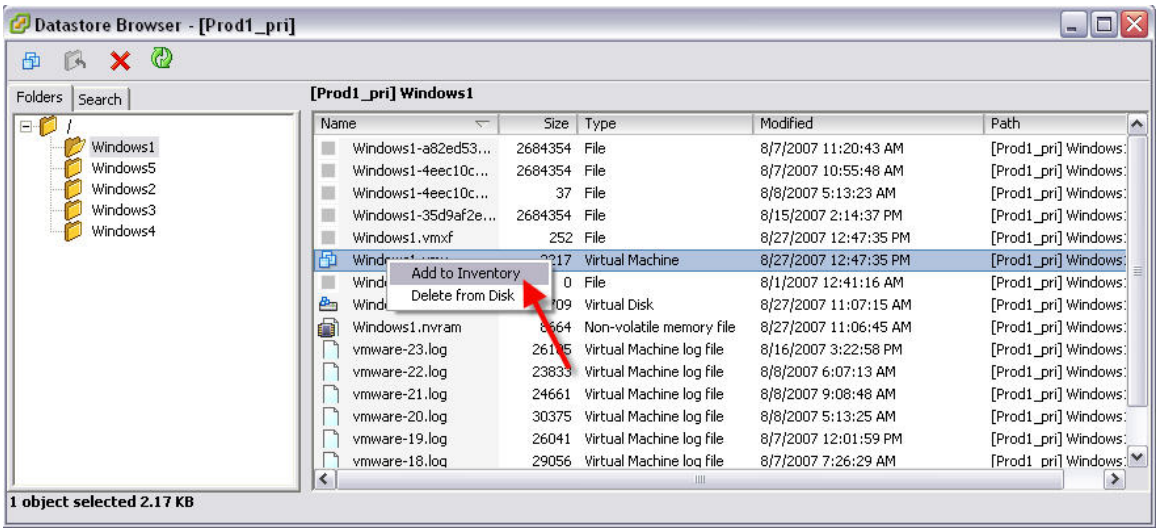

Repeat this process for each virtual machine.

7. Once you have added all the virtual machines to inventory, you can power them on.

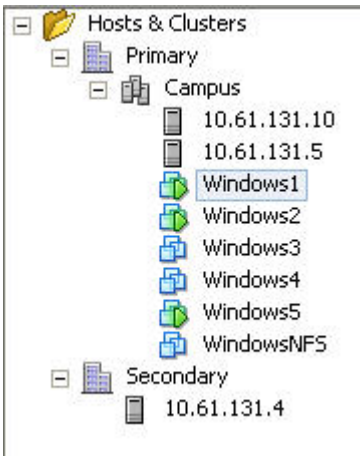

Expected Results	SnapVault will restore the entire LUN to a new qtree and the data will be available for VI to import.
Actual Results	Data was restored and the virtual machine could be powered on.

4.3 RECOVERY OF AN ENTIRE NFS DATASTORE

In the event an entire datastore on a primary system becomes corrupt, it's important to be able to restore that datastore in an acceptable amount of time. The following shows the procedure for restoring an entire datastore in an NFS environment. In this scenario, the entire primary volume was destroyed and is being restored from the secondary system.

Task

1. Verify that the virtual machines are no longer accessible.

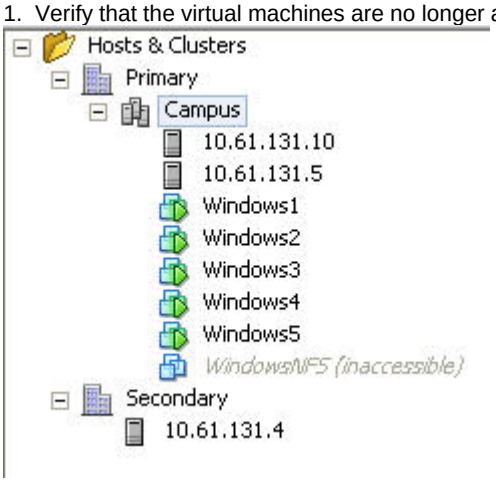

2. Restore the SnapVault destination qtree on a new qtree on the primary system.

```
ESX-PRIMARY> snapvault restore -S 10.61.131.101:/vol/VMNFS_SV/sv_dest /vol/VMNFS/sv_source
Restore from 10.61.131.101:/vol/VMNFS_SV/sv_dest to /vol/VMNFS/sv_source started.
Monitor progress with the 'snapvault status' command.
Abort the restore with ^C.
Made qtree /vol/VMNFS/sv_source writable.
Restore to /vol/VMNFS/sv_source completed successfully.
```

3. Once the restore is complete, you must export the file system again. You can do this either through FilerView® or by modifying /etc/exports on the NetApp storage system. After modifying /etc/exports, be sure to re-export the file systems again.

```
ESX-PRIMARY> exportfs
/vol/VM_VOL      -sec=sys,rw,nosuid
/vol/VM_TMP      -sec=sys,rw,nosuid
/vol/vol0/home   -sec=sys,rw,nosuid
/vol/vol0        -sec=sys,rw,anon=0,nosuid
/vol/VMNFS       -sec=sys,rw,anon=0
```

4. After you export the file systems, log back into the VI client and rediscover the NAS file system.

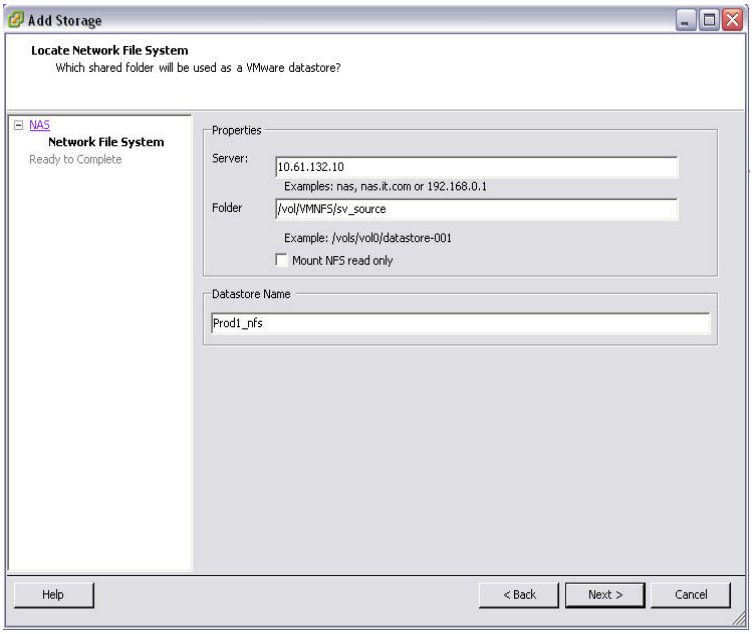

5. Now verify that the NFS datastore is available through the VI Client.

Identification	Device	Capacity	Free	Type
storage1	vmhba0:0:0:3	26.25 GB	25.64 GB	vmfs3
Prod1_pri	vmhba40:0:0:1	99.75 GB	69.13 GB	vmfs3
Prod1_tmp	vmhba40:0:1:1	9.75 GB	2.13 GB	vmfs3
Prod1_nfs	10.61.132.10:/vol/VMNFS/sv_source	80.00 GB	74.98 GB	nfs

6. The virtual machine can now be added back to the inventory and powered on.

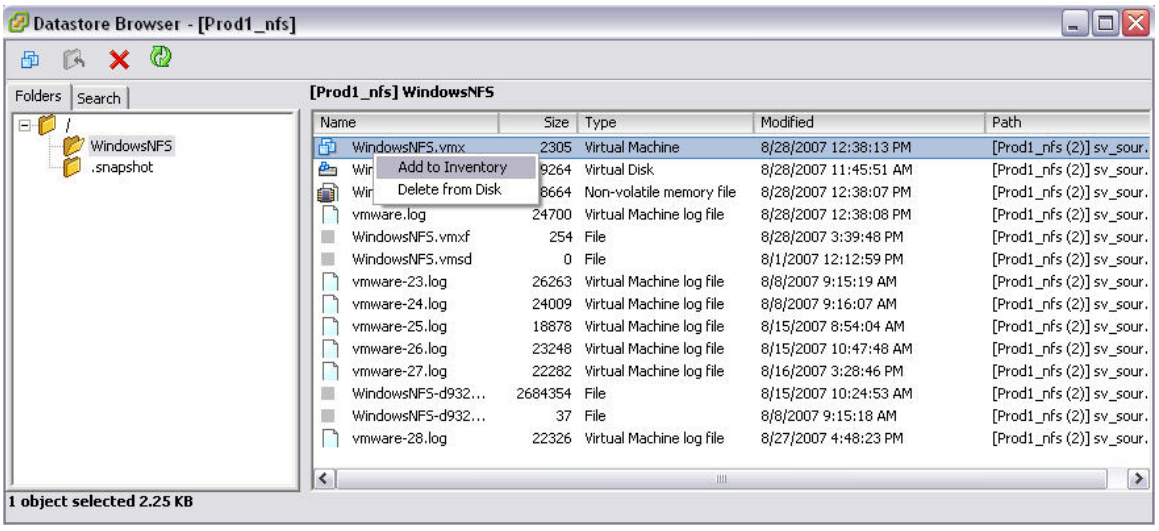

Expected Results	It is expected that the virtual machine will be restored from the secondary system and then successfully powered on.
Actual Results	Once we restored the NFS datastore the virtual machine could be powered on.

4.4 RECOVERY OF A SINGLE FILE (WITHIN A VIRTUAL MACHINE)

One of the more difficult tasks in replicating entire virtual machines is recovering a single file. One method for recovering a file that resides within the vmdk for that virtual machine is taking a FlexClone copy from the Snapshot in which the data resides and mounting that vmdk to the virtual machine.

Task

1. First we delete some data so we have something to recover.

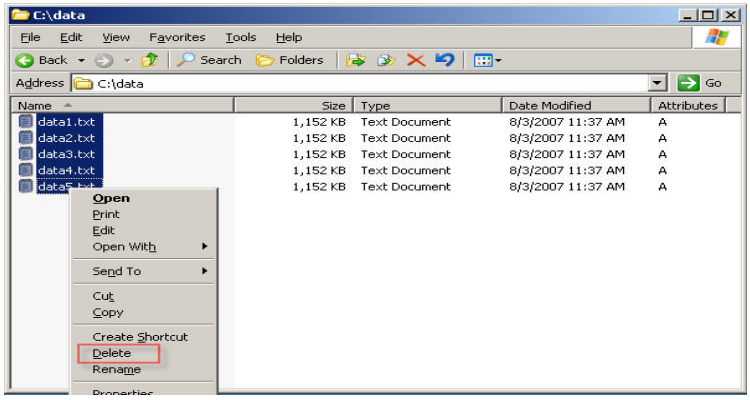

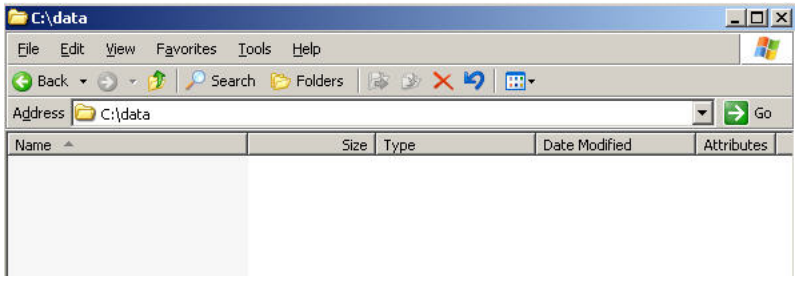

2. Once the data is deleted, we create a FlexClone copy of the last Snapshot copy (in this case sv_hourly.0) using the `vol clone create` command.

```
ESX-PRIMARY> vol clone create NFS_RECOVERY -b VMNFS sv_hourly.0
Wed Aug 29 19:30:46 GMT [ESX-PRIMARY: wafl.snaprestore.revert:notice]: Reverting volume NFS_RECOVERY to a previous snapshot.
Creation of clone volume 'NFS_RECOVERY' has completed.
```

3. When the clone is created, it should already have been exported with the same settings as the export of the parent volume. If it was not, you need to export the clone via NFS.

4. Once the clone is created and exported, mount the cloned file system inside the VI Client. This can be done by clicking on “Add Storage” from the storage configuration in the VI Client.

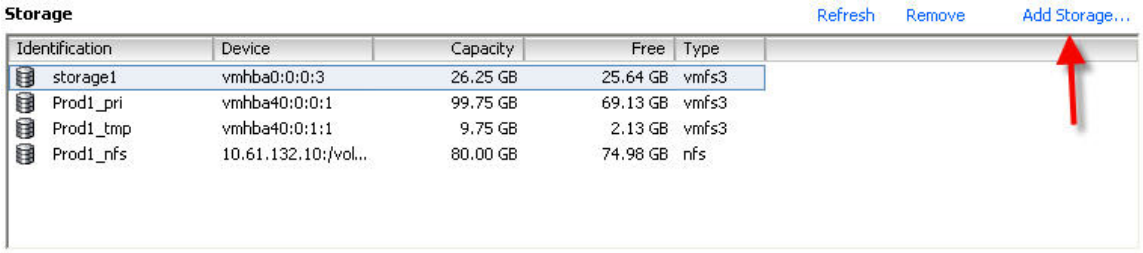

5. Be sure to give the NFS datastore a unique name since it will be removed once the recovery is complete.

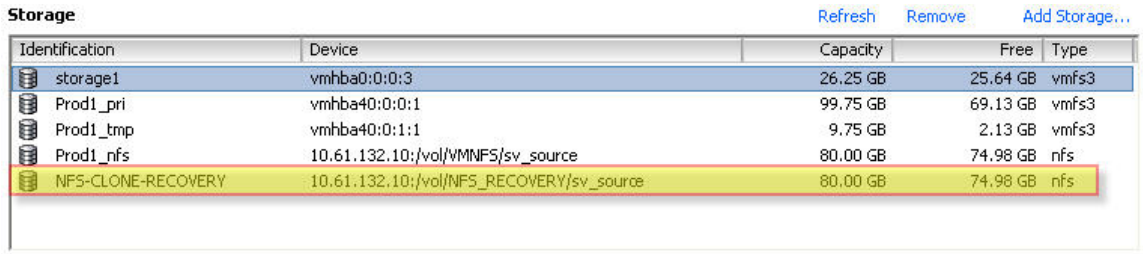

6. You can now add the vmdk from the Snapshot copy to the Windows virtual machine. To do this, select the virtual machine, select “Edit Settings,” and choose to add an existing hard drive.

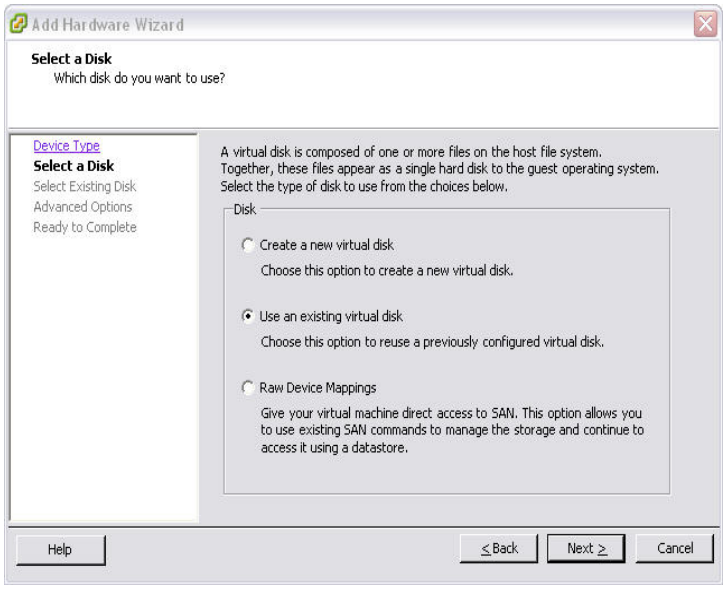

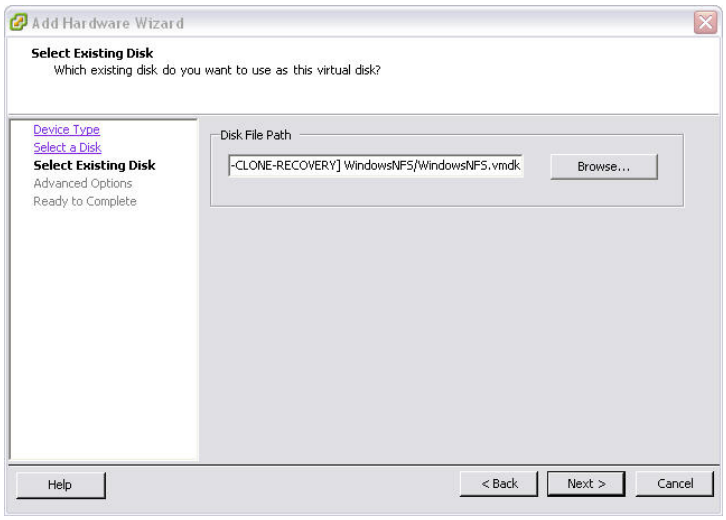

7. After you add the drive, log onto the virtual machine and re-scan to pick up the newly added drive.

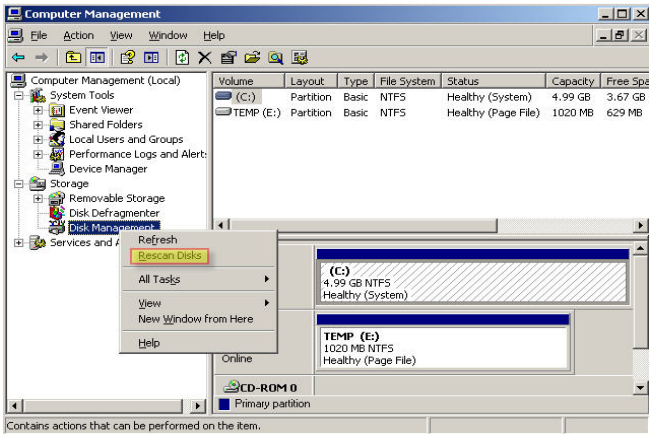

	8. After the drive is discovered, mount the drive (in this case we mounted using the drive letter “R”) and browse through the drive to find the data that was deleted. 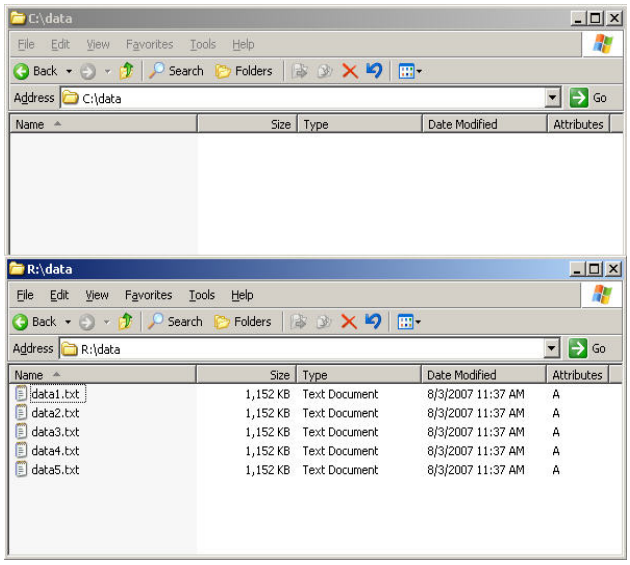 9. The data can now be restored to the correct local drive or copied to another system. 10. Once the data has been restored, be sure to remove the vmdk (this requires that you power down the virtual machine), remove the cloned datastore from the VI Client, and destroy the clone.
Expected Results	It is expected that a FlexClone copy can be created from an existing Snapshot copy. Once the FlexClone is created and the datastore discovered, a single file should be recoverable by mounting the vmdk to a running virtual machine.
Actual Results	A single file was restored with no service interruption to the running virtual machine. The virtual machine needs to be powered on in order to remove the vmdk used for the restore process.

4.5 USING BACKUP FOR TEST/DEV (LUN)

One of the many benefits of NetApp Snapshot technology is the ability to use the FlexClone technology to create a writable version of a Snapshot copy. Pairing FlexClone capabilities with SnapVault provides the ability to verify and test against any backup without compromising the data. The following procedure focuses on using FlexClone copies in a LUN environment.

Task

Since the SnapVault Snapshot copies are read-only on the destination system, we need to create a FlexClone copy to use these backups in a test/dev scenario.

1. First create a clone of the volume that contains the LUN with the datastore. Note that the cloned LUN is offline and the existing LUN is "r/o."

```
ESX-SECONDARY> vol clone create VM_VOL_CLONE -b VM_VOL_SV sv_hourly.1
Wed Aug 29 20:52:18 GMT [ESX-SECONDARY: wafl.snaprestore.revert:notice]: Reverting volume VM_VOL_CLONE to a previous snapsho
t.
Wed Aug 29 20:52:18 GMT [ESX-SECONDARY: wafl.qtree.qsmBreak.base:error]: Breaking snapmirrored qtree 1 in volume VM_VOL_CLON
E: base snapshot no longer exists. Use snapmirror resync or initialize to re-establish the snapmirror.
Creation of clone volume 'VM_VOL_CLONE' has completed.
ESX-SECONDARY> Wed Aug 29 20:52:19 GMT [ESX-SECONDARY: lun.newLocation.offline:warning]: LUN /vol/VM_VOL_CLONE/dest/vm_lun h
as been taken offline to prevent map conflicts after a copy or move operation.

ESX-SECONDARY> lun show
        /vol/VM_TMP/vmtmplun          10g (10737418240)    (r/w, online, mapped)
        /vol/VM_VOL/sv_source/vm_lun    100g (107374182400)   (r/o, online)
        /vol/VM_VOL_CLONE/dest/vm_lun    100g (107374182400)   (r/w, offline)
        /vol/VM_VOL_SV/sv_dest/vm_lun    100g (107374182400)   (r/o, online)
        /vol/vol1/.jimlun              5g (5368709120)     (r/w, online, mapped)
```

2. Since the LUN is "offline," you must first bring it online.

```
ESX-SECONDARY> lun online /vol/VM_VOL_CLONE/dest/vm_lun
ESX-SECONDARY>
```

3. After you bring the clone online, map the LUN and add it to the Initiator group. This can be done using FilerView.

Add New LUN Hide Maps

LUN	Description	Size	Status	Maps Group : LUN ID
/vol/VM_TMP/vmtmplun		10 GB	online	dr : 1
/vol/VM_VOL/sv_source/vm_lun	Main Datastore for VMware	100 GB	online	No Maps
/vol/VM_VOL_CLONE/dest/vm_lun	Main Datastore for VMware	100 GB	online	No Maps
/vol/VM_VOL_SV/sv_dest/vm_lun	Main Datastore for VMware	100 GB	online	No Maps
/vol/vol1/.jimlun	test lun	5 GB	online	dr : 2

4. Once the LUN is mapped to the Initiator group for the ESX Server, log into the VI Client to discover the LUN.

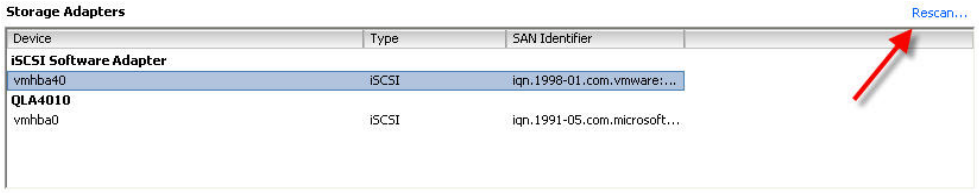

After the scan (the LUN was mapped with an address of 10).

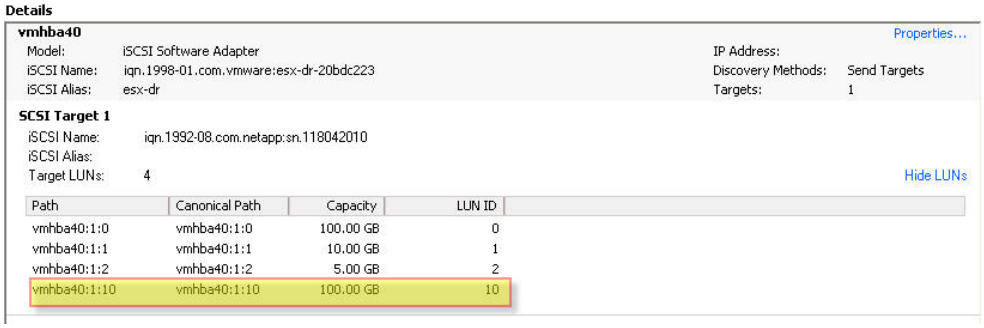

5. After the LUN has been discovered by the VI Client, verify that the VMFS datastore is now visible.

NOTE: The datastore will come in with a “snap” prefix in the name. This can be changed, but since it’s for test/dev, it may not be necessary.

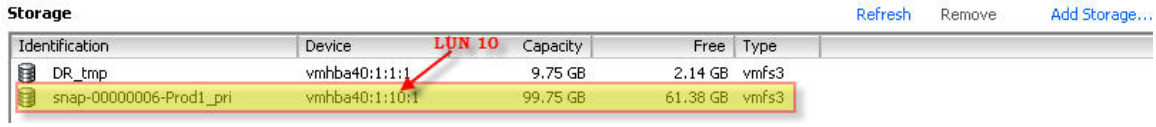

6. In order to import the virtual machines, you must add them to the inventory by browsing the datastore. You can do this by right-clicking on the datastore and selecting “Browse datastore.” Once the contents of the datastore are visible, drill down to the virtual machine you wish to test and select the vmx file to add it to the inventory.

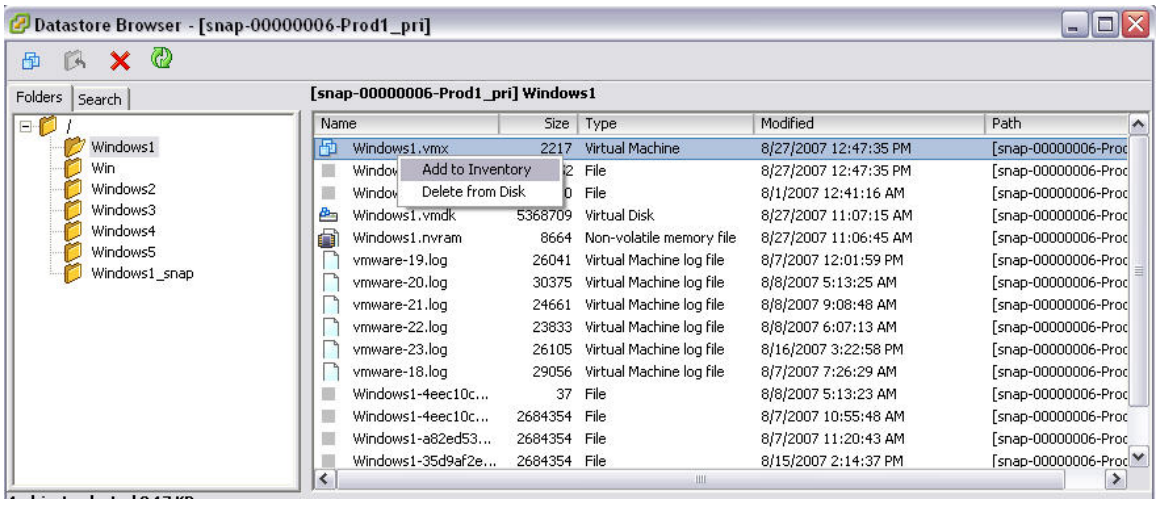

7. When the “Add to Inventory” window appears, be sure to select the secondary data center.

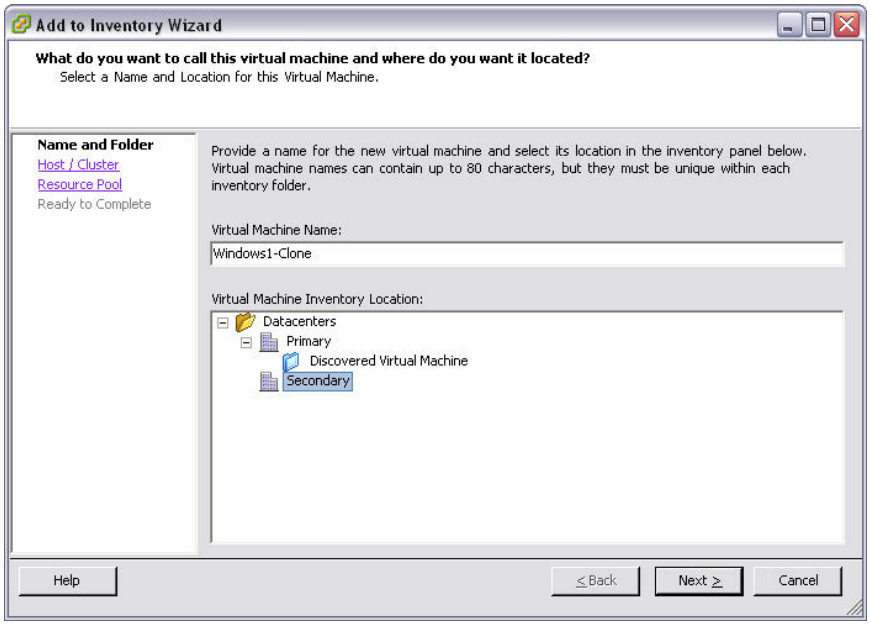

8. Disk2 will show an invalid disk name for the TEMP datastore.

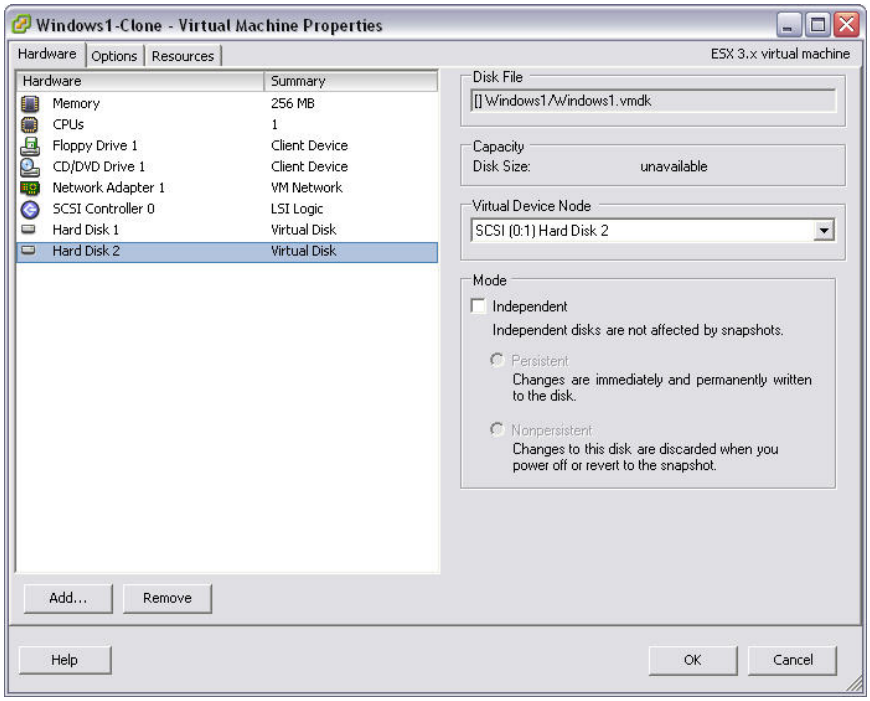

9. Remove Hard Disk 2.

	10. Add a new Hard Disk 2, pointing to the correct TEMP location (be sure to select the option “Use Existing Virtual Disk”). Repeat for each virtual machine. 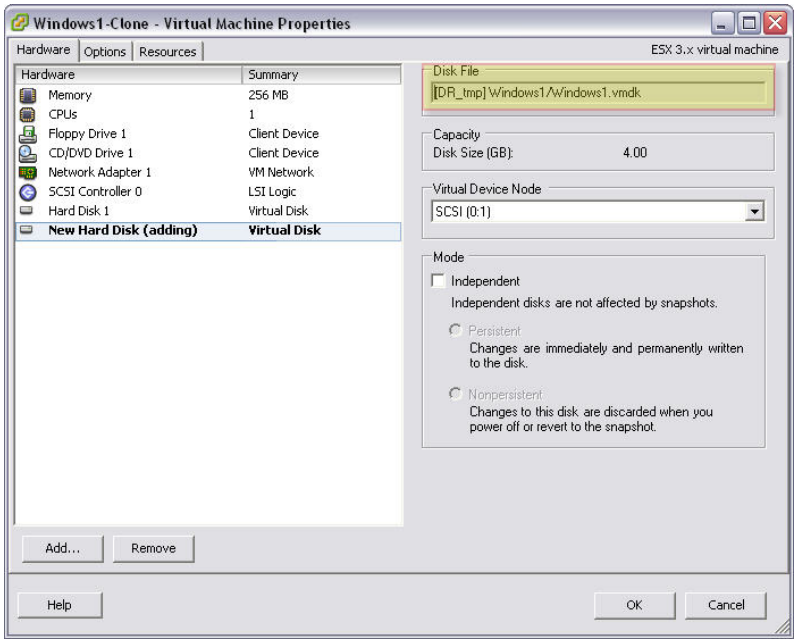 11. You can now power on the virtual machines. 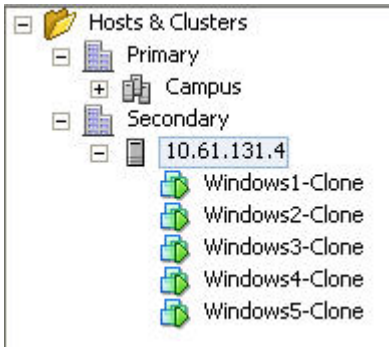
Expected Results	A FlexClone copy will be created and the virtual machines can be powered on.
Actual Results	The FlexClone copy was created and all virtual machines could be powered on.

4.6 USING BACKUP FOR TEST/DEV (NFS)

One of the many benefits of NetApp Snapshot technology is the ability to use the FlexClone technology to create a writable version of a Snapshot copy. Pairing FlexClone capabilities with SnapVault enables the ability to test any backup without compromising the data. The following procedure focuses on using FlexClone copies in a LUN environment.

Task

1. Create a clone of the volume with the NFS datastore. It's important to note that you can't use sv_hourly.0 (the most recent Snapshot copy) because it will retain the r/o attribute of the SnapVault destination.

```
ESX-SECONDARY> vol clone create VMNFS_SV_CLONE -b VMNFS_SV sv_hourly.1
Thu Aug 30 17:48:54 GMT [ESX-SECONDARY: wafl.snaprestore.revert:notice]: Reverting volume VMNFS_SV_CLONE to a previous snapshot.
Thu Aug 30 17:48:54 GMT [ESX-SECONDARY: wafl.qtree.qsmBreak.base:error]: Breaking snapmirrored qtree 2 in volume VMNFS_SV_CLONE:
 base snapshot no longer exists. Use snapmirror resync or initialize to re-establish the snapmirror.
Creation of clone volume 'VMNFS_SV_CLONE' has completed.
```

2. Verify that the export has the correct permissions.

```
ESX-SECONDARY> exportfs
/vol/VMNFS_SV_CLONE      -sec=sys,rw,anon=0
/vol/VM_VOL      -sec=sys,rw,nosuid
/vol/VM_VOL_CLONE      -sec=sys,rw,nosuid
/vol/vol1_dest   -sec=sys,rw,nosuid
/vol/VM_TMP      -sec=sys,rw,nosuid
/vol/vol0/home   -sec=sys,rw,nosuid
/vol/VMNFS_SV    -sec=none,rw,anon=0
```

3. Add the NFS datastore from the Conguration tab for the ESX Server.

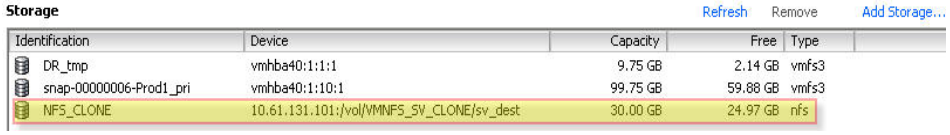

4. Browse the datastore and add the virtual machine to the inventory by right-clicking on the vmx file and selecting "Add to Inventory."

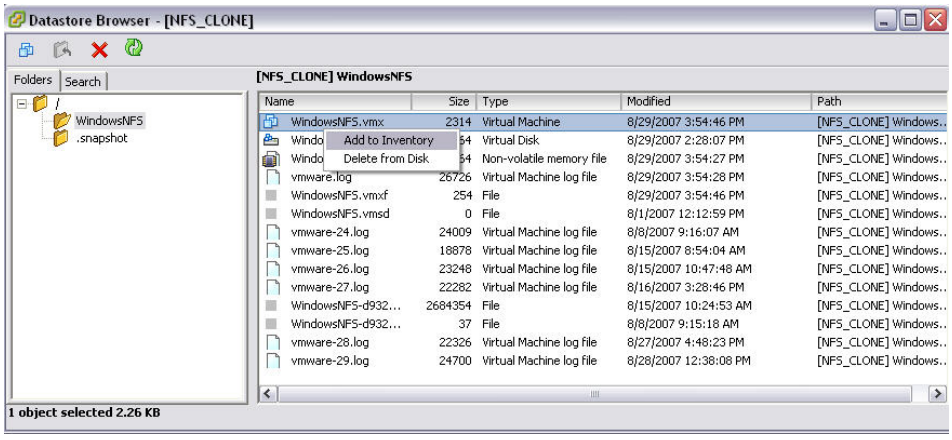

	5. Remove the old Hard Disk 2 since it points to a false location. 6. Add a new Hard Disk 2, pointing to the correct TEMP location (be sure to select the option “Use Existing Virtual Disk”). Repeat for each virtual machine. 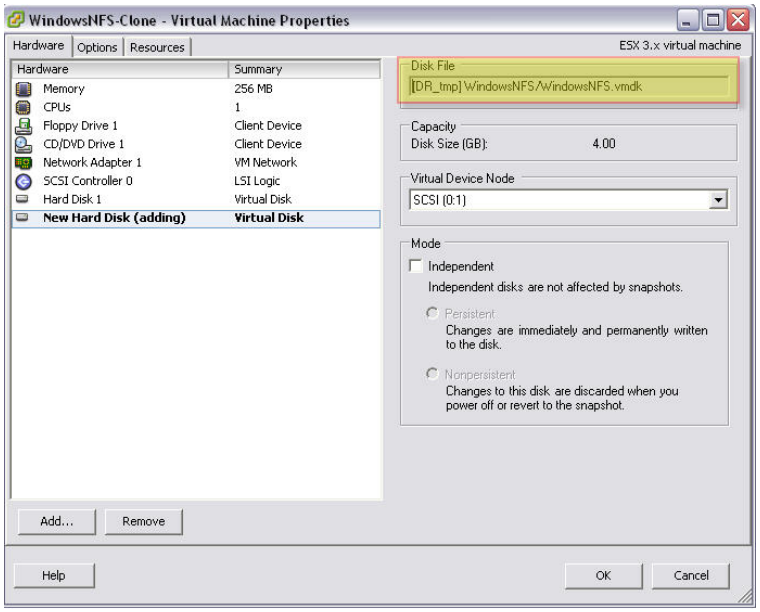 7. You can now power on the cloned virtual machine. 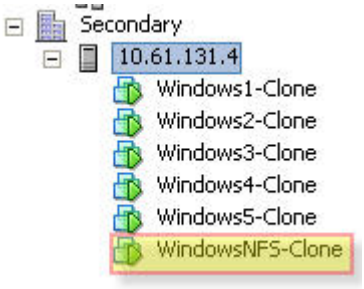
Expected Results	A FlexClone copy will be created and the virtual machine can be power on.
Actual Results	The FlexClone copy was created and all virtual machines could be powered on.

5 CONCLUSION

After running through the various operational scenarios in a VMware and NetApp environment, you can see that SnapVault is a very effective tool in backing up a virtual environment. Using SnapVault as the disk-to-disk solution provides fast, efficient, and reliable backups. In addition, SnapVault provides space savings by performing a single level 0 backup and then sending only the 4KB blocks required for the incremental backups. SnapVault brings robust restore capabilities to the virtual environment by allowing users to create read-writable copies of their data for backup verification and in test/development situations.

6 REFERENCE MATERIAL

NetApp

NetApp and VMware ESX Server 3.0

NetApp and VMware ESX Server 3.0 - Storage Best Practices

SnapVault Best Practices Guide

Open Systems SnapVault Best Practices Guide

VMware

VI3.0.1 Installation and Upgrade Guide

VI3.0.1 Resource Management Guide

VI3.0.1 Server Configuration Guide

APPENDIX A MATERIALS LIST

Table A-1) Materials list.

Hardware				
Storage	Vendor	Name	Version	Description
	NetApp	FAS3050C		SnapVault Primary System
	NetApp	FAS6070		SnapVault Secondary System
Hosts	IBM	IBMX335		
	IBM	IBMX306		
	IBM	IBMX335		
Software				
Storage	NetApp	Data ONTAP	7.2.3	Operating System
	NetApp	SnapVault	7.2.3	Data Protection
	NetApp	FlexClone	7.2.3	Test/Dev
Hosts				
	VMware	Infrastructure 3	3.0.1	
	Microsoft	Windows Server 2003 Enterprise Edition-SP1 (x86)	2003	Operating System
Host Utilities	NetApp	VMware Host Utilities (WHU)	3.0	

APPENDIX B PLATFORM SPECIFICATION

B.1 FAS STORAGE CONTROLLER

CONFIGURATION

The FAS Storage System in the ESX-Primary

STORAGE CONTROLLERS

Name	Description	IP Address
ESX-Primary	FAS3050	10.61.132.10
ESX-Secondary	FAS6070	10.61.131.101

AGGREGATE LAYOUT

Table 5-2) Aggregate layout.

Controller	Aggregate Name	Options	# Disks	Purpose
ESX-Primary	aggr0	RAID_DP, aggr mirrored	3	Root volume
ESX-Primary	aggr1	RAID_DP, aggr mirrored	10	Datastores
ESX-Secondary	Aggr0	RAID_DP,	3	Root
ESX-Secondary	Aggr1	RAID_DP	39	Backup datastore destination

VOLUME LAYOUT

Controller	Volume Name	Qtree	Options	Total Volume Size	Purpose
ESX-Primary	Vol0		RAID_DP, Flex	191GB	Root volume
ESX-Primary	VM_VOL	sv_source	RAID_DP, Flex	200GB	VMDK datastore
ESX-Primary	VMNFS	sv_source	RAID_DP, Flex	100GB	NFS datastore
ESX-Primary	VMTMP		RAID_DP, Flex	20G	VMDA temporary (page file, etc.)
ESX-Secondary	Vol0		RAID_DP, Flex	268GB	Root volume
ESX-Secondary	VM_VOL_SV	sv_dest	RAID_DP, Flex	200GB	VMDK datastore
ESX-Secondary	VMNFS_SV	sv_dest	RAID_DP, Flex	100GB	NFS datastore
ESX-Secondary	VMTMP		RAID_DP, Flex	20G	VMDA temporary (page file, etc.)

B.2 HOST SERVERS

Software Configuration

The hosts in the cluster are installed according to the procedure documented in the *VMware ESX 3.0.1 Installation and Upgrade Guide* with:

- Windows 2003 Enterprise Edition
- VMware Infrastructure 3

Network Settings

The following tables provide the network settings for the ESX Servers:

Hostname	IP Address
ESX-PROD1	10.61.131.5
ESX-PROD2	10.61.131.10
ESX	10.61.131.4

iSCSI / LUN Setup

Several qtrees have been created for the LUN files. These LUNs have the following attributes:

Purpose	Drive	LUN Size	LUN File
VMDK	C:	200G	/vol/VM_VOL/sv_source/vm_lun
VMDK	D:	20G	/vol/VM_TMP/vmtmplun